

Is Computing an Experimental Science?

Robin Milner, *Laboratory for Foundations of Computer Science, Edinburgh University*

Is computing an experimental science?⁽¹⁾

At the Laboratory for Foundations of Computer Science at Edinburgh we are beginning an ambitious programme of research. The particular programme which we have put forward is a new kind of exercise. What makes it new is a central commitment to a double thesis: that the design of computing systems can only properly succeed if it is well grounded in theory, and that the important concepts in a theory can only emerge through protracted exposure to application.

If we take this commitment seriously, then we must make an organized attempt to unite the development of theory with its experimental application. That is why I chose a title for this paper which, in its truly philosophical sense, I am not well equipped to answer. All the same, I think people involved with a developing subject ought to ask themselves difficult questions about what it is they are developing, even if the final answers are beyond them. Personally, I became committed to the inseparability of theory from experiment in the course of doing computer science, not through philosophical enquiry. In this paper I want to justify the commitment by looking at some examples of how ideas and applications have developed side-by-side. After all, however deep a philosophical study you wish to conduct, you must always support it by examples of scientific practice. After having read this paper, I hope you will at least agree that the question 'Is computing an experimental science?' is both fascinating and intricate. Whatever the answer, I hope even more that you will be convinced that our commitment to unite theory with application and the research, which we plan on the basis of this commitment, are both inevitable in the process of establishing firm foundations for computing as a science.

The nature of computation theory

Computation has a hard theoretical core which few people would regard as experimental. This core is closely linked with logic, and is the foundation

stone of theoretical computer science; it concerns what is computable, which is strongly connected to what is deducible. Around 1900, part of Hilbert's programme was to show that every mathematical truth would turn out to be a deducible from a set of axioms, and it was vital to find those axioms. Of course this theory was exploded by the work of Godel, Kleene, Church and Turing; it was found that any axiomatic system which is strong enough to entail enough of the things we want to be true will also, inescapably, be either inconsistent or too poor to entail other things which we regard as true. There will always be truths which are not deducible, and functions which are not computable. Out of this came several characterizations of what is computable, and they are agreed; Church's thesis states that they will always agree. (2)

The hard core of computing theory ramified, with this foundation, into a classification of what is computable. This ramification began with recursion theory, showing that certain things are essentially harder to compute than others. Modern complexity theory carried on this programme; we now know that a lot of useful things can be computed in linear time (time which is a linear function of the size of the problem), others take non-linear but polynomial time, and unfortunately still other — undeniably useful things! — almost certainly take exponential time and a few of them definitely do.

For some computer scientists, this difficult and unfinished study is computation theory. To them, everything else is engineering of one kind or another; either hardware engineering (what configurations of matter do our computations reasonably fast) or software engineering (what configurations of text express our computations reasonably clearly).

I want to argue that this is a restrictive view. Of course, any engineering discipline employs mathematical theories in the service of design; theories which help in doing good design and in validating artifacts, i.e. establishing that they do

the job they were designed to do. Computing is no exception in this, although agreement about which mathematical theories are most appropriate to inform design is as yet limited. It might be argued however, that these helpful mathematical theories are auxiliary, in the sense that they fail, unlike the hard core theory, to tell us anything about the nature of computation; all they do is to help us in building complex computing systems and in managing their complexity.

This is the position which I want to undermine. I intend to show that these apparently auxiliary theories, for managing the things which we build, are just as essential as the hard core; that is they tell us just as many things, though different things, about the very nature of computation. Not only do they help us to organize big systems more effectively; they also provide structural concepts for us which are indispensable when thinking about our big systems. Finally, I want to show that we arrive at these concepts, and assess their claim to centrality (or their inevitability) not just by thinking about big systems but by building and working them.

Engaging theory with experiment

Britain has always played a prominent part in fundamental computing research (Turing's role has already been mentioned). At Cambridge and Manchester the first stored-program computers (inspired probably by Von Neumann's ideas) were developed by Wilkes and Kilburn, around 1950 (3). At Oxford in the late 1960s, Christopher Strachey and the American logician, Dana Scott, inspired one another in beginning the mathematical study of computer program semantics (4). This fusion of engineering and logical insights generated an energy which has fired the qualitative, rather than quantitative, study of computation ever since. For some reason the US has, in the last two decades, been largely preoccupied with the quantitative aspect of computation theory, called complexity theory, while in Britain the emphasis has been more upon semantics. In general, this semantic study has been most strongly represented here in Edinburgh, notably in the work of Plotkin (5); of course we acknowledge pioneering work such as that of Hoare in proposing special logics for programming (6), and of Kowalski in uniting logic with programming (7).

In parenthesis, it is interesting to remark that AI — with its demand for a kind of programming vastly different from FORTRAN — has provided some of the impetus towards a study of semantics. Perhaps this was because a large part of AI's subject matter was always, and is now, automated inference. Perhaps it was also because the languages inspired by AI — LISP at MIT and POP2 at Edinburgh —

were themselves challenging objects of semantic study, since they were informed by the lambda-calculus which Church had established at the foundation of computation theory.

Whatever its origins, the work here in Edinburgh on the foundations of computing theory has flourished over the past decade. One of its strongest contributions was to complexity theory, in the work of Valiant and his followers (8). Otherwise it has concentrated on semantics: on program transformations which preserve meaning; on the model theory of programming and in particular for non-determinism; on semantically-directed design of functional programming languages; on the methodology of computer-assisted reasoning, both in general and with particular specifications; and on the conceptual basis of parallel computation.

In effect, then we already have a Foundations Laboratory. Why are we gilding the lily, by bothering to establish one? Why should we give the impression that we want to establish something, when in fact we are already doing it? I began to answer this question by claiming that one arrives at concepts (at any rate in this subject), not only by thinking about systems but by building them. In fact, the time is right to consider this attitude really seriously, because suddenly almost everybody who builds computing systems is convinced that all system design — hardware or software — needs to be done within a rich conceptual frame, which is articulated by rigorous methodologies, which are in turn assisted by computer-based reasoning tools. One key point here is that a conceptual frame only provided through a programming language is too narrow; a wider frame is needed in which to understand the specifications of the systems, and the methodology which articulates this understanding must be based on some form of logic.

It is excellent that this conviction is now widespread; but there is an attendant danger. Much use has been made recently of the term 'formal methods' in relation to system design. The choice of this term (wherever it came from) rather than the term 'theory' suggests that the methodology is paramount; I fear that it also reflects a mistaken assumption — that is, it suggests that the conceptual frame for design already exists and that we only need to animate it by the right formal methodology.

One need only recall that most of mathematics has arisen from some kind of application demand, often from the physical sciences, to see how unlikely this assumption is to be valid. Let us assume that it is invalid and that the conceptual frame for system design is still incomplete. How do we develop it? Well, much as a physicist or

chemist tests and refines his theories by carefully controlled experiments, so it should be with us. I believe this analogy is close and that the word 'experiment' is also correct for computer science. However, for a physical scientist an experiment will reinforce (or undermine) his conceptual grasp of what is true; for the computer scientist, the experiment will reinforce (or undermine) his conceptual grasp of how to design. It takes a philosopher of science to decide whether a 'theory of what' is really different from a 'theory of how', although later I boldly contest the existence of such a distinction.

For now let us return to the question of experiment. What is an experiment for us? I claim that it employs a prototype methodology constructed upon some conceptual basis, and that its outcome determines whether we adopt (or reject) the methodology, thus reinforcing (or undermining) the conceptual basis – the 'theory of how'. Notice that it won't usually be all or nothing. But at least one philosopher of science, Imre Lakatos, argues that no experiment ever completely confirms or denies a 'theory of what', either.

In Edinburgh we have contributed the conceptual basis for some such experiments and are beginning to assess the outcome. I would like to explain two experiments that I have been involved with.

1. In the mid 1970s we designed LCF (Logic for Computable Functions), a machine-assisted reasoning system (8) based on Dana Scott's model of computation, together with certain notions of proof tactic or strategy with which a user could conduct his proof. At Edinburgh and then at Cambridge, Michael Gordon (9) specialized this tool towards verifying the design of a computer; he called his tool LCF-LSM. Recently at RSRE, Malvern, Dr John Cullyer has succeeded, using Gordon's work, in proving real microcomputers correct (or, more often, faulty!) and he strongly advocates that no micro should be installed in any situation where humans are at risk unless such a proof has been done. This outcome reinforces at least part of our 'theory of how'; other parts were not so much undermined as found unnecessary, but then this application was very specialized.
2. In 1980 we put forward a Calculus of Communicating Systems (10), an algebraic system for reasoning about concurrent computation informed by a particular concept of process. In the last two years, two Dutch scientists, Chris Vissers and Ed Brinksma, have designed a methodology based upon our work, which they call LOTOS (11) (Language of Temporally Ordered Specifications), for describing 'Communications Protocols'; a protocol is a discipline to be adopted by digital systems which need to exchange messages reliably,

even in the presence of faulty media. The design has also been strongly influenced by the seminal work of Tony Hoare on communicating sequential processes. This language LOTOS is well-advanced as a candidate for an International Standard, and it has been found rather convenient for its purpose. At the same time, it identified the need for at least one construction which was not derivable from our theory. It has also served as the arena for a satisfyingly heated debate about the central concept of process! This has been criticized as too coarse in some respects, too fine in others, and the outcome is not yet clear. But one thing is clear: it is hard to imagine a better way of testing the process concept as a 'theory of how'.

These two examples should indicate that theoretical innovation is necessary in computer science, that it can be tested, and that we certainly need more of it. Perhaps it has not persuaded you that we should change the way we work as a theoretical group, by founding a Laboratory; after all, in these two examples the experiments were done (or are being done), and the outcome is there for us to assess.

A large part of the problem is the scale of it. It is the essence of computer science that if a methodology does not scale up, then it isn't a methodology at all. So a theory which does not support a scaling methodology is not the right theory. Tackling large experiments inevitably demands software tools. Building these tools alongside the theoretical development is the only way to make theoretical research really sensitive to potential use and this joint pragma-theoretic approach demands both more research effort and more co-ordination than is normally supported by an academic department. Another part of the problem is that the people best able to conduct experiments with these tools are usually not academics, but industrial researchers. To induce them to do so requires an environment in which industrial researchers can both attend short courses and (in smaller numbers) reside for longer periods, to discover what theories there are, and to help us see what theories are lacking. And if we are to get feedback from experiment in a controlled and intimate manner, then the environment must house both joint experimental projects with industry, and also be organized to respond to experiments conducted within the companies.

At the end of this paper, I shall outline how we propose to develop our Laboratory to make all this possible. I hope now to have persuaded you that some special initiative is required in order that theoretical research should be closely engaged with its possible application. I now want to

examine two new kinds of theory being developed in Edinburgh and elsewhere which promise to help in the design process.

New theories

Types and Specifications

One of the most far-reaching experiments which computer scientists are now conducting aims to provide the end user with a specification-calculus — that is, he is provided with a way of building large program systems not using programs as a conceptual unit, but instead using specifications — which describe *what* is to be done (not *how* to do it). The term 'type' here is synonymous with 'specification' — this reflects the insight that the concept of type familiar in logic can be aligned with what programmers like to call a specification. These types, and the properties which relate to them, are essential in his 'theory of how'; the experiment will succeed, and thereby reinforce this theory, just to the extent that the conceptual frame of types provides users with a vehicle for their design purposes.

A simple example illustrates how rich these types can be. Let us suppose that we are working with tables of various kinds. An employee table, giving the employment grade of each employee, would have type

(NAME, GRADE)TABLE

This is a simple data-type, but not only data may be specified by type; operations on data may also be specified by type. For example, a common operation in database systems is to join two tables. A salary scale may have type

(GRADE, SALARY)TABLE

Then the operation which will join tables of these two types into a single table can be specified by a type JOIN, which may be written

given input1 of type (NAME, GRADE)TABLE
 JOIN = and input2 of type (GRADE, SALARY)TABLE
 yields output of type (NAME, GRADE, SALARY)TABLE

Of course more succinct mathematical notation can be used here. Notice that the type JOIN does not say how to compute anything; it merely says how the output type is related to the input types. Also, the type JOIN is not very specific. There are many programs with this type which would fall short of our requirements, which are to produce an output table which exactly combines the information of the two input tables. For example, it might produce an output table with just one entry:

BLOGGS LECTURER £8,000

whatever the input tables were! A better specification can state what the output-input dependency should be, thus:

given input1 of type (NAME, GRADE)TABLE
 and input2 of type (GRADE, SALARY)TABLE
 PROPERJOIN = yield output of type (NAME, GRADE, SALARY)TABLE
 such that project(output)
 (NAME, GRADE) = input 1
 and project(output)(GRADE, SALARY) = input2

Notice that this is tighter; anything of this type will also have the type JOIN, which we may express as follows

PROPERJOIN implies JOIN

Nevertheless, the new type does not say *how* the joining operation should be done; it only gives conditions which it must satisfy.

It may appear from this example that a type is really a logical formula, and that the calculus of types is a logical system. This is indeed how it is often treated although it has a strong algebraic character too. This is clearer if we take the example just a single step further. Notice that joining tables is an operation which makes sense whatever the elements in the tables are; it does not work just for the element types NAME, GRADE and SALARY. In other words, it is correct and much more useful to treat JOIN as a parametric type, where the parameters are types themselves. This idea is much more general: it turns out that with some help from category theory, we can build a hierarchy of specifications, where higher members in the hierarchy take lower members as parameters. This is just what Rod Burstall at Edinburgh, and others elsewhere, have done (12). What emerges then, is a highly structured conceptual frame in which designers can think and reason about their designs before they express the fine detail in the form of programs. It is the structure among these type objects, not the notation (logical or otherwise), which informs the designers' understanding. It is exciting, but hardly surprising that the Swedish logician Per Martin-Lof has shown us that his Constructive Type Theory (13), which he developed purely as a theory of reasoning, gives us an excellent framework in which to handle types and specifications.

This experiment, dealing with specifications, is mainly about the Man-Machine Interface rather than the ergonomics of this interface (a subject which is of great importance and which was chosen as one of the foci of concentration of the

Alvey Programme (14)). It concentrates on the things a human would like to think about when he is controlling a machine. I would argue that however refined an ergonomic study is made in MMI, it cannot really succeed unless it is firmly based on a structured conceptual frame such as we have proposed for specifications.

The notion of process

Let us now turn to the other kind of interface – the Machine-Machine Interface. When one of these machines is a computer and the other is an aircraft, then the communication discipline involves a host of problems to do with real-time, estimation, feedback and fault tolerance which are both complex and highly specialized, and which I do not want to discuss here. What I do want to discuss is the symmetry of the interface, which is often obscured. When an aircraft or chemical plant is controlled by a computer, then the interaction is definitely two-way; it helps to think of the computer as a process, and each of the two processes evokes action from the other. So it leads to an integrated theory if we can model each agent in such a pair by the same conceptual notion of process.

There is more to it than this, because we are not restricted to two parties interacting. Consider three power plants, each consisting of a computer controlling a generator and with the computers linked to each other.

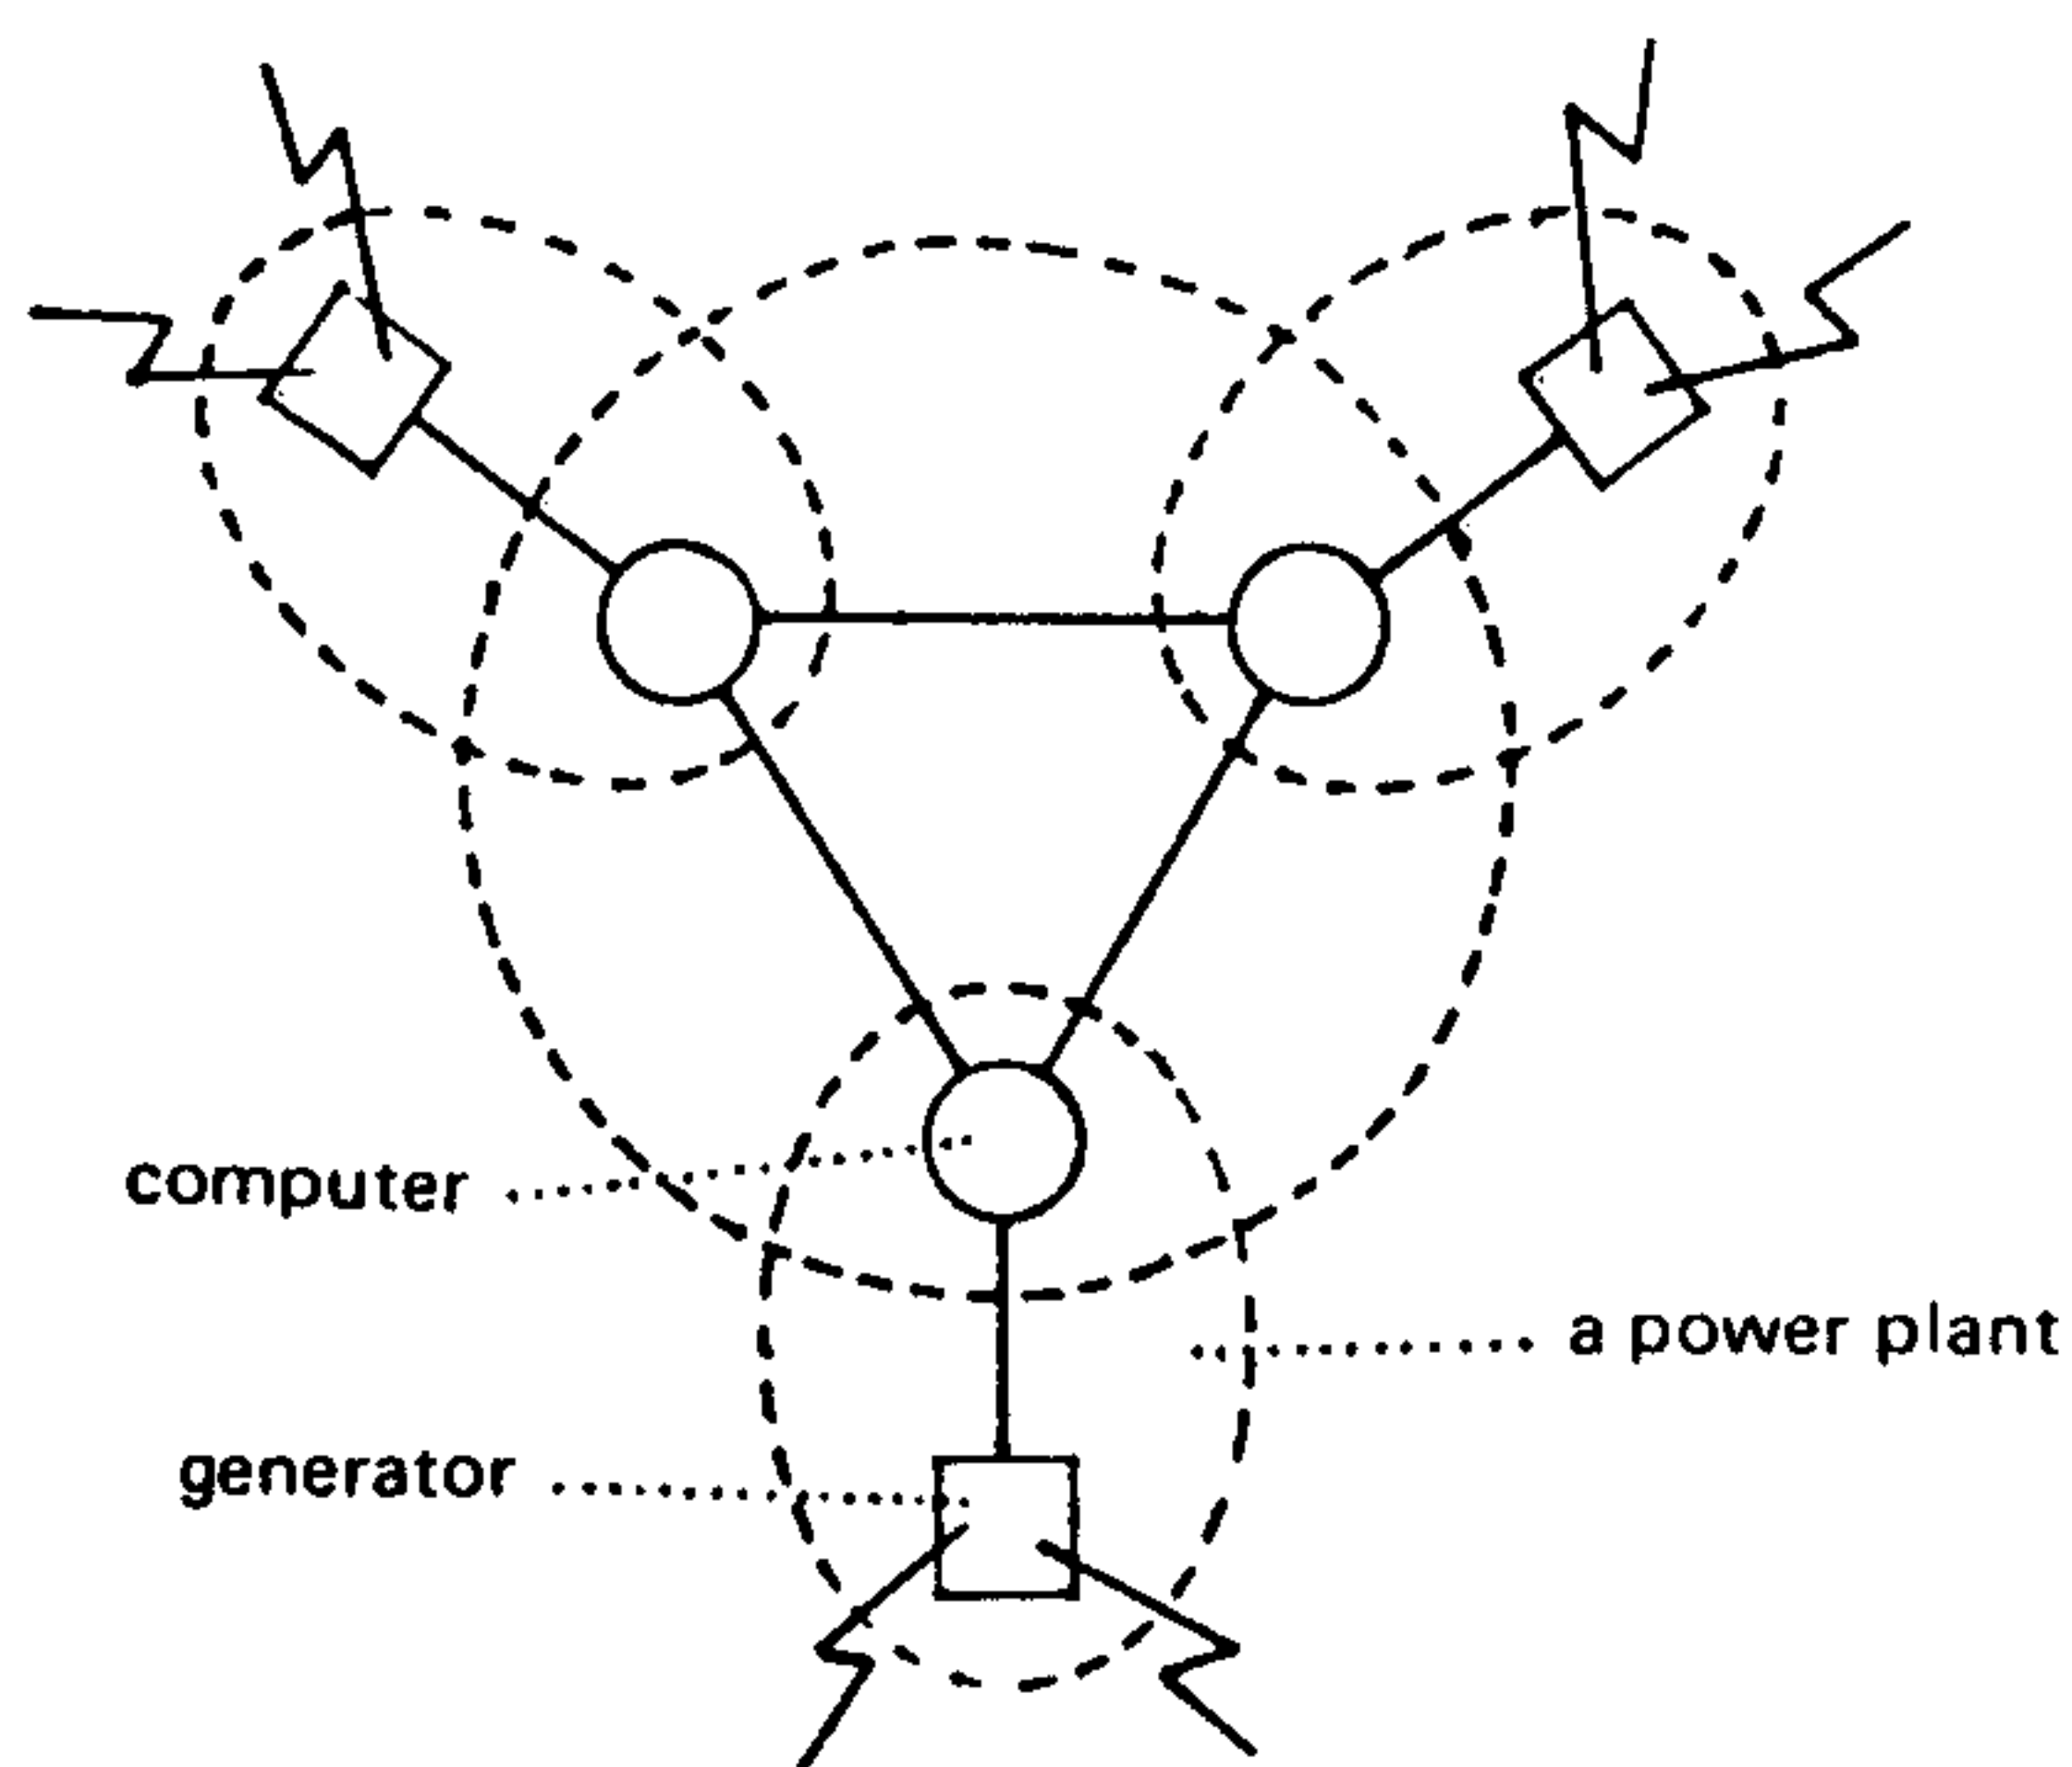


Figure 1. Three linked Power Plants

Evidently we can think of this in at least two ways, as the diagram suggests; either as three power plants linked together, or as three generators each linked to a process consisting of three interlinked computers. However we choose to sub-divide the system, it is most important that we are able to

treat any collection of processes, interacting with each other, as a single process at a higher level of abstraction. In other words, just as with specifications, we must be able to compose and decompose processes hierarchically if we are to have a useful theory.

What I now want to show is that there is a notion of process which is so general and so mathematically elementary that it rivals a notion which has become fundamental to a large body of mathematics – the notion of a set.

Perceived from the outside, all we can see of a process in a state P is that it can perform any one of a number of possible actions, and move to a new state P' say:

$$\begin{array}{ccccc} P & \xrightarrow{a} & P' \\ \text{state} & \text{action} & \text{state} \end{array}$$

Since it can have several alternative actions in any state, its full capability can be arranged as a tree-like diagram:




Figure 2. Action tree of a Process

Notice that some of these branches may just terminate (the process stops) while others may proceed infinitely. In fact with the help of David Park at Warwick, we have worked out a theory in which a process is almost exactly a tree of this kind.

Now consider the mathematical notion of set, which may contain members which are atoms or individuals and other members which are again sets. Just as a process can be described by a tree-like diagram which tells its *action* story, so a set may be described by a tree-like diagram which tells its *containment* story:




Figure 3. Containment tree of a Set

The analogy is that process corresponds to set, and process action corresponds to set containment (we can handle the slight difference that actions are of many kinds, while containment is just of one kind).

Take the analogy a step further. When are two sets S and T equal? Answer: when they contain equal members. Some naturally say that two processes P and Q are equal when, for each action, their successors upon that action are equal.

Now we come to a fascinating difficulty. We feel uncomfortable if our test for equality of sets gets into an infinite regress, which happens if there are containment chains which go on for ever.

S contains S_1 contains S_2 contains ...?

Such sets, where we can go on for ever taking a member without reaching an atom, are called non-wellfounded; because of the infinite regress, mathematicians have usually excluded these sets. But we don't want to exclude infinite processes! All sorts of simple objects, like parking-ticket machines or computers, are perfectly modelled by infinite processes (if we ignore physical decay). Recently a Manchester logician, Peter Aczel, argued that there is no need to exclude non-wellfounded sets either; he proposed what he calls the antifoundation axiom, which explains their nature in a very natural way. Within the last year or so, partly during his visit to Edinburgh, he discovered that the theories of infinite processes on the one hand, and non-wellfounded sets on the other, are essentially one and the same theory (15).

For me this discovery is exciting. True, it doesn't alter the theory in any way, nor does it provide very detailed suggestions on how to develop the theory of processes. In fact, the pragmatic ramifications of the theory of processes are likely to look very different from the ramifications of set theory which are important for the foundations of mathematics. It indicates that the concepts which are useful in a theory of how to build complex concurrent systems are not firmly established in traditional mathematics, and may even have an influence on the development of mathematics. I cannot see why school textbooks in the year 2000 should not treat processes with as much familiarity as present day textbooks treat wellfounded sets. In other words, processes may make as much claim to existence as wellfounded sets; this is what I meant earlier by questioning the distinction between a 'theory of what' and a 'theory of how'.

This philosophical digression may strike fear and doubt in the minds of people who hope that we shall do useful practical work and would like to use the results. But it should not; already this

process work has shed light upon useful practical systems such as the communications protocols mentioned earlier. We should not be surprised if the concepts which inform a new kind of practice also arouse a new kind of theoretical interest; perhaps we should rather expect it to be so.

Composing calculi

Having given some examples of concepts, like types and process, around which we can build theories of how to design systems, I would like to consider how we can work with these concepts in practice, in a way which permits serious design experiments, which will in turn reinforce (or undermine) the theories.

Simple experiments can be done with pen and paper; anything more complex which will test whether a theory scales up, is going to need more sophisticated tools to help us manage all the details. Now we have seen that types of specifications can be expressed in a logical system and the same is true of processes. It would indeed be nice if some single logical language, with a single axiomatic base, were to satisfy all our needs, because then we could set about providing a tool kit of programs to assist designers in using this single calculus. Work along these lines is being done, and will certainly be a valuable experiment; a noteworthy example is the specification language Z designed by Abrial, which is being developed and put to experiment by Hoare and his group at Oxford (16).

There is evidence to suggest that this approach is limited. For all the success of logicians in founding a large part of mathematics upon a single logic, they have never succeeded — and seldom claimed that they would succeed — in getting mathematicians to do their work in this one medium. The different branches of mathematics are too various, and our linguistic invention too fertile for this kind of uniformity to be sustained. The subject matter of computation is equally various, so we would expect the same need for plurality; indeed Kim Larsen and I recently did a small experiment in process validation and found that, in the space of three or four pages, we needed three or four distinct formal calculi to express the natural proof succinctly.

So how do we set about building computer systems which will help people to do different design experiments in whatever combination of calculi is most appropriate? The broad answer follows the general pattern which we have already followed, but we are now working one level higher. We wish to design systems which will help people to design systems; we therefore need a 'theory of how' again: a theory of how to design

these meta-systems. Here are three analogous questions:

1. What is a specification, and how can specifications be composed hierarchically?
2. What is a process, and how can processes be composed hierarchically?
3. What is a calculus, and how can calculi be composed hierarchically?

We looked at the first two questions earlier, and we are now confronted with the third question. We are using the term 'calculus' to mean any systematic way of reasoning; it embraces both the theoretical and the pragmatic aspects of logic.

Consider a simple example. The calculus of matrices involves certain objects — scalars, vectors, matrices; it involves elementary operations like sum and product over these, and more complex operations like partitioning; it involves basic algebraic properties; it involves methods of calculation (e.g. of the inverse); at a more practical level it involves convenient visual means of display. Most important, it is a parametric calculus working over any fields of elements. This shows that we must implement calculi in such a way that we can compose the calculus of matrices and the calculus of real numbers (with all their pragmatic aids) to obtain the calculus of real matrices. The same story can be told, almost word for word, for a calculus of specifications or of processes.

This relentless pursuit of hierarchy is now becoming familiar! All the same, its application to logical calculi is something rather new, and is strongly guided by the need for really practical reasoning tools. Some recent work at Edinburgh will help us towards a pragmatic definition of logic, which will allow logics to be built for component logics. Burstall and his group have used category theory to define a notion called institution (17) which explains the composition of algebraic theories. At a more basic level, Plotkin (with some ideas from Paulson at Cambridge) is looking at how the apparently simple idea of inference rule should be presented to a computer (18).

This subject is too taxing to pursue in a general paper; perhaps I have already gone too far for comfort. There is, however, a need to expose this problem of implementing calculi in a hierarchical way, because it is so pervasive to the kind of system-analysis that we wish to promulgate into practical use. Indeed the projects which have already been funded for our Laboratory are all seriously involved with building calculi which can be used in earnest; we intend to grapple with the theory of how to do this, side-by-side with solid and specific experiments in which industrial designers can participate.

One ingredient which is crucial to this work is the implementation medium. About ten years ago we designed ML, a programming metalanguage for implementing logic. There was no idea at that time of composing logics hierarchically but in the intervening years the technology of structured functional programming has advanced impressively. Due to a continual flow of ideas about parameterization from Burstall, MacQueen, Plotkin and others, we now have a vastly superior version of ML, whose power in handling parametric structures exceeds that of any other language in serious use. The language has probably been submitted to more theoretical analysis than any other, and is now reaching its last stages in implementation, in a development project funded initially the Science and Engineering Research Council (SERC) Software Technology Initiative and then adopted by the Alvey Directorate. The language is widely known and its implementation is highly portable; this means that experiments can be conducted on a much wider basis (19).

The Laboratory

I would like to conclude by arguing that now is the right time and Edinburgh is a good place from which to mount a vigorous attack on computation theory founded on experiment, and by outlining where we have got to and what more is needed to mount this attack.

First, the Edinburgh environment is extraordinarily propitious for the exercise. As explained, we already have successful experience in building theories, with encouraging feedback from application. The Department of Computer Science fully supports the Laboratory, and we have never doubted that the Laboratory should grow within the department rather than as an independent entity. First, the Laboratory has a commitment to teaching — both undergraduate and postgraduate — and everything must be done to meet the huge demand for scientists and engineers who are well-grounded in the foundations of the subject. Second, we reject any artificial barrier between theory and engineering in computer science, and anticipate a healthy flow of ideas between the Laboratory and the department as a whole. This is already exemplified by the work of George Milne in applying algebraic design methods to VLSI (20), and should extend to an interaction between theory and more general systems design, both hardware and software, as the department intensifies its activity in this area.

Beyond our own department, we are lucky to have both the Artificial Intelligence Department and the Centre for Cognitive Science as neighbours. The AI department has a strong research effort in mechanized reasoning with which we have always interacted, and in our relationship with industry

we hope to form useful links with the AI Applications Institute. Cognitive science, on the other hand, represents the pervasive role played by logic in all branches of IT. I suspect that we all find it baffling to try to demarcate between the sciences of computing, intelligence and cognition; fortunately there is so much to do that no-one finds any value in boundary disputes! The problem is how to benefit as fully as possible from each other's pursuits; we are grateful for the efforts of the Committee for Information Technology in giving some co-ordination to this IT research community.

If Edinburgh is a good place for our Laboratory, then now is the right time. What makes it right is the strong synergy between academia and industry; they are pursuing the same goal of organized (rather than *ad hoc*) understanding of system design. Industry's motive is to produce systems which sell — ours must be to understand the subject for its own sake. I have tried to show not only that the same kind of understanding is needed for both purposes, but also that the method of reaching this understanding is largely the same — it lies in protracted and organized experiment with theories. For this reason we are confident that Industry will be interested in taking part in our experiments; indeed, our joint projects with ICL, Software Sciences and BP already give firm evidence of this.

The Alvey Directorate, as well as Industry, have already given us considerable financial support. Alvey have funded the post of Assistant Director of the Laboratory; they are supporting our ML development project, and have recently bought us a Pyramid computer. Finally a very strong commitment has come from SERC who have approved further funding for our research projects over the next three or four years. This has placed us in a solid position; the theoretical backbone of the Laboratory is now assured. These projects are built around the implementation of logical calculi as a focus. What we shall be doing over the next few years is to surround this backbone with an increasing volume of interaction with applications; for this we look to Industry for increasingly active support, and to the University to help us find the space in which to exploit this support in the fullest possible manner. If we do our job properly we shall also have contributed something permanent to our subject and made the case for establishing our Laboratory as a permanent entity.

Notes


(1) This paper is based on the text of the inaugural lecture of the Laboratory for the Foundations of Computer Science at Edinburgh University. The editors have

made minor modifications and have added these notes to help the reader of the Journal.

- (2) The classical paper that expresses this is by Church, A. An unsolvable problem of elementary number theory. *American Journal of Mathematics*, 58, 345-63, reprinted in: M. Davis (ed), *The Undecidable*, Raven Press, New York, which also contains other classics.
- (3) A very readable and remarkably foresighted paper is Burkes, A.W., Goldstine, H.H. and von Neumann, J. (1946) Preliminary discussion of the logical design of an electronic computing instrument. US Army Ordnance Department Report, reprinted in: Bell, C.G. and Newell, A. (1971) *Computer Structures: Readings and Examples*. New York, McGraw Hill.
- (4) See, for example, Scott, D.S. and Strachey (1971) Towards a mathematics semantics for computer languages. *Technical Monograph PRG-6*, Programming Research Group, University of Oxford.
- (5) A structural approach to operations semantics. Report DAIMI FN-19, Computer Science Department, Aarhus University 1981. For a current overview of the whole field the reader is referred to Schmidt, D.A. (1986) *Denotational Semantics: a methodology for language development*. Allyn and Bacon.
- (6) Hoare, C.A.R. (1969) An Axiomatic Basis for Computer Programming. *CAM. ACM*, 12, 576-80pp.
- (7) Kowalski, K. (1979) *Logic for Problem Solving*. North Holland, Artificial Intelligence Series.
- (8) Gordon, M.J., Milner, R., and Wadsworth, C.P. (1979) Edinburgh LCF. *Lecture Notes in Computer Science*, 78, Springer-Verlag.
- (9) Gordon, M.J.C. (1983) LCF-LSM: A system for specifying and verifying hardware. *Technical Report 41*, Computer Laboratory, University of Cambridge.
- (10) Milner, R. (1980) A calculus of communicating systems. *Lecture Notes in Computer Science*, 92, Springer Verlag.
- (11) LOTOS. A formal description technique based on the temporal of observation behaviour. ISO draft proposal 8807, March 1985.
- (12) Burstall, R.M. and Goguen, J.A. (1980) The semantic of CLEAR, a specification language. *Lecture Notes in computer science*, 86, Springer Verlag.

- (13) Martin-Lof, P. (1982) Constructive mathematics and computer programming. Cohen, L.J., Los, J., Pfeiffer, H. and Podewski, K.P. (eds) *Proc. 6th Int. Congress for Logic, Methodology and Philosophy of Science*. pp.153-75, North Holland.
- (14) The Alvey Directorate (Millbank Tower, Millbank, London SW1P 4QV), oversees the British National Research Programme in Information Technology.
- (15) Aczel, P. (1987) The antifoundational axiom. *CSLI Lecture Notes*. Stanford University, Chicago University Press.
- (16) The Z handbook (Sufrin, B (ed), Programming Research Group, Oxford University, 1986) provides an overview of the notation and Specification Case Studies (Hayes, I. (ed) London, Prentice-Hall Int., 1986) provides a useful perspective; see also Working with Formal Methods Nicholls, J., this volume.
- (17) Burstall, R., and Goguen, J.A. Introducing institution. Clarke, E., Kozen D., (eds) *Proc. Logics of Programming Workshop*. Carnegie-Mellon University, Springer-Verlag Lecture Notes in Computer Science, Vol. 164, pp.221-56.
- (18) Harper, R., Honsell, F. and Plotkin, G.D. (1987) A framework for defining logics. *Proc. 2nd LICS Conference Ithica*. New York.
- (19) Harper, R., MacQueen, D. and Milner R. (1986) Standard ML. *Report ECS-LFCS-86-2*. Department of Computer Science, University of Edinburgh.
- (20) Milne, G.J. (1984) A model for hardware description and verification. *Proc. 21st Design Automation Conference*. IEEE Computer Science Press.

Biographical notes



Robin Milner has worked in the Theory of Computation for the past 24 years, at City University in London, at University College Swansea, at the Artificial Intelligence Laboratory in Stanford, and for the past 14 years at Edinburgh University. His main achievements are in semantics and in the methodology of computer-assisted reasoning. He devised LCF (Logic for Computable Functions), an interactive proof system, which employed a programming meta language called ML now in widespread use. He also introduced the Calculus of Communicating Systems, a theory of concurrent and distributed computation. He is now Professor of Computation Theory at Edinburgh, where he also directs the Laboratory for Foundations of Computer Science.

Address for correspondence: Laboratory for Foundations of Computer Science, Department of Computer Science, JCMB, The King's Buildings, Edinburgh EH9 3JZ.