A Working Party under the Chairmanship of Dr. J. L. Alty was set up by the Computer Board in October 1977 with the following terms of reference:

> "to report to the Computer Board (by January 1978 if possible) on the effect in the next ten years of microprocessors and other microtechnology on the provision of university computing services (especially as regards computer centres and their communications)".

The following were invited to serve on the Working Party and agreed to do so:-

> Professor D. Aspinall,
> Mr. I. M. Barron,
> Mr. M. B. Williams,
> Dr. B. Zacharov.

The Working Party has met on four occasions and there has been extensive communication by correspondence. Because of the difficulty of arranging meetings in a short time scale it was not possible to report to the Board by January 1978.

J. L. ALTY
8 November, 1978.

Section A The University Computing Scene

A1 At present local and remote batch computing requirements in Universities are reasonably well provided for, but interactive provision for both teaching and research varies markedly between universities. It is the Working Party's opinion that existing interactive facilities are inadequate and that the Computer Board will fail to meet its new objectives for providing facilities for teaching and research by using existing hardware.

A2 The Working Party expects that there will be a shift in emphasis towards more interactive working and the on-line control of experiments and data collection. At present such a trend is already in evidence in individual University departments and there is considerable user pressure for such a change. The Science Research Council Interactive Facility provides an example of this.

A3 In like manner the pressure for more interactive teaching systems will continue to increase. Departments have found such techniques very beneficial, and the Working Party is convinced that the need for simple, but effective interactive teaching systems will increase significantly, and that more emphasis will be required on this teaching remit of the Board in the future.

A4 The University computing population is at present dominated by the Engineering and Science disciplines, though there is now an increasing trend for research workers and undergraduates in the Social Science and Medical faculties to utilise computers in their work. The working party expects this trend to accelerate (as has already happened in the USA). Such a trend will bring an increasing number of naive computer users in contact with the computing machinery, and successful exploitation will depend upon the simplicity and effectiveness of the user-computer interface.

Section B The Approach

B1 The Working Party decided to seek the views of a number of organisations on the likely impact of the new technology. Four of the major computer manufacturers and one semi-conductor manufacturer were consulted and agreed to provide information on the understanding that all such information would be treated as "commercially in confidence". This approach enabled us to have full and frank discussions, though it obviously meant that all the information presented to us could not be included in this report. Information was sought from individuals, and the Working Party received a number of personal contributions from Universities. Additionally the opinions of the UK Scientific Advisor in Tokyo, the Computer Services Association and Universities in the USA were canvassed.

Finally the members of the Working Party themselves drew upon their many formal and informal contacts to provide additional input to our discussions.

B2 The Working Party decided early in its investigations that any statements concerning the impact of the new technology would need to be general and tentative in view of the rapid advances taking place. It therefore decided to issue an interim report to the Computer Board which would concentrate upon general issues and trends.

B3 In view of the importance of feedback of information in such an exercise as this, the Working Party decided that its interim report ought to be widely circulated for comment and that any views expressed as a result of such action should be thoroughly taken into account before any final report was issued.

Section C The New Environment

C1 Any discussions of the University computing environment must look in turn at peripheral costs, software costs, maintenance issues and communication costs as well as the obvious effects of the new technology - the dramatic fall in the cost of storage and processing power. Whilst the latter have been repeatedly emphasised in the press, and highlighted by the rapid growth in personal computing, they cannot be viewed in isolation, and it is the Working Party's view that too many pronouncements have recently been made which do not fully take into account all aspects of the computing environment. Statements such as "a 370/168 on a chip" and "the new technology will render computer centres and regional centres obsolete" are somewhat premature.

C2 That processing costs are falling is not in doubt. A consensus view of the organisations who talked with us was that processing costs would fall at a rate of about 33% per annum, providing an order of magnitude reduction every five years. Main storage costs are expected to fall at about the same rate.

C3 Another interesting and relevant parameter is the number of elements on a chip. This parameter limits the capability set of a microprocessor. For example, a typical minicomputer at the present time has about 4-6K of elements within its basic logic. In 1972 the new technology could only support about 700 elements per chip. By 1978 this had risen to about 10K elements per chip, and this figure is expected to rise at a staggering rate of 50% per annum. Such technology will be able to support a very powerful instruction set, and the concept of a 370/168 processor on a single chip is certainly feasible in the near future.

C4 In many cases the cost reductions will manifest themselves in a reduction of cost per function rather than in unit cost. For example, although we do not expect terminal prices to fall significantly, their functional capability will be much greater.

C5 Most existing peripheral devices, magnetic tapes and disks, printers and readers are mainly electro-mechanical in their operation. Thus one would not expect, and indeed has not seen, any dramatic effects on the cost/performance ratio of such devices arising out of the new technology. Rather, there has been an upsurge in new devices aimed at particular requirements such as floppy disks, cassettes and small printers which have reduced the minimum capital required for peripherals considerably. Yet the cost of such devices is still considerable (when compared to the processor and storage costs) and it is unlikely to fall further by any dramatic amount. For example, floppy disks and printers often cost nearly an order of magnitude more than the processor which supports them and it is clear that over the next few years the price of a personal computer will be dominated by its peripherals, the processor and storage being a relatively insignificant part of the total cost.

C6 The floppy disk has reduced the entry point for use of disks on small systems but it has not offered a better price-per-byte-ratio than is currently offered by a 200M byte disk. Mass Storage devices offer a price-per-byte ratio which the new technology will not be able to match for a number of years. Similarly the new small printers which are now available at relatively low prices offer considerably reduced performance when compared with their traditional mainframe counterparts.

C7 The microprocessor manufacturers are making some efforts to improve the quality of the software offered on their systems. At present most systems support an assembler, and a high level language such as PLM. However, fault finding and debugging are time consuming tasks on the present generation of microprocessors, a task which can only be eased by purchasing a considerable amount of support equipment such as logic stateanalysers or development systems. The overheads for software production are at present very high.

C8 The present 8-bit generation of microprocessors will soon by joined by their 16-bit counterparts, and early in the next decade we may possibly see the introduction of the 32-bit microprocessor. At this stage these devices will be able to support most of the functions currently restricted to large minis or mainframes. However each generation will require new investment in hardware and software testing equipment and a rewrite of software which ran on the previous generation.

C9 The microprocessor chip manufacturers need a huge market to sustain their high volume production lines. The present installed computer base represents a trivial fraction of this market. Chip manufacturers are not interested in usurping the present computing base. It represents one month of their manufacturing capability. Thus these manufacturers are looking for totally new markets to keep their production full. This is important since it means that one cannot assume that chip production will be followed up with full software support as happened for second and third generation computers. Such software development will probably be left to users and to software houses.

C10 There can be little doubt that software houses will fill the gap. However such software, although probably cheaper than existing software, will still be relatively expensive and such costs could easily outstrip all the hardware costs particularly for exotic application packages. Any in-house software developments will be very expensive in human effort.

C11 Low cost processing and RAM memory hardware will permit the proliferation of processing elements dedicated to particular constituent tasks. The possibility therefore exists of computing solutions based upon a multitude of separate interconnected elements.

In computing systems of intercoupled components, the probable pattern for the future will involve components from different manufacturers. Thus the problems of intercoupling and compatibility will assume a major importance. Work will be urgently needed to ensure that industry is prepared to exploit the possibilities of such inhomogeneous intercoupled systems.

Section D Computing in the 80s

D1 It is the Working Party's view that the new technology will have an impact both at the high and low ends of the Computer Market. Developments in personal computing and single-application chips will proceed alongside developments in the traditional data processing market. The former will be characterised by low costs and high volumes whilst the latter will involve high costs and relatively low volumes.

D2 There will be a proliferation of low cost, high volume systems which will be used in new application areas. Such processors will normally be programmed in Read-only-memory (ROM) and will therefore be unalterable. Such chips will normally have a single purpose and the software will be relatively simple. The production of such chips will be measured in millions and the end user will often not realise their existence as a computing element. We expect that the majority of such devices will be dedicated to single applications.

D3 The new technology will be rapidly utilised in existing mini computers. Where applications involve few peripherals (i.e. process or network control) prices will drop markedly. Such systems will normally have access to further data processing capability via some form of computer network.

D4 The entry point for the interactive mini computer will remain high, at about £25K-£50K but such systems will provide a vastly superior service compared with present-day systems. For example, a 32 or 64 interactive terminal system may cost about £50K (excluding terminal costs). The Working Party expects performance to rise rather than costs to fall because the new technology will remove the traditional bottle-necks of interactive computing namely power and main storage, but will not reduce peripheral or software costs significantly. Additionally, performance will take on the new character of usability.

D5 The new logic will be utilised by the large computer manufacturers and will lead, in the long term, to significant improvements in price/performance. Such manufacturers will endeavour to hold price levels as far as possible but the impact of the high performance minicomputer systems will force prices down. Usability however will assume a key role and will justify higher costs.

D6 The real problem in utilising the new technology in the traditional mainframes will be the present length of the design cycle which is typically 3-4 years. This cycle will have to be reduced to 1-2 years in order to be competitive. We know that many of the present large mainframe manufacturers are aware of this problem and are making strenuous efforts to overcome it.

D7 The initial effects of the new technology are already apparent. An examination of the IBM 3030 series for example reveals that Grosch's law no longer holds - the relationship between power and price is now linear - and price/performance improvements of about 30% are being realised compared with the 370 series. This has the obvious implication that the use of even larger computers to process increasing volumes of simple jobs may not be the best strategy for the future. However, increasing emphasis on usability will make large systems more attractive, even though their throughput cost/performance may not be as good.

D8 Special purpose engines such as the Distributed Array Processor, or Data Base Access Facilities, will become available. Although these will be relatively expensive one might expect a really powerful engine with say 50-100 times the power of a CDC 7600 at a comparable cost in the future.

D9 Thus in the next decade the Working Party expects the impact of the new technology to be felt throughout the range of computing applications, from the personal computer market to the parallel vector processors. By 1985 we might expect capabilities and prices to be distributed as in Table 1.

Specials (dedicated)	£10
32-bit 64 kbyte processor (no peripherals)	£200
Mini: (100-200 kips)	£5,000
Office Processor (1-20 Mips)	£10,000
Interactive Mini (64 terminals)	£25,000
Large Processor (50 Mips, 50 Mbyte store)	£200,000
Super processor (1000Mips) with 500 Mbyte store, 1 Gbyte disk	£4,000,000

Table 1 A possible spectrum of systems available in 1985

D10 Software production costs like hardware costs will depend upon the market volume. There will be a similar High Cost/Low Volume and Low Cost/High Volume split. The high volume software will reside in the single application chips often in ROM. It will be simple, reliable and unalterable. The low volume software will be similar to the software of today, - complex, multipurpose and probably as unreliable. High Volume software will be unsupported.

D11 Many software houses will attempt to break into the micro-software market and we expect that much hardware and software effort will be based upon existing architectures particularly those of IBM and DEC. The current size for such compatible software will render such an approach very attractive.

D12 The new technology will enable peripheral devices to be interfaced easily to a variety of systems and the traditional restrictions on the range of input/output devices that can be connected to existing mainframes will be much reduced. Again the large existing IBM and DEC mainframe and minicomputer base will make the early adaption of peripheral interfaces to such systems attractive.

D13 Because of the inherent cheapness of the hardware there will be a temptation for users to embark on system design and write their own software. This approach will be fallible and very costly particularly in the early part of the decade. Each new generation of microprocessor will require its own testing aids, such as simulators, emulators and logic state analysers. The investigation costs would be very high, and equipment bought may well be obsolescent on delivery.

D14 Because processing and storage will be far cheaper than peripherals the Working Party expect to see rapid developments in dispersed computing which will rely on networks for peripheral and secondary storage support. The hardware for network interface processors and switches should become very cheap, and will encourage this trend, but their success will depend upon the availability and adoption of standards.

D15 The availability of relatively cheap storage and processing will enable software to be developed which is sympathetic to the naive user. New software techniques will be required to make this possible.

D16 The division between software and hardware will become diffuse and will shift. There will be a need for changes in the education of programmers to encompass more hardware aspects.

D17 The traditional large computer manufacturer will aim at broadening its market and will cease to be interested in catering for specialised, small sectors of the market where each order presents unique problems. These manufacturers will not be interested in the super-powerful machine for large scientific calculations and will be unwilling to offer any other than standard products (both as regards hardware and software). Competition for defined market ranges will be severe.

D18 A small number of firms will concentrate on particular specialised requirements, including military/scientific needs for exceptionally-powerful machines. U.K. firms are unlikely to offer comparable products, and foreign purchase without full competition, may be inescapable.

D19 A large number of suppliers of small microcomputer systems will arrive in the market place. They will not offer support for systems design, and will be largely incompatible with one another in both hardware and software.

Section E University Computing in the Next Decade

E1 The development and exploitation of the new technology must be seen in the context of University funding. Universities obtain funds for computing from three major sources - the Computer Board, the University Grants Committee and the Research Councils. Board money is normally used to enhance central facilities for the provision of a general service, UGC money is provided to departments for equipment grants, and Research Councils grants are normally only given for specific research projects.

E2 Because of the variety of funding sources and the autonomy and individuality of the various departments in a University, the growth in usage of the new technology will be haphazard and lack central control. Furthermore because funding tends to be spent on equipment rather than recurrent needs (since all funding bodies try to avoid long term commitments) the Working Party believe that there will be a tendency to see the new technology as solving a capital expenditure problem and that too little thought will be given to the longer term recurrent implications.

E3 As a result of this emphasis on capital expenditure individual university departments will often not realise the full financial implications of software development and maintenance. Indeed there is already evidence that such an approach is being adopted at some sites. The absence of properly trained personnel together with a lack of software and hardware testing aids will retard development and place heavy demands upon the time of research workers which could be better spent.

E4 Users and departments will identify tasks that can be carried out on independent local small machines as an alternative to sharing the resources of a large central system. In some cases this view will be justified because a sufficiently well-defined market will have been identified for a supplier to offer machines complete with specific applications programs and data storage to function as independent entities. In other cases a true stand-alone machine may not be feasible or efficient because access will be required to remote files, additional computing power and a wide range of applications programs or special peripherals.

E5 Departments will have to be particularly careful in comparing costs. An independent machine may require significant support by skilled personnel in program or data preparation or in users' advisory services. These needs, and their cost, are not taken into account in a simple comparison of machine cost against the costs of using a well-supported central computing service.

E6 In the University computing environment there are many thousands of files in various stages of development. Such a large collection needs careful maintenance and housekeeping. The Working Party viewed the possibility of a large number of user files spread over many independent departmental systems with some concern - not only from a security point of view but also from simple economics. It is the view of the Working Party that the maintenance and protection of the majority of users files will best be carried out in a central facility. This conclusion is an example of the importance of usability which has been mentioned earlier.

E7 It is expected that the well-established types of mini-computers will continue to be marketed as complete systems at improving cost/performance ratios but with the range of performance and facilities extended upwards and downwards. As the cost of the processor will fall much more rapdily than the cost of terminals or external storage, these items will effectively determine the cost of a machine. An opportunity will therefore exist for significantly improving the interactive facilities in Universities via a series of departmental, or faculty systems, and connection to a central filestore would enable the full potential of such systems to be realised.

E8 Thus, departmental and personal machines should be evaluated in the context of distributed processing and should conform to appropriate standards for communication with other computing and storage services. It is therefore essential that the central University computing service should be consulted when any equipment or software is being purchased.

E9 One area where the new technology is liable to have an immediate and economic impact is that of data collection and on-line control of experiments. Since in such applications the main needs are for processing power and providing storage, departments should be able to realise real financial savings provided the required software effort is available. Such software effort will often have to be carried out by departmental personnel, but since, by its very nature, this software is of the one-off variety, such effort will be worthwhile. Again communication with a larger central filestore could be of significant advantage.

E10 The Working Party was concerned about the amount of time and effort that could be spent, particularly by inexperienced users on microprocessor software development. It was felt that departmental users should have access to both hardware and software development aids and educational facilities. Since the environment is developing rapidly such aids would be via some central pool, which at the same time could act as a centre for expertise and training.

E11 Developments in the Universities will be retarded if departments concentrate on low level language software development. The experience on mainframes over the past decade points to the use, as far as possible, of high level programming languages. If cross compilers are provided on the University central system departmental users would have access to powerful development facilities. Such an approach however implies a certain degree of central control in deciding which microprocessors and high level languages should be supported.

E12 The Working Party identified four distinct areas requiring training facilities

- in-depth hardware education
- software training in both high and low level languages
- minimal hardware familiarity necessary for the writing of software
- general appreciation courses.

Additionally the expertise of the University ought to be utilised in an extra-mural manner for local industry. Such training could cover all four categories outlined above.

E13 Because of the varied nature of University departments the Working Party would not wish to lay down stringent guidelines as to which departments in the University ought to coordinate the courses outlined above in E12. As a general guide however it was felt that courses of an academic (in-depth) nature would be the responsibility of Electrical or Control Engineering departments, whereas service courses and general appreciation courses would be given by the Computer Centre. A co-ordinated approach between the above departments would be ideal.

E14 Many departmental requirements will be met by a "black box" approach whereby a cabinet containing a processor store and input/output interfaces for a range of peripherals would be configured to user requirements. Such a box would need to be programmed but a minimal amount of hardware expertise would be needed by the user. The Working Party felt that such a service might best be provided by a service orientated organisation such as the Computer Centre.

E15 On the other hand the Working Party recognised that some user departments will require in-depth hardware knowledge in order to fully exploit the new technology. Such users will need to be fully trained in hardware design probably by the Electrical or Control Engineering department.

E16 Thus the Working Party expect that the role of the University computer centre will change as the new technology enables computer power to be spread round the campus. Instead of being the main, or sole, supplier of computing power the centre will increasingly concentrate on co-ordinating the computing resources available to the users whether from on-site or remote sources. It will have important functions as custodian of central filestores as the gateway to external resources and as the focal point for a campus network. Its personnel will increasingly be concerned with

- local advisory services
- retailing external resource centres and facilities
- general support on applications programs
- communications and network management
- coordination of campus computing
- management of central file storage and special peripheral machines.

E17 The skills required by Computer Centre staff will shift gradually from large scale operating system maintenance and operation to a support and educative function dealing with a very wide range of user expertise.

E18 At the larger end of the market the new logic will enable really large processors to be developed at modest cost. Such processors will be both serial and parallel in nature. The Working Party envisages that one or more really large scale centres for such computing will be established and will be accessed over national networks.

E19 Similarly the new logic, as applied to secondary storage controllers will enable really powerful data base engines to be produced. Such centres will be few in number and will be accessed over a national communications network.

E20 No diminution in the role of national centres is foreseen in the provision of computing facilities to deal with large jobs (i.e. those requiring massive computing power) on the manipulation of large data bases. Furthermore national facilities will also still be required to satisfy the demands for special and costly facilities which cannot reasonably be provided locally.

E21 Communications costs are likely to become less sensitive to distance so that the cost of access to a national centre may not vary significantly for alternative locations. Although the relative costs of small and large machines is expected to favour the use of local machines for small and medium size jobs, a true comparison of the cost of local or centralised computing as alternative ways of meeting a particular university's needs will require detailed and continuing study as the balance of costs alters. Significant changes are likely in the relative costs of machine, applications support and communications.

E22 Dispersed computing, as envisaged above, will demand that a network architecture be defined and adopted and that standards be laid down for command language and for transfer protocols. The heterogeneous nature of university computing is certain to continue since, unlike commercial or industrial organisations, there is no identifiable corporate requirement against which a single manufacturer could construct an integrated system. Specialised systems made up from elements and modules are, however, likely to be a significant feature in computer manufacturers' market strategies. Manufacturers' systems and procedures may be incompatible one with another; until the BSI/DPS20 concept of "open-system interconnection" is adopted generally it may be necessary for the university community to impose its own standards.

E23 In summary therefore, the Working Party sees an emphasis in the new decade on networking (local and national), and a requirement for continued user support in both hardware and software. What will be needed is a mechanism for ensuring that such an environment exists and is properly supported.

Section F Recommendations

The main issues that have been highlighted in the previous sections are:

a) The likely impact of a continued fall in the cost/performance of processors, and the dominance of peripherals and support costs.

b) An inevitable growth in dispersed computing in the University environment; with a continuing scope for large and special facilities to be available nationally.

c) Existing funding mechanisms for University computing will encourage a proliferation of small machines in an uncoordinated way.

d) A lack of expertise and training facilities concerning the new technology in the University environment.

e) A need for both academic and service teaching.

f) A requirement for supporting facilities probably provided from a central unit.

g) The importance of campus and national networks.

h) The use of computers in universities will change rapidly and the customer base will extend rapdily to areas outside science and engineering.

Provided that difficulties outlined above are satisfactorily resolved the Working Party believes that the impact of the new technology will be beneficial to University research and teaching. What is needed is an approach which provides the proper environment for its exploitation.

Whilst the Working Party has definite views about the way in which such issues should be resolved, it would welcome input from Universities and other bodies concerning the issues raised. The following recommendations should therefore be regarded as in no way immutable. They are there to stimulate thought and encourage discussion. As a result of such discussions the Working Party would then hope to issue a report with firm recommendations.

F1 The Computer Board should recognise that dispersed computing will shortly become common place in University computing. It should therefore take responsibility for the provision and maintenance of campus networks whose focus will normally be the University computing centre.

F2 In order to avoid the duplication of effort which has occurred in regional networking the Board should take responsibility for the design and implementation of a standard campus networking system. Such a system ought to be the result of a joint design study between University personnel and members of the Network Unit and should be based upon the open network principle. Whenever possible "off-the-shelf" solutions should be adopted.

F3 The campus networks should be able to attach a variety of minicomputers, microcomputers and terminals and be able to connect to a variety of mainframes. Each local site will need to provide the effort to establish connection to the local mainframe at the University Computer Centre, but such effort ought to be coordinated by machine type over a number of universities.

F4 The campus network should provide for efficient file transfers between systems, and be able simultaneously to carry interactive loads to and from the mainframe, and between other systems on the campus.

F5 The interconnection between campus networks and national networks should be such that a local campus user can easily obtain access to national facilities.

F6 The campus network should provide connection for Research Council systems which are situated in University departments since in the future there could well be a division of responsibility for service between say an SRC provided microprocessor and a Board provided central system.

F7 Universities should be encouraged to develop centres of microprocessor expertise which will provide advice, guidance, and education to user departments, for both academic and service needs. The Board should concentrate its efforts in the service area.

F8 At the expertise centres equipment should be provided to enable users to familiarise themselves with the new logic. Additionally extensive testing facilities (both software and hardware) should be provided.

F9 One possible way of achieving F7 and F8 would be via joint funding between the Board and a particular University. For example the Board could provide capital funding for testing equipment, compilers and familiarisation equipment provided that the University covers the staffing costs. A scheme to achieve this objective is outlined in Appendix A.

F10 One effect of the new technology will be to shift expenditure from the centre to the periphery. It is vital that the Board exercises some control over equipment connected via campus networks to its central or national systems. Whilst it is impracticable for the Board to be formally responsible for periphery developments, it should (via delegation to the local Computer Centre Director) insist on full consultation when equipment is being evaluated for departmental use.

F11 Universities have traditionally sought to solve their operational reports by the use of monolithic mainframes. In the new environment such requests could be met by a complex of processing elements each dedicated to a particular task yet working in concert to satisfy the conflicting needs of a University workload. Universities will need guidance on how to draw up such requirements and in evaluating technical solutions. The Working Party considered that there might be a role for an advisory body to ensure standardisation in the language of requirement definition, and to assist in technical evaluation.

F12 Every effort should be made to avoid excessive wasting of effort in the software field. A software coordinating committee (like the Satellite User Liaison Group) could ensure the dissemination of software information among Universities and minimise duplication of effort.

F13 The Board and the Research Councils should work closely together to maximise the exploitation of the new technology. Adoption of standards for programming, architecture and communications would be an important factor.