COMMISSION OF THE EUROPEAN COMMUNITIES
COMMISSION DES COMMUNAUTES EUROPEENNES

EMBARGO : VENDREDI 15/2/1991 A 12H15

Report of the EEC Working Group on High-Performance Computing

Mandate

Prof. Carlo Rubbia was encharged by the Commission of the European Communities to constitute and direct a Working Group with the mandate to prepare a recommendation on a European policy for high-performance computing. The main questions to be answered by this Working Group are:

- What is the present and foreseeable situation in Europe of the high-performance computing domain, both for users and suppliers?

- Is such a situation satisfactory, considering the impact of this domain for the future developments of science, economy and society in Europe and in the world?

- Should this situation not be fully satisfactory, which actions have to be recommended on a European level and which resources would these actions entail?

Working Group Members

***Prof. Carlo Rubbia** (Chairman)*
CERN

***Dr. Federico Carminati** (Secretary)*
CERN

Dr. Renzo Allaria
FIAT Auto

Prof. John Argyris
University of Stuitgart

Prof. Herman Berendsen
University of Groningen

Dr. Tor Bloch
Advanced Computer Research Institute

Dr. David Burridge
ECMWF

Dr. Jean-Marie Cadiou
Commission of the European Communities

Dr. Michel Carpentier
Commission of the European Communities

Prof. Jack Dongarra
Univ. of Tennessee

Dr. Herve Gallaire
BULL S.A.

Prof. I.H. Hillier
University of Manchester

Prof. Jacques-Louis Lions
CNES

Dr. Jean-François Omnes
Commission of the European Communities

Dr. Pierre Perrier
DGTDEA, AMDBA

Prof. F. Troyon
École Polytechnique Fédérale de Lausanne

Dr. Roberto Vio
FIAT

Prof. Paolo Zanella
CERN

Contents

Executive Summary

The field of computing is on the verge of a new revolution. During recent years the amount of computational power at the disposal of scientists and engineers has increased dramatically. This has enabled them to envisage approaches that will revolutionize all fields of science, as has already taken place in the structural analysis of astronautical and other more mundane engineering systems. A new philosophy of conceiving scientific theory is about to be born in the so-called computer laboratory which may be considered to stand half-way between theory and experiment. This new approach, in fact a novel understanding and a re-formulation of the basic principles governing the behaviour of physical, economic and social systems, provides ultimately in concert with powerful computing devices a potential to simulate reality to a degree of accuracy which was unthinkable until now. Scientists and engineers will soon be able to model reality and study it, generating new knowledge and opening the door of the future for science and technology. Complex and costly experimentations will be replaced to a great extent by computer models revealing the finest details beyond the limits of any real experiment. Imaginative solutions to most complex problems will be feasible. Researchers will be effectively free from material limitations in their exploration of reality, constrained only by their own imagination. As the industrial revolution was triggered off by the creation of steam-power, the limiting factor to the new revolution is the computing power that is available. **Scientific and societal progress, industrial competitiveness, the understanding and control of environmental factors necessary to human well-being will be governed by the availability of adequate computing power.** This will affect in particular those activities where both simulation and optimisation play a leading role. Chemical, pharmaceutical, aerospace, automotive industries, as well as physics and mathematical sciences are among the most promising examples. To develop the full potential of the new approach a thousandfold increase in the power of high-performance machines is needed by the end of the decade. This will lead within 7 to 10 years to new computer devices, capable of at least one Teraflops, that is 1000 billion arithmetical or logical operations per second, instead of a few Gigaflops (1 billion operations per second), which the best supercomputers are now able to deliver. As already indicated, such an increase may reasonably be considered within human inventive capacity.

This has been well recognised in the USA and Japan, where, notwithstanding the existence of a flourishing computing industry, governments are investing huge resources in high-performance computing bringing together

manufacturers, universities, users and other relevant participants. The situation created by the prospects of quantum-leap changes in the computing power availability causes one to reflect on the European situation and on the crucial urgent decisions which have to be taken now. In the face of these massive strategical actions in other countries and in spite of the technological success of many European information technology programmes, European supply is totally absent even though Europe represents 30% of the world market for high-performance computing. **The results of this startling situation is that Europe's economical and industrial future increasingly depends on foreign suppliers if this is perpetuated.** This situation puts at risk the ability of the whole European industry to seize the advantages and opportunities offered by this new revolution which will dramatically reduce R&D costs by replacing expensive and sometimes dangerous experiments and substantially shorten the time for product development.

To overcome this unacceptable situation, a substantial investment in the field of high-performance computing is needed in order for Europe to become, by the end of the decade, an active participant and a recognised partner at the level of the leading countries. A global action plan composed of five inter-related actions is proposed. The creation of a competitive European high-performance computing industry in software, architecture and hardware should be pursued vigorously. The knowledgeable and influential community of European users of high-performance computing machines must be organized. An advanced and reliable Pan European High Speed Network should link the major European computing centres giving user access to the most advanced services. Education and training should be strengthened and the imaginative use of these tremendously productive machines and systems should be encouraged. New investments should be gradually introduced during a first five year period, matching the growth of the European ability to make profitable use of the available resources. Around 1995, various sources (Community, national, industrial, business and telecoms) should aim to contribute to a total annual investment of the order of 1000 MECUs. **A close collaboration between industry and users in Europe is the key to the success of the programme, but it requires a corresponding commitment to continuity in action.** The success will be measured by the increased presence of Europe in all key areas of high-performance computing, the emergence of many new enterprises in the field and above all it will enable European industry to maintain and improve its competitive position at the start of the next century.

Applications of high-performance computing

Computer modelling of a large variety of physical systems has reached the point where **reality can now be simulated** with a high degree of reliability. Physical models of real systems, be it the atmosphere, turbulence, chaos, combustion in chemical/mechanical systems, aerospace and automotive vehicles, engines, protein molecules, factory operations or economics, can be sufficiently detailed to be used for realistic predictions. Therefore high-performance computing is more and more instrumental to the development of society and to the competitiveness of industry at large, and not any more limited to a very specific industrial segment.

Simulation using high-performance computer systems is increasingly replacing experiments and scale modelling in a wide range of scientific, engineering and commercial applications. In some **industrial areas**, such as in the design of aeroplanes, space vehicles, ships and cars this is already a inescapable necessity. In others, scientists and engineers are on the verge of large-scale breakthroughs in the design of drug molecules, enzymes, catalysts and new materials. The time elapsed from the conception of a new product to its introduction on the market and the cost of its design process can be dramatically reduced with the help of computer simulation. This, combined with the accuracy reached in reproducing reality, has the effect to increase drastically the competitiveness of industry by reducing costs and improving quality. At the same time there are areas where computers must be used for establishing the most efficient physical simulation methods and models.

Optimisation is similar to simulation, insofar as it depends on the availability of accurate mathematical models of reality and powerful computers. Optimisation is a rather new application for high-performance computers and it is becoming increasingly important in the operations of large industries and service-providing companies. The solution to large optimisation problems is often time critical: decisions must be taken in real time to respond to a fast evolving situation. One of the main limitations to a wider usage of optimisation techniques is the lack of computer systems fast enough to process all the data involved and provide an answer in time. This is, for example, true in the application of optimisation techniques to achieve efficient design in engineering. A very challenging optimisation involves the search for the most efficient aerolastic configuration of aircraft. This is of the utmost importance for economic designs of modern civil and military aircraft. In many areas of business, the allocation of resources such as people, raw materials, capital or time can be among the most challenging of problems. The difference between *good* solutions to these problems and a *near-optimal or optimal* solution may mean the saving of huge sums of money. The main application areas for optimisation include design of complex structures, investment and portfolio management, production planning and scheduling, distribution planning, vehicle routing and crew scheduling. The help of supercomputers through a better man-computer interaction gives a dramatic enhancement to the **creative power** of men. It helps them to incorporate in their design choices a broader set of precise elements of technical, economical and managerial importance.

Artificial intelligence (AI) will benefit from symbolic processing on very high-performance systems. AI produces applications of certain although limited commercial value, but it holds the promise of far more commercially attractive applications if the required fast processing, allied with sophisticated software and flexible database systems, can be made available at acceptable cost. Examples of such applications would be real time natural language systems, decision support tools in combinatorial problems such as real-time scheduling problems which cannot be adequately modelled numerically only, tools for supporting financial analysis and many more.

The **impact on society** of high-performance computing systems is not limited to its benefits to industry, business and services. Be it the spread of diseases, the recognition and translation of natural languages, the global climate changes or the dynamics of complex economic systems, it is well recognized that the major problems affecting our society are planetary in their nature, and they need to be studied and solved on that scale. In many cases the absence of complete data, such as for the atmosphere and biosphere, or for the human population, forces us to develop criteria for significant predictions. These necessitate the understanding of very complex systems, the behaviour of which can only be fully understood and precisely predicted by detailed modelling using high-performance computers.

It is probably daring, but also challenging, to speculate upon the possibility of creating machines which could accept high level specifications for a desired product or service, and, by combining the above techniques, provide a set of optimal solutions. Such systems could be used, for instance, not to simulate the effect of a drug, but to find in an inverse procedure the composition of the drugs which have the desired effects. Many of the elements of such a system are already there, and many still need to be developed, but the availability of adequate computing equipment is one of the necessary pre-conditions. In this connection, we have to inquire also into the design of custom orientated parallel computers conceived for a specific purpose.

Advances in all these areas need new important developments **now**, not only leading to necessary major increases in computer power, but providing for a revolutionary improvement of the whole infrastructure of large-scale computing. This will also involve improving hardware and software for visualization, databases management and high-performance networking facilities.

The issues involved may be clarified by three examples chosen for their relevance to the European situation: molecular design, the forecasting of climate changes and applications in the automotive and aerospace industry. Many examples of equal or greater impact on society, industry business and services could be added, such as applications in nuclear fusion research.

The first example, the **design of new molecules or materials**, is basic to the chemical and pharmaceutical industry. Combined with developments in biotechnology, new proteins will be designed and produced, leading to new products having an impact on everyday life. Traditional strategies for the design of new compounds involve intelligent trial-and-error methods. Computer modelling is just starting to make an impact on the design process. However, the design is still very time consuming and expensive. The considerable increase in computer power that we propose is needed to speed up the design of new products, which will allow the European fine-chemical industry to remain pre-eminent. When fully developed, modelling methods will incorporate a combination of simulation, database search and visualization, and will enable the designer to by-pass costly and time consuming laboratory experimentation. They will require a thousandfold increase in the presently available computer power. It is foreseen that the application of molecular modelling techniques at this level will be mandatory for the chemical, pharmaceutical and biotechnological industries in order to remain innovative and internationally competitive.

Our second example is the monitoring and simulation of the characteristic statistical **behaviour of the earth/atmosphere/ocean eco-system** over long periods of time. Governmental organisations are already investing heavily into new techniques to observe the atmosphere and ocean and into new research programmes aimed at assessing man's effect on the environment. Here numerical models of the atmosphere and ocean play a crucial role in the study of climate. However, the thorough evaluation of the global and regional response to increases in the concentration of greenhouse gases from anthropogenic sources, for example, will require the use of considerably more sophisticated climate models than has been possible hitherto. Significant increases in resolution and the use of much improved representations of the physics and chemistry of the atmosphere are required now.

Our third example refers to the **aerospace and automotive industry**. Here the availability of simulation and optimisation techniques on high-performance computers allows to achieve substantial reductions both in the development cycle-time for a new model and in the overall costs of the model, from the moment of its conception to the moment of its phase-out from the market. High-performance computing is critical in providing an *excellent response* to the project requirements and in producing a *robust solution* in terms of reliability and functionality. The application of high-performance computing in these industrial areas can be classified in the following categories:

- simulation of advanced technological processes, like sheet metal forming, forging, injection moulding, ceramic composite materials;
- crash analysis of aircraft and cars;
- simulation of fluid dynamics, combustion, lubrication;
- aerolastic response and manoeuvrability of aircraft;
- real time simulation of the behaviour of critical and vital components;
- automatic research of the best and soundest solution, obtained with the combination of all the simulation techniques depicted above and of sophisticated optimisation techniques.

Some of these applications for realistic models with real time response need a 10^2 to 10^3 increase in computing power. The real benefit of the availability of increased computing power and user friendly software is the possibility of simulating product properties and of optimising production methods. This strategy will result in by-passing substantial parts of costly experimentation in the design process and form an essential part of the path toward a *total quality approach*.

Key technologies

In order to meet these major **scientific and technical challenges** in a timely manner, new computer architectures, comprising both **hardware and software** will have to be developed. In assessing these computational needs, four areas can be highlighted:

- High-performance computing hardware and software, innovative architectures, algorithms and applications;
- Storage and retrieval techniques for large databases;
- Visualization;
- Networking.

Currently the sustained **processing speeds** that can be achieved with today's generation of supercomputers is typically of the order of 1 Gigaflops (billion arithmetic results per second). For many applications, such as those indicated before, the foreseen requirement by major scientific and industrial experts for this decade is of the order of a Teraflops - a 1000 fold increase in sustained speed.

The requirements for **larger data stores and faster retrieval methods** of many different types of objects such as numbers and images, are increasing rapidly in many important areas, including engineering, high energy physics, oil industry, biotechnological and medical sciences and the earth and atmospheric sciences. Comparable efforts to support the development of hardware and software for the management of very large data bases are also urgent.

Graphical representation is the only effective way of evaluating the results from large-scale computation involving large data volumes. For most of the applications stressed here, two and three dimensional display and real time animation is required. To achieve this, individual scientists and engineers will require the ***personal*** use of fast computer hardware and improved graphical displays. Indeed, it is likely that in many instances scaled-down versions of the fastest machines would provide ***on-the-desk*** power for such applications.

Networking for large amounts of data at high speeds over large distances as well as local area communications will require major improvements in order to meet the requirements of wide scale collaboration efforts in key application areas, combined with the need to access large-scale computational resources and major data bases.

In summary high-performance computing is strategic to the development of society through progress of science and technology to increase the competitiveness of industry and human well-being. As high-performance computing is becoming a motivating technology for a wider segment of industry, there is a shift of emphasis in the market towards high-performance computing. For example Electronic Trend Publications also expects strictly commercial applications of supercomputers to grow from the current 6% to 14% by 1993 (of cumulative shipments).

In the next section we will discuss the main features of high-performance computing in Europe together with the reasons for our grave concern with the current European position in this area.

European situation

High-performance computing is international in its requirements and its user community. Yet the European market has a number of features which distinguish it from the US and the Japanese markets. The economic requirement for state-of-the-art information technology in Europe is clearly shown by the fact that Europe alone represents 25% to 30% of the world market. There are hundreds of high-performance computing installations in Europe, not unlike the situation in the other leading countries, Japan and USA.

The European information technology **market** is very open as can be seen by the 20% market share obtained in Europe by Fujitsu supercomputers distributed by Siemens and Amdahl. This should be compared to the 20% market share in Japan of CRAY - who has a totally dominant position everywhere else - and to the single Japanese-made supercomputer sold in the US so far. In spite of or because of this openness there is no major European supplier of high-performance computers. Almost all existing machines come from American or Japanese manufacturers. A number of start-up companies exist, but they experience difficulties becoming international and growing fast enough compared to their US counterparts.

High level **research** and development activities in such areas as parallel computing or scientific and technical software applications writing are taking place in Europe, but they do not, normally, generate highly successful European commercial products on the world market. In contrast non-European major high-performance computer manufacturers are using this large scientific potential, and lack of competition, to improve their products and establish themselves better in the world market.

The **lack** of commercially **successful companies** in Europe has resulted in a lack of state-of-the-art industrial teams in the key technologies used in high-performance computers today: commercially produced high-density ECL and GaAs semiconductors, RISC, superscalar and VLIW[1] architectures, FORTRAN and C optimising compilers, etc. The consequence is that the European entry in the high-performance computing industry cannot be entirely based on present European technology. We see as necessary, at least at the beginning, the **acquisition** of **technology** and **components** from abroad, possibly within the framework of industrial collaborations and joint ventures.

---

[1] Very Long Instruction Word

The situation whereby **Europe is not a producer** of one of the key elements in the economic and scientific development of its society is unacceptable. The high-performance computing industry provides tools which are crucial for maintaining competitiveness in almost all advanced industrial sectors. In this respect information technology is at least as relevant as the automotive or aerospace industries, and far more than steel, in which massive investments are being made in order to maintain and improve Europe's role. The absence of Europe contrasts painfully with its growing political and economic relevance and with the high technological level of academic research and developments in this field. Talents and the motivation exist for a strong European presence which could become a reality, provided the necessary financial and human resources are deployed and coordinated toward a common challenge. This would represent the natural evolution of the ambitions of Europe, especially after the success of the information technology research and development programmes over the last years.

The **consequence** of Europe's absence can be fully grasped considering that high-performance computing products are essential tools for the production of goods and services. Presently Europe has no control over the the design and production of these key tools. This is all the more surprising since significant research and development is being sustained in Europe. The immediate effect of this situation is a **political weakening**. To ensure its scientific and industrial future, Europe has no choice but to negotiate with foreign suppliers the acquisition of high-performance computing equipment. Consequently high-end information technology in Europe is acquired mainly in the form of packaged products rather than as the result of a close collaboration with the industrial research and development teams, as it is common in Japan and the US. There is a **delay** in the acquisition and efficient exploitation of hardware and software products. We must constantly adapt scientific and industrial applications to the availability on the market; rather than being able to tailor the supply, particularly of software packages, together with the suppliers, in response to the genuine needs of Europe. Regrettably the above considerations not only apply to the specialized market of high-performance computers, but also to the rapidly growing market of advanced graphic workstations.

Demand for high-performance computing is constantly rising. Europe does not get any economic benefit from this, rather it experiences an **economic loss**. This will become more and more serious as the market expands, and the need reaches wider and wider sectors of society. The negative effects on the trade balance will also increase.

The lack of an industrial infrastructure in the field causes the loss of **spin-offs**, both in information technology at all levels and in other related industries. This entails a global loss of industrial competitiveness and readiness to respond to market pressure in comparison to competitor American and Japanese industries. The process is prone to a positive cultural feed-back. The best researchers and experts in high-performance computing formed in Europe go and work where innovation is generated, typically the US. This produces a **brain drain** which has the effect of impoverishing the European information technology cultural environment and of creating obstacles to the establishment of a successful European industry in the area.

The **development** of the European cultural level and the ability to produce innovation in this field assumes a scientific and technology background that can be created only if both hardware and software are developed in Europe by teams at world leadership level. There is a need for close contact and interaction between scientists, developers and users.

What, and why now?

For Europe to become a recognized market force by the end of the decade, more industry-oriented research and development must start now. The efficient design and exploitation of new architectures requires a substantial hardware and software effort, both at the system and application level and a close collaboration between computer architects, engineers, system designers and application scientists. The achievement of such goals is not only constrained by technology but also depends heavily on political commitment and determination. This is clearly recognized in Japan and North America, where governments are mobilizing large resources in bringing together manufacturers, national laboratories and universities.

In many fields, substantial computational resources are an essential tool to gain a deeper theoretical understanding of the basic principles of science. This is the case, for instance, of simulation, where the general principles of the modelisation theories are beginning to be recognized and studied in greater detail. Progress in these areas requires initiatives at the European level to make available to scientists and engineers the highest-performance machines in a production environment. Training and consultancy support for applications and communications should be provided. Research in the areas which are at the boundary between applied Computer Science and the specific disciplines should be encouraged.

Even when resources are available, users may not exploit them, either because they are not aware of the potential benefits or because they lack sufficient training. Potential users should be educated in the opportunities offered by high-performance computing. This cannot be achieved via training and education of professionals alone: the culture of high-performance computing should be introduced in the curricula of study of Universities, higher education institutions and perhaps even secondary schools, and this should not be limited to scientific disciplines. Students should be taught what a high-performance machine can do for them and how to use such a tool.

The problem of high-performance computing goes beyond the capability to produce fast machines. High-performance computing resources must be accessible via fast and reliable network connections, particularly to allow such applications as interactive graphics and animation.

System software and human interfaces should evolve toward the definition of a user environment which must be seamless, user friendly and open. Existing and future scientific and engineering applications must be easily transportable from one computational resource to another. At the same time, the software must offer advanced users or authors of applications the tools to fully exploit the parallelism offered by the hardware to obtain the required performance. Excellent compilers and debugging environments are as important as the availability of the hardware itself. Completely new approaches must be pioneered to support efficient program development in distributed environments with a high degree of parallelism. The efficient implementation of leading industrial and scientific applications should be seen as a priority.

The European effort cannot be limited to a single aspect, either hardware or software, neglecting the other. Software of the highest quality is indeed already produced in Europe, but, at present, it only has a significant impact on the market when supported by a large non-European computer company. New application software should be written to be portable and run efficiently on different machines. **Europe** should make a substantial **effort to develop** and market system **software,** compilers and interfaces which exploit the features of the machine architecture hiding them from the user at the same time, so that the same application can run without modification on different systems. Only software that exploits new architectures and new modelling capabilities will survive. This implies a close relationship with the producers to take advantage in a timely fashion of the hardware developments. This requires mastering the architectural conception and the production of systems locally in Europe. Europe may, at present, lack the necessary industrial infrastructures in certain areas, but these limitations should be overcome in the medium term.

From the **architectural** point of view an increase in performance cannot come from the enhancement of known architectures alone, be these super-scalar or vector, but only from combining these with a more intensive exploitation of parallelism. The three main architectural approaches which should be explored in order to progress toward the Teraflops goal are:

- Tens of very powerful CPUs, shared memory and fast access to it;
- Hundreds of CPUs with local memory and a variety of communication networks between them - mainly used as MIMD (Multiple Instructions and Multiple Data streams) machines;
- Thousands of relatively cheap CPUs with a combination of shared and distributed memory, mainly used as SIMD (Single Instruction and Multiple Data streams) machines.

There are no elements at the moment to privilege one of these three directions over the others and experimentation should proceed on all of them. The range of application and validity of each approach can only be assessed via the realization of machines exposed to the judgement of the market. The economical and technological viability of special purpose hardware should be studied and assessed. Intermediate solutions should be explored, like for instance a general purpose machine which could be optionally equipped with special purpose hardware for a specific task. We have examples of this in vector co-processors or floating point units, but the concept could be extended. The possibility of customizing a machine via the addition of special purpose, custom design modules, especially produced for a certain class of applications should be explored to assess its economical feasibility. This demands a steady close liaison between researchers, practitioners and hardware producers.

Much more than for other complex industrial products, a substantial improvement of the effective performances will come from the interaction of **hardware, architectures and software** in the design of systems. The design of systems, from the handling of instructions and data flow to the architecture and systems software is directly influenced by the need of an easy and efficient implementation of basic and applied software and viceversa. This requires that, in the long term, Europe must be able to master hardware, architecture and software, developing them in close connection with the users.

A unique opportunity is arising now. New basic technologies and architectures are just emerging. The start of an effort in Europe can be immediately effective without having to overcome a large technology gap. The new technology, mainly relying on highly parallel processing, can offer dramatic reduction in the cost of large computation for industry and a substantial increase in computational capacities. This can lead to a breakthrough in many scientific areas in the next decade. An additional reason for the present window of opportunity is the emergence of open standards in the operating systems, languages and user interfaces areas (UNIX™, X-WINDOW™, etc.). This allows user applications to be introduced in heterogeneous environments, where different hardware platforms can be exploited.

The breakthroughs in parallel architectures allow the use of non-exotic technologies to conceive very high-performance computers scaleable in capacity by a factor of more than one hundred. The high cost effort of pushing the cycle time to the extreme limits of 2 to 1 nano-seconds may be less essential now.

A number of the specialized building blocks and technologies to put together high-performance computers are available on the open market. The list includes optimising compilers, semiconductors and design automation tools, peripherals, networks, etc. It is therefore no longer either necessary or advisable to master the full spectrum of technologies. Successful development is more a matter of synthesis.

In spite of the lack of commercially successful enterprises, there are in Europe very well trained experts who could start immediately to work in the the above directions. These people are mainly scientists either building massively parallel special purpose machines for their own research activities or developing application software in national or European research laboratories. Their expertise will be valuable in the creation of a European high-performance computing infrastructure.

Proposal

The Working Group believes that Europe should make a **substantial investment** in the field of high-performance computing in order to become an active **participant** and a recognized **partner** at the level of the leading American and Japanese industries (both vendors and users) by the end of the century.

To achieve this objective, the Working Group recommends a **global action plan** articulated in 5 interrelated actions.

1 Promotion of High-Performance Computing

The European Authorities should stimulate the creative use of the most powerful computing systems available. We recommend that all practical measures be urgently taken to help spread the know-how associated with their use and to encourage the further imaginative development of such devices. Scientists and Engineers should be made aware of and become familiar with techniques such as physical modelling, simulation, multidimensional interactive graphics and scientific data visualization, vector and parallel programming in order to be able to tackle large, complex problems and applications.

To this effect the development of an **advanced Pan-European High-Speed Network** is of strategic importance. The present generation of data communications networks should be rapidly evolved into a large scale high-performance (multi-Megabits/sec.), multi-protocol backbone, while the necessary research and development is to be encouraged in order to allow Europe to compete in the Gigabits links race. This involves substantial investments both in hardware and software. This will permit the formation of a European high-performance computing community and would identify the real user's needs and present them to the suppliers.

2 Development of a European High-Performance Computing Industry

The design and production of advanced high-performance machines should be pursued vigorously. In order to obtain the required computer performance (Teraflops speed by the year 2000) priority should be given to parallel architectures. In these fields there is a high potential for discovery and innovation. A significant effort should be extended to the design and production of innovative computer architectures for specific applications in research and industry.

The development of a competitive and credible European industry in this domain will foster the ability of designing and producing all the equipment and software which is required for a leading-edge computing environment, such as high-bandwidth communications, high-performance graphics, data storage systems and workstations.

3 Development of European Software

A major European effort should be directed at the inventive development of novel software. In view of the revolutionary changes taking place in the architecture of high-performance computer systems as well as in their interactive real-time use, either locally or via networks, entirely new software concepts need to be developed and implemented. Existing application software, as well as the underlying support software, often needs to be adapted or redesigned to take advantage of the rapidly evolving hardware. We recommend that major efforts be deployed to face this challenging software evolution and to exploit the opportunities offered by the emerging computing devices. Part of this effort is the design, enhancement and application of standards which could reduce this effort to the minimum for the end user. Technology transfer mechanisms should be studied and employed to successfully market the high quality software produced by European academic and scientific research. Such mechanisms must, amongst other things, encourage the use of new software by the scientific and industrial users.

4 Research and Development

We recommend that the existing competence in Industry, Universities and Large Research Laboratories be effectively mobilized to carry out the basic and applied research necessary to raise the competitive level of European industry in the domain of high-performance computing. It is critical to ensure that close collaboration be achieved between the leading edge users and the emerging industry. The aim should be to supplement the current effort on basic components by an increased emphasis on systems integration.

The European Authorities should promote advanced pilot projects involving the best European Institutions, in order to develop the theoretical and practical understanding of the various aspects of high-end computing applications. These projects should actively contribute to the design and implementation of all the software related to the exploitation of parallel machines in Science and Engineering.

5 Promotion of Education and Training

In order to alleviate the critical shortage of skilled engineers and scientists for the design, development, production and intelligent use of high-performance computer systems, we recommend that Education and Training in all areas connected with this field be strongly enhanced.

The high-performance computing culture should be spread among scholars of all disciplines as well as in industrial, commercial and financial environments. This can be enhanced via the introduction of the subject in the curricula of Universities and Engineering schools. In addition, access to leading-edge machines should be encouraged and facilitated.

Preliminary estimate of the required investment

This Working Group believes that a substantial effort over at least **ten years** is necessary in order to establish a competitive European high-performance computer industry.

Although we felt that it was not part of the mandate to make a detailed financial plan, it is however, appropriate to provide preliminary investment estimates. The working group considers that the totality of the new investments over the next ten years should be commensurate with that made in US and Japan. **In order to implement the proposals as outlined above, it is estimated that, by 1995, various sources of financing (Community, national, industrial, business and telecoms) should contribute to a total annual investment of the order of 1000 MECUs.** The specific mechanisms by which this money would be controlled and allocated to projects will have to be examined in considerable detail. During the first five year period, beginning in 1991, there will be **a gradual increase of investments**. The rate of the expenditure increase in each of the different action areas will largely be governed by the availability of the necessary competence and knowledge.

The heavy **involvement** and **commitment** of the **European users** of high-performance computing techniques is the key to the success of the programme. A strong political commitment to **continuity** of the action is necessary to obtain this involvement.

The ultimate **success** of the programme will be judged on the increased **market presence** of European industry in the key sectors of high-performance computing: software and applications packages, advanced architectures, high-speed networks, high-speed workstations, etc.

High-performance computing equipment running European-developed software should be acquired and made available on a broad scale via a rapid and reliable network with services and access rules, to the applications groups in Europe. This will help in supporting an emerging European industry and in creating the cultural environment necessary for a unique market in the area: knowledge, communication, norms, feed-back to industry, educational programmes.

A first estimate of the needs in each of these areas indicates that about 200 MECU would be necessary annually to support high-speed networking development between about 100 centres of service and 100 to 200 user groups. Another 200 MECU is estimated to be necessary per year for user group projects in advanced high-speed computing and for the corresponding services and European hardware acquisitions in about 20 advanced pilot centres.

On the industrial side the support of emerging industry in key areas of high-speed computing such as advanced architectures, system software, data storage and retrieval, graphics and high-speed workstations, highly parallel computers etc. is estimated to require about 300 MECU annually, of which half should be earmarked for software, an essential part of any such effort.

The effort on the applications side: industrializing and marketing in particular, of packages which only exists in a pre-industrial state is estimated at 200 MECU per year. Optimisation efforts and development work in new, promising directions such as AI, physical modelling and simulation, optimisation and so on, are also included in this figure.

Finally the educational and cultural aspects of high-performance computing (publications, scholarships,etc.) should be developed with about 50 MECU per year.

During the early phases of the programme great profit should be realised by the existing programmes, ESPRIT and RACE in particular, in order to encourage emerging industrial efforts both in systems and in applications. In this context, of great importance is the effort currently going on to link computer centres in the framework of RACE.

The review of the quality of the contract execution by computers centres, applications groups, emerging industry and network operators will have to take into consideration both technical and economic criteria.

The user community (industry as well as research) has to be largely represented in all the instances of the programme office, committees and advisory bodies. Economic and marketing expertise should also be widely activated in order to evaluate and help projects.

Appendix

Working Group Members Curricula

The working group was composed as follows:

***Prof. Carlo Rubbia** (Chairman)*
CERN

Carlo Rubbia is Director-General of CERN (the European Laboratory for Particle Physics in Geneva) since 1 January 1989. Rubbia has been working at CERN as Senior Physicist since 1961. After completing High School, Rubbia was admitted to the Scuola Normale of Pisa where he completed his University education with a thesis on Cosmic Ray Experiments under the guidance of Marcello Conversi. From 1970 to December 1988 Rubbia has spent one semester per year in Cambridge, Massachusetts, at Harvard University where he was Higgins Professor of Physics. Early in 1983 at CERN, an international team of more than 100 physicists headed by Rubbia and known as the UA1 Collaboration, detected the intermediate vector bosons, a triplet of particles, the W^+, the W^- and the Z^0, which had become a cornerstone of modern theories of elementary particle physics, long before they were observed by Rubbia and collaborators. The revolutionary techniques developed with Simon van der Meer were the motivation for awarding them both as Laureates for the 1984 Nobel Prize for Physics. Carlo Rubbia has also been awarded many other important prizes and Honorary degrees from Institutions and Universities all over the world and is also member of Scientific Academies, such as The Soviet Academy of Sciences.

***Dr. Federico Carminati** (Secretary)*
CERN

Federico Carminati received his Degree in Physics at the University of Pavia in 1981. He has then been working at various High Energy Physics experiments both in Europe (European Laboratory for Particle Physics - CERN) and in the Unites States (Los Alamos National Laboratory). Since 1985 Federico Carminati is working in the Computing and Networks Division of CERN, where he is now responsible for the detector simulation software activities.

Dr. Renzo Allaria

Sistemi e Informatica

Applicazioni e Tecniche

FIAT Auto

Dr. Renzo Allaria received his degree in Physics at the University of Genova, Italy, in 1973. His experience is in the analysis, design and implementation of Computer Aided Engineering and Computer Graphics Systems in industrial environments. He is among the authors of the APPLE-SAP system, a Computer Aided Engineering (CAE) program for structural analysis. At present he is in charge of CAE applications in FIAT-Auto, in the Information Systems Department. He is responsible for the definition of hardware and software architecture for CAE and CAT applications, as well as for the analysis of engineering problems and the related software developments.

Prof. John Argyris

Institute for Computer Applications

University of Stuttgart

Prof. John Argyris, obtained his Degree in Engineering at the University of Munich; he is Fellow of the Royal Society and is recipient of its Royal Medal for the creation of the finite element method and for leadership in engineering sciences. He is a Honorary Fellow of the Royal Areonautical Society and a Fellow of many International scientific Institutions and Academies. He is a holder of a great number if international scientific distinctions in Europe and in the USA. He has been Professor of Aeronautical Structures in the University of London, at Imperial College of Science and Technology, in 1955-75, Visiting Professor in 1975-78, and he is now now Emeritus Professor. He has been director of the Institute of Statics and Dynamics of Aerospace Structures in Stuttgart from 1959 to 1984 and he is now director of the Institute for Computer Applications in Stuttgart, since 1984. He is principal Editor of the Journal of Computer Methods in Applied Mechanics and Engineering since 1972-. Prof. Argyris was awarded several honorary professorships in Universities in Europe, U.S. and Asia. He is the authors of several books on Structural Analysis and finite element methods applied to aerodynamics and other physical subjects. These publications appeared in Ingenieur Archiv, Reports and Memoranda of Aeronautical Research Council. Journal of Royal Aeronautical Society and Aircraft Engineering. CMAME, Jl of AIAA etc.; over 310 scientific publications.

Prof. Herman Berendsen

Physical Chemistry Dept.

University of Groningen

Herman J.C. Berendsen (1934) is professor of physical chemistry at the University of Groningen, The Netherlands. He was educated at the University of Utrecht where he obtained his degree in physics in 1957. After service as officer of the Royal Navy Reserve, he spent two years at the Massachusetts Institute of Technology in the Neurophysiology group of Warren S. McCulloch. In 1962 he obtained his Ph.D. at the University of Groningen on a thesis about biophysical applications of nuclear magnetic resonance. In 1963 he was appointed associate, and in 1967 full professor of physical chemistry at the University of Groningen. In the 1970's he shifted his interest from experimental biophysics to molecular simulations and organized a series of meetings and workshops in molecular dynamics at the Centre European de Calcul Atomique et Moleculaire (CECAM) in Orsay, France. The first simulation of a protein molecule was carried out during such a workshop in 1976. His group in Groningen, now part of the BIOSON Research Institute, has pioneered in the development of methods for simulations of complex molecular systems. The program package GROMOS, developed in Groningen, is in use in more than 250 laboratories in over 40 countries. The latest project of the group is the construction of an application-oriented parallel computer system for molecular simulations. Professor Berendsen is member of the Royal Netherlands Academy of Arts and Sciences since 1979. He serves on several editorial boards and has acted as chairman of the governmental Computer Commission, advising the Minister of Education on major computer equipment for Universities and research institutes. He is (co)author of over 100 scientific publications.

Dr. Tor Bloch

Advanced Computer Research Institute

Tor Bloch is Vice President Marketing and Planning of A.C.R.I. Tor Bloch has his Degree in Mathematics from the University of Copenhagen in 1964. He has worked at CERN, the European Organization for high Energy Physics, as a systems programmer, project leader (CDC 7600 installation) and head of systems, deputy division leader and later as responsible for the Advanced Computing Group. Tor Bloch consulted with the UN, with the European Centre for Medium range Weather Forecasting (ECMWF) and visited for periods of several months the Lawrence Berkeley Laboratory, the High Energy Physics Research Laboratory (KEK) in Tsukuba (Japan). He helped to create from scratch and directed from its start in 1983 the CCVR, Centre de Calcul Vectoriel pour la Recherche, a Supercomputer centre available for use by French scientists through a scientific committee and running the French Meteorological forecast every night. A CRAY-2 was installed in late 1986. Tor Bloch has then worked with BULL as Director of the Scientific and Educational Market and responsible for sale, marketing and strategy.

Dr. David Burridge

Research Department

ECMWF

Dr. David Burridge is the Deputy Director and Head of Research at the European Centre for Medium-Range Weather Forecasting (ECMWF) which is an International research and operational centre located in Reading U.K. His background is in Mathematical Physics and he holds a Doctorate in Applied Mathematics (Bristol University). Apart from a short Post-doctoral appointment in Florida State University he has spent the whole of his professional career either developing of managing the development of operational weather forecasting system initially at the U.K. meteorological office (1970 - 1975) and subsequently at ECMWF (1975 - present). As the ECMWF Head of Research he is responsible for formulating and directing research in Numerical Weather Prediction. During his time at ECMWF he has always been involved with the supercomputer developments and acquisitions and he has chaired and been a member of many computer evaluation boards. His research interests and publications are mainly in the fields of numerical weather prediction, numerical methods for solving partial differential equations and the application of supercomputers to such problems.

Dr. Jean-Marie Cadiou

Commission of the European Communities

Jean-Marie Cadiou graduated from the Ecole Polytechnique in Paris in 1962 and from the Ecole des Mines de Paris in 1966. He began his career as a civil servant in the French Ministry of Industry, working on environmental issues. He then switched to computers and went to Stanford University where he received his Ph.D. in Computer Science in 1972, and then he returned to France to become Scientific Director at INRIA. In 1975, he joined IBM in the San Jose Research Laboratory, working in the area of relational data-bases and natural language translation. In 1977 he was appointed manager of the IBM France Scientific Centre in Paris and returned to California in 1979 to take responsibility for Programming Technology Research at the IBM San Jose Laboratory. In 1981, Dr. Cadiou left IBM to become Director of New Information Technologies in the Commission of the European Communities. His current position is Director of Information Technologies and ESPRIT.

Dr. Michel Carpentier

Commission of the European Communities

Michel Carpentier is Director-General for Telecommunications, Information Industries and Innovation at the Commission of the European Communities. As Director-General he is responsible for programmes in the development of technology, electronics and computer science such as ESPRIT and RACE (telecommunications) and more generally, policy and action programmes concerning telecommunications, information industries and services, standardization in these fields, and innovation. Mr. Carpentier has extensive experience both in R&D and industrial policy. After having been Administrator in the French Atomic Energy Commission, he became Head of the R&D contracts division at EURATOM, and later Head of Division at the Directorate-General for Industrial, Technological and Scientific Affairs at the newly-created Commission of the European Communities. He then created and headed for a period of ten years, first as Director and later as Director-General, the Environment and Consumer Protection Service. He received several international prizes, was President of the Administrative Board of the European Foundation for the Improvement of Living Conditions in Dublin, and a Member of the International Commission to fight the pollution of the Rhine. He later became Director-General at the Energy Directorate-General, where he was not only responsible for Community programmes related to new energy sources and energy-saving, but also for relations and agreements with developing countries, particularly with China and Latin American countries. Mr. Carpentier is a graduate of the Ecole des Hautes Etudes Commerciales (business administration) and of the Ecole des Sciences Politiques of Paris. He also has a degree in law and economics. He is doctor en Science honoris causa of Loughborough Technological University and Foreign Member of the Swedish Royal Academy of Engineering Sciences.

Prof. Jack Dongarra

Dept. of Computer Science

Univ. of Tennessee

Prof. Jack Dongarra is presently Distinguished Scientist at the Department of Computer Science at the University of Tennessee and of Mathematical Sciences at the Oak Ridge National Laboratory. He obtained his Ph.D. in Applied Mathematics at the University of New Mexico in 1980. From 1975 to 1980 he has been Senior Computer Scientist at the Argonne National Laboratory and from 1980-1989 he was Scientific Director of the Advanced Computing Research Facility at the Argonne National Laboratory. He is member of the Society for Industrial and Applied Mathematics (SIAM) and of the Association for Computing Machinery (ACM). He is also member of the editorial board of several leading Computer Science publications. He is holding positions of adjunct professor at the Rice and Northern Illinois Universities.

Dr. Herve Gallaire

BULL S.A.

Dr. Gallaire received his Ph.D. at the University of California, Berkeley in 1968. He has been Professor of Mathematics and Computer Science at ENSAE from 1970 to 1980. From 1972 to 1980 he held the position of Head of the Research Department in Computer Science at ONERA-CERT. From 1980 to 1983 he has been Head of the Research Department in Computer Science of CGE. He then moved on to be Director of the European Computer - Industry Research Centre, a laboratory jointly owned by BULL, ICL and Siemens. From 1989 he is Vice President, R&D for Applications, Distributed Services, Communications and Networking at BULL S.A. Dr. Gallaire is also President of the Association for Logic Programming and European Representative of ACM.

Prof. I.H. Hillier

Department of Chemistry

University of Manchester

Ian H. Hillier is a Professor of Chemistry at the University of Manchester. His research interests are in the application of accurate modelling methods to study chemical problems and he has a number of collaborative projects with the U.K. chemical industry in this area. His work involves the extensive use of vector and parallel computers and he is the author of 280 papers in the scientific literature.

Prof. Jacques-Louis Lions

CNES

Prof Lions graduated at the Ecole Normale Supériore de Paris in 1950. He has held several professorships in Universities and he is now Professor at College de France. He is President of the Centre Nationale d'Etudes Spatiale (CNES), president of the Scientific Council of EDF and he will take over the presidency of the International Mathematical Union from 1991. He is member of Accademia of Sciences in several countries and he has a number of Honoris Causa degrees. His activity has been mainly in the fields of the methods for system analysis and control, of the mathematical, numerical and informatic treatment and of the applications, with particular attention to the mechanics of solids and fluids and to the climatological modelization. He has also been working in the area of econometrical modelling. He is the author of more than twenty books translated in English, Russian, Chinese, Japanese and Spanish.

Dr. Jean-François Omnes

Commission of the European Communities

Dr. Jean-François Omnes joined the Commission of the European Communities in 1984. He was responsible for the Information Processing System Architecture sector, within the ESPRIT Programme. He is currently Deputy Head of Unit for the Information Processing System Division. J-F. Omnes got his engineering degree in 1967. He joined the Centre National d'Etudes des Télécommunications where he worked on operating systems, computer evaluation and provided a support for application development. In 1978 he became Head of a Research Department on system engineering and programming. Up to 1981 he lead the CNET work on multiprocessor system developments and programming environment. In 1982 he was responsible for the department Simulation and Protocol Validation, where he worked on specification languages, simulation and evaluation of telecommunication protocols, in particular ISDN protocols. At that time he was a member of the CCITT working group for language specification standardisation. J-F. Omnes was president of AFCET Informatique from 1982 to 1984 and a member of the AFCET board from 1983 to 1985.

Dr. Pierre Perrier

DGTDEA, AMDBA

Dr. Pierre Perrier received his Diploma from the Ecole Nationale Supériore of Areonautics in 1958. His Ph.D. work was on rarefied gas dynamics and hypersonics. After that he worked continuously at Dassault Aviation, where he made developments in the area of Computational Fluido-Dynamics and other critical computer simulation and applications of informatic tools for the design of all military and civil Dassault aircraft from the Mirage to Hermes, including research in the field of aerodynamics. He is the author of many papers in fluid mechanics and computational mathematics and he is a corresponding member of the French Academy of Science.

Prof. F. Troyon

Centre de Recherches en Physique des Plasmas

École Polytechnique Fédérale de Lausanne

Prof Troyon received his Degree of Engineering Physicist from the Ecole polytechnique de l'Universite de Lausanne in January 1957 and his Ph.D. from the University of Rochester (N.Y.) in May 1962. He joined the newly created Centre de Recherches en Physique des Plasmas (CRPP) in Lausanne. From September 1973 to August 1974 he has been Research physicist at the Plasma Physics Laboratory in Princeton. He was titular professor of controlled fusion in 1974 at Ecole Polytechnique Federale in Lausanne (EPFL) and Director of CRPP since 1982. Ordinary professor of Physics at EPFL since 1983, since 1980 he is Swiss representative in the Euratom Fusion Programme; in particular chairman of the JET Scientific Council, member of the Fusion Technology Steering Committee-Programme and representative of Euratom in the ITER Scientific and

Technological Advisory Committee. His main research interests are in the field of the stability of toroidal magnetic confinement systems. Major contributions include the development of large numerical codes to study operational limits of tokamaks which are now of widespread use in the fusion community, the derivation of a scaling law giving the maximum pressure that can be stably confined in a tokamak through extensive numerical experimentation. Long experience since 1962 with all the successive generations of Computers.

Dr. Roberto Vio

Sistemi e Telecomunicazioni

Sviluppo, Coordinamento e Controllo

FIAT

Dr. Ing. Vio graduated in Electronic Engineering at the Polytechnics High School in Torino (Italy). He has now 18 years of experience in planning, design and realization of Information Technology applications in various industrial environments. Presently he is in charge of the Information Technology Development and Coordination Department at FIAT Headquarter, whose main goal is to define architectures, guide-lines and standards valid throughout the FIAT group. His task is to assure a homogeneous development of Information Technology applications in all the operating divisions of the FIAT Group, having a constant comparison with external developments: vendors, competitors, public research. In this context he is in charge of the technical relationships with the European Commission.

Prof. Paolo Zanella

CERN

Prof. Paolo Zanella obtained his degree in Mathematics and Physics at the University of Bologna in 1959. Later he was in charge of computing at the SACLANT ASW Research Centre, La Spezia, Italy from 1960 to 1962. He is member of the CERN Scientific Staff from 1962 where he has been doing pioneering work in computing for High-Energy Physics. Currently he is Advisor to the Director General of CERN, after having directed the Computing and Data Handling Division of CERN from 1976 to 1988. He also holds the position of Associate Professor with the Informatic Dpt. of the University of Geneva (CH), where he teaches Computer Architecture since 1981. He is also lecturer on Information Technology and its application to natural and human sciences at the

Scuola Normale Superiore, Pisa (Italy) and director of the CRS4, a Research, Development and Education Institute on high-performance computing and mathematical modelling, being established in Sardinia (Italy).