PROGRAMMING
IS AN
ENGINEERING PROFESSION

C.A.R. Hoare

Technical Monograph PRG-27
May 1982

Oxford University Computing Laboratory
Programming Research Group
45 Banbury Road
Oxford OX2 6PE.

Oxford University Computing Laboratory
Programming Research Group
45 Banbury Road
Oxford OX2 6PE.

Summary.

The ideals of professional practice for programmers engaged in large scale computing projects are based on a sound understanding of the underlying mathematical theories, and they should follow closely the traditions of engineers in better established disciplines. But there are a number of important differences which must not be neglected. It will not be easy to attain the ideals described, but the key to success is found in an improvement in the education of programmers.

PROGRAMMING
IS AN
ENGINEERING PROFESSION

In earlier times and in less advanced societies, the welfare of a community depended heavily on the skill and dedication of its craftsmen -- the millers and blacksmiths, spinners and weavers, joiners and thatchers, cobblers and tailors. A craftsman possesses special skills, not shared by his clients, which he has acquired by long and ill-paid apprenticeship to a master of his craft. He learns exclusively by imitation, by practice, by experience, by trial and error. He knows nothing of the scientific basis of his techniques, nothing of geometry or even of drawing, nothing of mathematics, or even of arithmetic. He cannot explain how or why he does what he does: and yet he works effectively, by himself or in a small team, and can usually complete the tasks he undertakes in a predictable timescale and at a fixed cost, and with results that are predictably satisfactory to his clients.

The programmer of today shares many of the attributes of the craftsman. He learns his craft by a short but highly paid apprenticeship in an existing programming team, engaged in some ongoing project; and he develops his skills by experience rather than by reading books or journals. He knows nothing of the logical and mathematical foundations of his profession. He does not like to explain or document his activities. Yet he works effectively, by himself or in small teams, and he sometimes manages to complete the tasks he undertakes at the predicted time within the predicted costs, and to the satisfaction of his client.

In primitive societies of long ago we hear of another class of specialist on whom the welfare of the community depended. Like the craftsman, he is dedicated to his task; like the craftsman he is regarded with respect, perhaps even tinged with awe, by his many satisfied clients. There are several names given to such a man -- a seer, a soothsayer, a sorceror or wizard, a witch doctor or high priest. I shall just call him a high priest.

There are many differences between the craftsman and the high priest. One of the most striking is that the high priest is the custodian of a weighty set of sacred books, or magician's manuals, which he alone is capable of reading. When he is consulted by his client with some new problem, he refers to his sacred books to see whether he can find some spell or incantation which has proved efficacious in the past; and having found it, he tells his client to copy it carefully and use it in accordance with a set of elaborate instructions. If the slightest mistake is made in copying or in following the instructions, the spell may turn to a curse, and bring misfortune to the client. The client has no hope of understanding the nature of

the error or why it has evoked the wrath of his deity -- the high priest himself has no inner understanding of the ways of his god. The best the client can hope is to go right back to the beginning, and start the spell again: and if this does not work, he goes back to the high priest to get a new spell.

And that is another feature of the priesthood -- when something goes wrong, as it quite often does, it somehow always turns out to be the ignorance or stupidity or impurity or wickedness of the client: it is never the fault of the high priest or his god. It is notable that when the harvest fails, it is the high priest who sacrifices the king, never the other way round.

Programmers of the present day share many of the attributes of the high priest. We have many names -- coder, systems analyst, computer scientist, informatician, chief programmer; I shall just use the word 'programmer' to stand for them all. Our altars are hidden from the profane, each in its own superbly air-conditioned holy of holies, ministered to night and day by a devoted team of acolytes, and regarded by the general public with mixed feelings of fear and awe, appropriate for their condition of powerless dependence.

An even more striking analogy is the increasing dominance of our sacred books, the basic software manuals for our languages and operating systems which have become essential to our every approach to the computer. Only thirty years ago our computers' valves and tanks and wires filled the walls and shelves of a large room, which the programmer would enter, carrying in his pocket his programming manual -- a piece of folded cardboard known as the FACTS CARD. Now the situation is reversed: the programmer enters a large room whose walls and shelves are filled with software manuals; but in case he wants to carry out some urgent calculations he carries in his pocket -- a computer.

The rise of Engineering.

In recent centuries with the advance of technology, we have seen the emergence of a new class of specialist -- the professional engineer. The most striking characteristic of an engineer is the manner in which he qualifies for entry into his profession, not only does he work out the long apprenticeship of the craftsman, not only does he undergo the brief graduation or initiation ceremonies of the high priest, but both of these are preceded by many years of formal study in schools and in universities. His education covers a wide range of topics, including the mathematical foundations of the differential calculus, the derivation and solution of complex equations, the physical principles underlying the science of materials, as well as the specific technicalities of a particular branch of his subject, and a large

Step 7.2. Let g be the <generation> immediately contained in abuf. If abuf contains a <key>, let k be this <key> and let csvb be its immediate component; otherwise let k and csvb be <absent>. Perform construct-record(g,k) to obtain kr.

Step 7.3. If fi contains <keyed> then if k is equal to any <key> in the <record-dataset>, designated by the <dataset-designator> in fi, or if k is unacceptable to the implementation, then perform raise-io-condition(<key-condition>,fv,csvb).

Step 7.4. Perfotm insert-record(kr,fv) to obtain pos.

Step 7.5. Replace the immediate component of the <current-position> in fi with pos.

Step 7.6. Perform free(g) and delete abuf from fi.

Step 8. Let dd be the <data-description> immediately contained in the <variable> of the <declaration> designated by cdp. Perform evaluate-data-description-for-allocation(dd) to obtain edd.

Step 9. Perform evaluate-size(edd) to obtain an <integer-value>,int. If int is unacceptable to the implementation then perform raise-io-condition(<record-condition>,fv,chs) and optionally perform exit-from-io.

Step 10. Perform allocate(edd) to obtain g.

Step 11. Let desc be a <data-description> simply containing <pointer> without other terminal subnodes. Let epsog be an <evaluated-target> containing the <generation> in the <evaluated-pointer-set-option> in els. Let agv be an <aggregate-value> containing <pointer-value>: g. Perform assign(epsog,agv,desc).

Step 12. Let d be the <declaration> designated by the <declaration-designator> in els.

Step 12.1. If the <aggregate-type> of g contains <structure-aggregate-type> then perform initialize-refer-options(g).

Step 12.2. Perform initialize-generation(g,d).

Step 13. Let abuf be an <allocated-buffer>: <generation>,g. If fi contains <keyed> then attach kk to abuf. Attach abuf to the <file-opening> in fi.

Step 14. Perform normal-sequence.

8.6.4 THE REWRITE STATEMENT

Purpose: The <rewrite-statement> causes replacement of an existing <record> or <keyed-record> in a <record-dataset>.

8.6.4.1 Execute-rewrite-statement

<evaluated-rewrite-statement>::= <file-value>
[(<key>] <evaluated-from-option>]

Operation: execute-rewrite-statement(rws)

where rws is a <rewrite-statement>.

Step 1. Let erws be an <evaluated-rewrite-statement> without subnodes.

Step 2. Perform Steps 2.1 through 2.3 in any order.

Step 2.1. Let f be the immediate component of the <file-option> in rws. Perform evaluate-file-option(f) to obtain a <file-value>,fv. Attach fv to erws.

Step 2.2. If rws contains a <from-option>,fr, then perform evaluate-from-option(fr) to obtain an <evaluated-from-option>,efo and attach efo to erws.

catalogue of known design methods and specific practical techniques... But this is only a start: during his professional career, the engineer will expect to continue his education, to expand his skills, and to keep pace with technological progress by continued study of new books and learned journals, and attendance at specialist orientation courses. Many engineers will even take a full year off work to bring themselves up to date, or to reorient themselves to a newly developed branch of technology. The older craftsmen will complain that the engineer already knows far more than he needs in the day-to-day practice of his profession; but his colleagues and clients will realise that the weight of background learning develops his good judgement and increases his competence and authority at all times; even if a recondite scrap of knowledge is used only once in his career, then the learning has paid for itself many times over.

We would like to claim that computer programming has transcended its origins as a craft, has avoided the temptation to form itself into a priesthood, and can now be regarded as a fully fledged engineering profession. Certainly, we have some right to this claim. Through our professional Societies we have formulated a code of professional ethics and a structure and syllabus of professional examinations. We discharge our duty to the community by giving evidence to government commissions on social consequences of computing, on privacy, on employment. Because of the great demand for our services, our clients and employers are willing to offer us professional salaries, and it is hardly likely we shall refuse them.

But more than this is needed for true professional status. What is the great body of professional knowledge common to all educated programmers? Where are the reference libraries of standard works on known general methods and specific techniques and algorithms oriented to particular applications and requirements? What are the theoretical, mathematical or physical principles which underlie the daily practice of the programmer? Until recently, these questions had no answer. Now the answers are beginning to emerge. We can point to the ACM curriculum for the study of Computer Science at University as a corpus of common knowledge for the programmer, though the proportion of Computer Science graduates in the programming profession is still low. Don Knuth's books on the Art of Computer Programming form an excellent encyclopaedia of known techniques -- but only three volumes have so far appeared, and how many programmers consult even those? And finally, we have only recently come to a realisation of the mathematical and logical basis of computer programming: we can now begin to construct program specifications with the same accuracy as an engineer can survey a site for a bridge or road; and on this basis we can now construct programs proved to meet their specification with as much certainty as the engineer assures us his bridge will not fall down. Introduction of

these techniques promises to transform the arcane and error-prone craft of computer programming to meet the highest standards of a modern engineering profession.

The Art of Computer Programming
Vols 1, 2, 3.
D.E. Knuth, Addison-Wesley.

Structured Programming
Dahl, Dijkstra, Hoare. Academic Press, 1972.

Systematic Programming
N. Wirth. Prentice-Hall, 1973.

Principles of Progrem Design
M.A. Jackson. Academic Press, 1975.

A *Discipline of Programming*
E.W. Dijkstra. Prentice-Hall, 1976.

The Architecture of Concurrent Programs
Per Brinch Hansen. Prentice-Hall, 1977.

Some books you may find on the shelf of
the well-read computer professional.

Let me expand on the nature and consequences of this discovery. It is like the Greek discovery of axiomatic geometry as the basis of the measurement of land, mapmaking, and its later use in plans and elevations for the design and construction of buildings and bridges. It is like the discovery of the Newtonian laws of motion and the differential calculus as the basis of astronomy as well as more mundane tasks like the navigation of ships and the direction of artillery fire. It is like the discovery of stress analysis as the basis for the reliable and economic construction of steel frame buildings, bridges, and oil platforms.

Large Programming Projects.

In future we may hope to see a radical change in the development and life history of large programming projects. The chief programmer, like the architect, will start by discussing requirements with his client. From education and experience, the programmer will be able to guide his client to an understanding of his true needs and avoidance of expensive features of dubious or even negative value. From respect for the professional status of the programmer, the client will accept and welcome this guidance. This kind of mutual understanding and respect is essential to any relationship between a professional and his client or employer.

The chief programmer at this time sketches out the overall structure of the specification of a product to meet his client's requirements. These sketches serve the same role as an architect's preliminary sketches of a building. Gradually, in orderly fashion, and in close consultation with the client, details of the design will be slotted into the appropriate place within the structure. This activity will culminate in a complete, unambiguous and provably consistent specification for the entire end product. It will serve the same role as blueprints in engineering or scaled plans and elevations in architecture.

The sign N means *number (positive integer)*.
The sign 1 means *unity*.
The sign a+1 means *the successor of* a
or a *plus* 1.
The sign = means *is equal to*.

1. $1 \in N$
2. $a \in N \;.\supset.\; a=a$
3. $a,b \in N \;.\supset:\; a=b \;.\equiv.\; b=a$
4. $a,b,c \in N \;.\supset:.\; a=b \;.\; b=c \;:\supset.\; a=c$
5. $a=b \;.\; b \in N \;:\supset.\; a \in N$
6. $a \in N \;.\supset.\; a+1 \in N$
7. $a,b \in N \;.\supset:\; a=b \;.\equiv.\; a+1 = b+1$
8. $a \in N \;.\supset.\; a+1 \neq 1$
9. $k \in K \;:.\; 1 \in k \;:.$
$x \in N \;.\; x \in K \;:\supset_x.\; x+1 \in k \;::\supset.\; N \supset k$

Extract from Peano.
Arithmetices Principia 1889.

Undoubtedly, the client will ask to see and check the full specification before he gives permission to go ahead with implementation. I'm afraid he will get a rude shock. Instead of pretty pictures and drawings, he will see a collection of definitions, mathematical formulae and logical proofs which he may be ill equipped to understand. One of the major problems of the programming profession is that our technical and structural decisions are almost invisible; there is nothing that can be seen in the finished program which can be illustrated beforehand by pictures. This sad fact explains simultaneously the persistent longevity as well as the basic futility of program flow charts.

A proper solution to this communication gap between programmers and client may be discovered by analogy from other professions. Before a building project goes into implementation, the architect produces from his specification a series of perspective drawings or even models which can be shown to the client, and carefully checked by him. Before a consumer product goes into mass production, an engineer produces a series of working prototypes, which can be subjected to severe and exhaustive test in a variety of simulated circumstances. In future, a chief programmer will be able with the aid of his programming teams to pursue both of these solutions at the same time.

Firstly, the formal specification is taken as the basis of a clear, complete and consistent set of user manuals and operating instructions, explaining exactly how to control the program and how it will behave in all circumstances, including when things go wrong. Of course, these manuals are illustrated by compelling examples, dealing with the main common cases; but the examples will be backed up by well structured and well indexed descriptions of the full range of the capabilities of the program, explaining why and when they are needed, how they can be successfully invoked in conjunction with other features, what can go wrong, and how to recover from failure. It is these manuals that give the customer a full understanding of what his program will look like and what it will do for him, long before a single word of code is written. They will be much clearer, much more complete, and much shorter than manuals of the present day, because they are firmly based on the simple mathematical model, in the same way that Newton's Laws of Motion are shorter and more illuminating than the planetary observations of Tycho de Brahe.

At the same time, the chief programmer or his colleagues may construct a prototype of the program as a whole, or of the more vital parts of it. Such a prototype may be cheaply programmed as a simulation, perhaps running on a small model of the data base held in the main store of a computer much larger and faster than the one on which the eventual program will run. These simulations are exact scale models of the final design, and can be used by the client to check the details

> "Computers are extremely flexible and powerful tools, and many feel that their application is changing the face of the earth... [But] their influence as tools might turn out to be but a ripple on the surface of our culture, whereas I expect them to have a much more profound influence in their capacity as intellectual challenge."
>
> E.W. Dijkstra, 1972.

of the design and suggest alterations before the project goes into the more expensive stage of design and implementation.

The construction of models and prototypes is not cheap; but in a large and important project is amply justified by the chance it gives to modify the design in the light of informed customer experience. Recovery from mistakes in design is much more expensive when they have been cast into the concrete of a million-line program.

Another important task can be completed at this stage. On the basis of the original requirements and formal specifications, it is possible to devise a series of rigorous and searching acceptance tests, which can be included in the contract between the client and the implementors. Some of these tests can be kept secret from the implementors, so that there is no temptation for them to orient their work towards passing the tests, rather than meeting the specification. This rigorous kind of secret acceptance test is made possible only by the corresponding mathematical rigour of the original specification: if the product fails the test, and the implementors claim that the test is unfair, any competent logician or mathematician would be able to decide who *is* right.

Implementation.

The next stage is to start work on the overall design of a program to meet the agreed and tested specification. The major components of the design are identified, and the interfaces between them are defined with mathematical precision. Some of the required components are selected or perhaps adapted from a library of existing components described in the engineering textbooks. The remaining components are specified with the same techniques and with the same care as used in the earlier design of the complete program. But most important: the chief programmer convinces himself and his colleagues by mathematical proof that if each of the components meets its specification, then when all the components are

assembled together, the overall product will meet the overall specification agreed by the client. In future this will be taken for granted, just as we now take for granted the fact that components of a bridge ordered to given measurements will fit together when they are assembled on site. So we hope to eliminate the so-called "system integration" phase of many current projects, in which bugs are painfully detected and laboriously removed from the interfaces between the components. This is the most expensive and unpredictable of all the phases of a large project: the fact that it is the final phase only increases the misery.

> "Program testing can be used to show the presence of bugs but never to show their absence."
>
> E.W. Dijkstra, 1972.

Why is debugging so expensive, particularly at the stage of system integration and afterwards in program maintenance? The reason is clear: the bugs involved are so subtle that they escaped the attention of the designer at a time when the design was still simple, and he still had all his options open. They escaped the attention of the programmer when he was devoting his best intellect to each line of code. Now they must be isolated in the context of a million-line program; and they must be eliminated under the additional and even more onerous constraint of changing as few of those million lines as possible! No wonder program maintenance during the whole life of a program is often many times as expensive as the original implementation. Using the new specification and design techniques of mathematics and logic, we hope to eliminate most of that cost, by never creating the bugs in the first place.

When the design has progressed sufficiently, it will be possible to build up teams, make plans and schedules, to estimate sizes and performance of the code, and above all to check preliminary estimates by calculation of the overall costs and timescales of implementation. This corresponds to the activity of quantity surveying in architecture, and requires experience and judgement at least as much as mathematical technique. Nevertheless, the estimates will be more accurate than they usually are nowadays, because they are based on complete and consistent and stable specifications and designs.

At last the project is ready to go into the construction phase. Now large teams of programmers can be engaged, perhaps from independent contractors or software houses, and all of them can work concurrently on different parts of the

design, without further consulting each other. Each programmer will use standard techniques of stepwise development to ensure that his code meets its specification, with minimal risk of the intrusion of error. When he has proved that his code is correct, both the code and proof will be signed off by a highly paid checker; and the code will then be typed into a computer.

Delivery.

When all the code is complete and compiled from its high level language, and loaded into the computer it will be subjected to the implementor's tests, which it will usually pass. It will then be delivered to the customer, and pass his secret acceptance tests as well. Since all manuals have been available for training, it will go into immediate service. Nothing can possibly ever go wrong. What never? Well hardly ever! On the rare occasion of failure there will be a full and independent enquiry, and the cause of the fault will be traced to the persons responsible. An independent assessment will be made to determine whether the fault is an isolated one, or whether it is a symptom of more serious and widespread flaws in the logic of the design or in the technique of implementation. In the latter case, large parts of the documentation and code and proofs will be rechecked by experts before the product is delivered again to the customer, and submitted to newly constructed secret acceptance tests. The payment of appropriate penalties to the customer will ensure that this kind of default is not too frequent.

In the years after the first delivery, it is very likely that the customer's requirements will change, and the program must be changed with it. Because of the clarity of program structure and the completeness of design documentation, it

> "It is reasonable to hope that the relationship between computation and mathematical logic will be as fruitful in the next century as that between analysis and physics in the last. The development of this relationship demands a concern for both applications and for mathematical elegance."
>
> John McCarthy, 1967.

will be quite easy to determine which parts of the design and coding need to be changed in order to meet a new requirement. Because all the assumptions and obligations of each piece of code have been made explicit, it is relatively easy to prove that a new piece of code which meets the same obligations can be safely inserted; and if the obligations can no longer be met, it is possible to identify all

other pieces of code which rely on these obligations, so that this code can be changed too. When a suggested change violates the fundamental structure of a program, the programmer will rack his brains to think of an alternative; and if he can't, he will know in advance that part or all of the program must be rewritten, and check that the cost is acceptable. Thus it is possible to escape the wild goose chase after consequential effects of each change made to a large program, which is common today.

That concludes my description of the life cycle of the large software project of the future. The description hardly makes reference to the most common feature of present programming practice, the program bug. I have left it out only because it won't exist. There will be no bugs. There will be no chance for a bug to germinate or to propagate. Every stage of the specification and design and coding will have been checked with mathematical rigour. It is an essential feature of the work of a professional in any discipline that he organises his working environment and his working methods to ensure that he does not make mistakes. Most pilots never crash a plane. Most surgeons never kill a patient. Most civil engineers never build a bridge which collapses. Until each programmer displays this kind of professional accuracy and responsibility, all our claims to professional status are subject to doubt. Every time a member of the public blames "the computer" for an error made by a programmer, it demeans our profession. Every time that a supplier of software writes a disclaimer of direct and consequential damages arising from its errors, it demeans our profession. We must always confess that it is the programmer who bears the responsibility for mistakes not the dumb but accurate machine; we must always point out that unfair disclaimers of responsibility are (or should be) forbidden by law.

Of course, my remarks apply only to large and important projects. In smaller less important projects many of the stages may be merged or omitted; and for the smallest projects (eg a program written for a single run by its own author), none of what I have said is relevant. One does not use structural engineering analysis to build a sandcastle. But neither does one choose the prize-winning builder of sandcastles as architect for a tower block of offices in a city.

Comparison with other engineering disciplines.

My description of the planning of large-scale programming projects follows closely the standard practices in more traditional branches of engineering. A conventional engineering design passes through the established phases of requirements analysis, specification, design, costing, production engineering, drawing office, prototyping, testing, toolbuilding, quality assurance etc.; it is many years before the design reaches the production floor. And indeed, many data processing

departments of the present day are organised on the basis of a similar division of labour between systems analysts, programmers, technical authors, coders, testers, and finally maintenance programmers.

But all too often this apparently logical division of labour leads to an awkward problem. Gradually, the size of the maintenance programming department increases until it outnumbers all the other groups put together. And it is increasingly difficult to recruit and retain computer programmers for this boring, ill-regarded and often poorly paid occupation. One likely cause for this problem is that the interfaces between the various groups of programmers have been less precisely defined than in a traditional engineering workshop, and that there is no proper quality control on the project documentation as it passes from one group to the next. As a result, each group does its best with what it gets, and it is the poor maintenance programmer at the end of the chain who has to pick up the pieces.

In my view, the standards that must be met by project documentation as it passes between groups are standards of logical accuracy and completeness which are characteristic of mathematics. A group which takes over such documentation should have the intellectual tools required to check its validity; they also should have the right, or rather the responsibility, to reject a project that fails to meet an adequate standard. Cases of dispute would be resolved by appeal to the line technical manager, who should be experienced and capable of resolving the dispute in a technically sound fashion. It is very unfortunate that many heads of data processing departments are promoted for achievements in accountancy, sales, or electronic engineering. They have little understanding of the nature of computer programming, and even less of the logical and mathematical techniques required for its control. It is the managers who could benefit most from the new disciplines: perhaps that is why they are sometimes most resistant to change.

Reliability.

In principle, we should find it much easier than other professional engineers to achieve the highest standards of quality, accuracy and predictability of timescale and cost, because the raw materials with which we work are much simpler and more plentiful, and much more reliable. Our raw materials are the binary digits in the stores and registers, disks and tapes of our computers. Our problem is that we have too many of them rather than too few. These bits are manipulated exactly in accordance with our instructions, at a rate of millions of operations per second for many weeks or months without mistake; when the hardware does go wrong, it is the engineer, not the programmer, who is called upon to mend it.

> "In order to use machines either to aid research or to aid teaching, the results, methods, and spirit of formalisation in mathematical logic are to play an essential role."
>
> Hao Wang, 1967.

That is why computer programming should be the most reliable of all professional disciplines. We do not have to worry about problems of faulty castings, defective components, careless labourers, storms, earthquakes or other natural hazards; we are not concerned with friction or wear or metal fatigue. Our only problems are those we make for ourselves and our colleagues by our overambition or carelessness, by our failure to recognise the mathematical and theoretical foundations of programming, and our failure to base our professional practice upon them.

Yet in some ways the engineers have an advantage over us. Because they are dealing with continuously varying quantities like distance, temperature, and voltage, it is possible for them to increase confidence in the reliability of an engineering product by testing it at the extremes of its intended operating range, for example, by exposure to heat and cold, or by voltage margins. We do the same in program testing, but in our case it is futile. Firstly we have to deal with impossibly many more variables; and secondly these variables take discrete values, for which interpolation and extrapolation are wholly invalid. The fact that a program works for value zero and value 65 535 gives no confidence that it will work for any of the values in between, unless this fact is proved by logical reasoning based on the very text of the program itself. But if this logical reasoning is correct, then there was no need for the test in the first place. That is why it is an essential prerequisite to the improvement of our professional practices that we learn to reason effectively about our programs, to prove their correctness before we write them, so that we know that they will not only pass all their tests, but will go on working correctly forever after

Structure.

Other engineers have a further advantage over programmers. When they split a complex design into a number of component parts, to be designed independently of each other, they can take advantage of the spatial separation of the parts to ensure that there can be no unexpected interaction effects. If the parts are wholly

unconnected, this is very easy to check by simple visual inspection. Thus when we turn our car to the left, we may be very confident that this will have no direct effect on the cigarette-lighter, the rear mirror, or the carburettor. When such interaction effects do occur, they are recognised as the most difficult to trace and eliminate.

But in the programming of conventional computers, there is no similar concept of spatial separation. Any instruction in a binary computer program can modify any location in the store of the computer, including those that contain instructions. And if this happens incorrectly only once in a thousand million instructions executed, the consequences for the whole program will be totally unpredictable and uncontrollable. There is no hope that a prior visual inspection of the binary content of store will enable us to check that such interaction cannot occur, or to find the cause of its occurrence afterwards. There is no structure or isolation of components in a binary computer program, other than that which has been carefully designed into it from the start, and maintained by the most rigorous discipline throughout implementation.

> "So then always that knowledge is worthiest.. which considereth the simple forms or differences of things, which are few in number, and the degrees and coordinations whereof make all this variety."
>
> Francis Bacon.

In spite of this, the programmer is often asked to include some feature into his program as an afterthought; and the only quick way to do this is to insert new instructions which cross all the boundaries between the carefully isolated components, and violate all the structural assumptions on which the original design was based. It would be repugnant to an engineer to introduce direct cross coupling effects between the steering and carburettor of a motor car, or the tapedecks and floating point unit of a computer. A programmer is all too willing to do his best, and his profession gets a bad name when unpredicted side effects occur.

A partial solution to this problem lies in use of a high level language like ALGOL 60 with secure rules governing the scope, locality, and types of variables. In such a language the programmer can declare the structure of his program and data, stating which groups of variables are to be accessed or changed by which parts of his program. An automatic compiler can then check that the appropriate disciplines have been observed throughout the whole of a large program, and can therefore give the same confidence to the programmer as the engineer gains by spatial separation of his components. Further confidence can be gained by running

the program on a machine like the Burroughs 5500 which makes similar checks while the program is running. In better established engineering disciplines, the observance of such elementary safety precautions has long been enforced by legislation. It is the law that dictates the measures that prevent unwanted interaction effects between an industrial machine and the body of its operator.

Tools of the Trade.

This brings me to the final disadvantage suffered by the programmer, the poor quality of the tools of his trade. I refer to his programming languages, operating systems, utility programs, library subroutines, all of which are supplied in profusion by the manufacturer of his computer. Many of these are so complicated that mastery of them absorbs all his intellectual efforts, leaving him little energy to apply to his client's original problem. Some operating systems are so poorly designed that they require twenty reissues (or "releases"), spread over a decade, before the original design faults have been rendered tolerable. And they are so unreliable that each issue has a thousand faults corrected by the next issue, which introduces a thousand new faults of its own. When finally the agony of reissues comes to an end, instead

of rejoicing, the poor programmer is cajoled or forced to accept an early issue of some "new" product. Such complexity, unreliability, and instability of basic tools were doubtless endured by engineers of each newly emergent discipline; but gradually the engineers developed better tool kits for their own use. That is a task which still faces the programming profession today -- the design of programming tools which are reliable, stable, convenient, and above all simple to understand, to control, and to use.

A crude measure of the simplicity of an engineering tool is the length of the manual required to give a full and complete account of how to use it and avoid

misusing it. At present our software manuals are both voluminous and inadequate. I believe that a solution to our problems can be sought in the design of software which can be completely described by shorter manuals. If an electronic engineer finds a method of satisfying with twenty components a need which has hitherto required thirty, the value of his discovery is immediately recognised and is often highly rewarded by fame or by money. When a software engineer designs a product that can be fully defined in twenty pages of manual, when the rival product has been inadequately defined in a hundred, his achievement is just as great, and possibly more beneficial; for he has achieved an economy in our scarcest resource -- not silicon or even gold, but our own precious human intellect.

How do we get there from here?

My description of the professional achievement of programmers of the future may seem to be nothing but an academic dream -- a pleasant one for our clients, but perhaps something more like a nightmare for us. How ever are we going to make such a fantastic improvement in our working methods? We are like the barber-surgeons of earlier ages, who pride themselves on the sharpness of their knives, and the speed with which they can dispatch their duties, either of shaving a beard or amputation of a limb. Imagine the dismay with which they greeted some ivory-towered academic who told them that the practice of surgery should be based on a long and detailed study of human anatomy, on familiarity with surgical procedures pioneered by great doctors of the past, and that it should be carried out only in a strictly controlled bug-free environment, far removed from the hair and dust of the normal barber's shop. Even if they accepted the validity and necessity for these improvements, how are they ever to achieve them? How could they re-educate all those hairdressers in the essential foundations of surgery? Clearly, a two-week course in Structured Surgery is all that we can readily afford. But more is needed, much more.

First we need good books, which can be studied by programmers and programming teams to familiarise themselves with the concepts of mathematical proof, and show how proof methods may be applied to the everyday practice of program specification, design, and implementation. Such books are beginning to appear in the publishers' lists. May I recommend the series edited by David Gries and published by Springer? May I even advertise the series edited by myself and published by Prentice Hall International?

Then we need a journal in which practising programmers can read the results of ongoing research, to keep themselves up to date with the most effective technology.

The Design of Well-Structured and Correct Programs
Alagic and Arbib. Springer-Verlag. 1978.

Programming Methodology
ed. D. Gries. Springer-Verlag. 1978.

Software Development: a Rigorous Approach
C.B. Jones. Prentice-Hall. 1980.

Structured Systems Programming
J. Welsh and R.M. McKeag. Prentice-Hall. 1980.

The Science of Computer Programming
D. Gries. Springer-Verlag. 1981.

Some recent books on professional aspects of programming.

A new journal of this kind has just been founded. It is called 'Science of Computer Programming', edited by Michel Sintzoff, and published by North Holland. I have high hopes for it.

Most of the books and articles on programming methods are of necessity illustrated only by small examples. Indeed, many of the programming methods advocated by the authors have never yet been applied to large programs. This is not a defect of their research, it is a necessity. All advances in engineering are tested first on small-scale models, in wave tanks, or in wind tunnels. Without models, the research would be prohibitively expensive, and progress would be correspondingly slow.

Nevertheless, I believe that the time has come to attempt to scale up the use of formal mathematical methods to industrial application. This can best be achieved by collaborative development projects between a university or polytechnic and an industrial company or software house. Such a project might be an entirely new program, or it might be a restructuring or redesign of some existing software product in current use, perhaps one which has lost its original structure as a result of constant amendment and enhancement. The great advantage of these joint projects is that they bring home to academic researchers some of the exigencies of working on much larger programs; and they give a practical training in formal methods to larger numbers of experienced programmers in industry. This is technology transfer in its best sense -- a transfer of benefits in both directions.

Education.

As I have emphasised already, the major factor in the wider propagation of professional methods is education, an education which conveys a broad and deep understanding of theoretical principles as well as their practical application, an education such as can be offered by our universities and polytechnics. Lecturers and professors regard it as their duty and privilege to keep abreast with the latest developments in their subjects, and to adapt, improve and expand their courses to pass on their understanding to their students. Many entrants to Computer Science courses have acquired a familiarity with the basic mechanics of progamming at their schools; and at university they are ready to absorb the underlying mathematical principles, which will help them to control the complexity of their designs and the reliability of their implementations.

Over the next decades, while the graduates of Computer Science courses are entering their profession, we will have an extremely awkward period, in which almost none of the senior professionals and managers will have any knowledge or understanding of the new methods, while those whom they recruit will seem to them to be talking academic gibberish. This could be a grave hindrance to the development of our profession. Furthermore, it would be a terrible wasted opportunity, because one of the major benefits of the technique of mathematical abstraction is that it enables a chief programmer or manager to exert real technical control over his teams, without delving into the morass of technical detail with which his programmers are often tempted to overwhelm him.

The solution to this problem is for the ambitious senior programmers of the present day to make the effort now to gain the necessary mastery of the subject, and so ensure that they will become in future the effective chief programmers, technical managers, and technical directors of their companies and institutions.

One way of acquiring a professional reorientation of this kind is to take a specialist postgraduate post-experience course in a new and important subject. Thus an electronic engineer might now be going back to university to study VLSI design; or an industrial chemist might be taking a Master's course in polymer science or genetic engineering, offered by some forward-looking university or polytechnic. I believe that ambitious programmers should not be reluctant to follow the example of the well established engineering disciplines. That is why at Oxford University we have instituted a new MSc course in Computation, devoted primarily to the objective of improving programming methods and ensuring their wider application. A similar course is offered at the Wang Institute in the U.S.A.

Conclusion.

In 1828, on the occasion of the grant of Royal Charter to the Institution of Civil Engineers, Thomas Tredgold defined civil engineering as "the art of directing the great sources of power in Nature for the use and convenience of man." Many branches of engineering have been established since that date. They have all been concerned with the capture, storage and transformation of energy, or with the processing, shaping and assembly of materials. Computer programmers work with neither energy nor materials, but with a more intangible concept. We are concerned with the capture, storage, and processing of information. When the nature of our activities is more widely understood, both within our profession and outside, then we shall be deservedly recognised and respected as a branch of engineering. And I believe that in our branch of engineering, above all others, the academic ideals of rigour and elegance will pay the highest dividends in practical terms of reducing costs, increasing performance, and in directing the great sources of computational power on the surface of a silicon chip to the use and convenience of man.

> "It has long been my personal view that the separation of practical and theoretical work is artificial and injurious. Much of the practical work done in computing, both in software and in hardware design, is unsound and clumsy because the people who do it do not have any clear understanding of the fundamental principles underlying their work. Most of the abstract mathematical and theoretical work is sterile because it has no point of contact with real computing. One of the central aims of the Programming Research Group as a teaching and research group has been to set up an atmosphere in which this separation cannot happen..."
>
> Christopher Strachey, 1974.

OXFORD UNIVERSITY COMPUTING LABORATORY
PROGRAMMING RESEARCH GROUP TECHNICAL MONOGRAPHS

MAY 1982

This is a series of technical monographs on topics in the field of computation. Copies may be obtained from the Programming Research Group, (Technical Monographs), 45 Banbury Road, Oxford, OX2 6PE, England. Prices include surface postage.

PRG-2 Dana Scott
Outline of a Mathematical Theory of Computation

PRG-3 Dana Scott
The Lattice of Flow Diagrams

PRG-5 Dana Scott
Data Types as Lattices

PRG-6 Dana Scott and Christopher Strachey
Toward a Mathematical Semantics for Computer Languages

PRG-7 Dana Scott
Continuous Lattices

PRG-8 Joseph Stoy and Christopher Strachey
OS6 – an Experimental Operating System for a Small Computer

PRG-9 Christopher Strachey and Joseph Stoy
The Text of OSPub

PRG-10 Christopher Strachey
The Varieties of Programming Language

PRG-11 Christopher Strachey and Christopher P. Wadsworth
Continuations. A Mathematical Semantics for Handling Full Jumps

PRG-12 Peter Mosses
The Mathematical Semantics of Algol 60

PRG-13 Robert Milne
The Formal Semantics of Computer Languages and their Implementations

PRG-14 Shan S. Kuo, Michael H. Linck and Sohrab Saadat
A Guide to Communicating Sequential Processes

PRG-15 Joseph Stoy
The Congruence of Two Programming Language Definitions

PRG-16 C. A. R. Hoare, S. D. Brookes and A. W. Roscoe
A Theory of Communicating Sequential Processes

PRG-17 Andrew P. Black
Report on the Programming Notation 3R

PRG-18 Elizabeth Fielding
The Specification of Abstract Mappings and their implementation as B^+-trees

PRG-19 Dana Scott
Lectures on a Mathematical Theory of Computation

PRG-20 Zhou Chao Chen and C. A. R. Hoare
Partial Correctness of Communicating Processes and Protocols

PRG-21 Bernard Sufrin
Formal Specification of a Display Editor

PRG-22 C. A. R. Hoare
A Model for Communicating Sequential Processes

PRG-23 C. A. R. Hoare
A Calculus of Total Correctness for Communicating Processes

PRG-24 Bernard Sufrin
Reading Formal Specifications

PRG-25 C. B. Jones
Development Methods for Computer Programs including a Notion of Interference

PRG-26 Zhou Chao Chen
The Consistency of the Calculus of Total Correctness for Communicating Processes

PRG-27 C. A. R. Hoare
Programming is an Engineering Profession