europhysics news — Computers
BULLETIN OF THE EUROPEAN PHYSICAL SOCIETY

J.A. Volume 19 Number 2 February 1988

Dear EPS Member,

Taking the opportunity of the appearance of this special computational physics issue of Europhysics News, *I should like to inform you about the activities of the Computational Physics Group (CPG) of EPS.*

The CPG was created in 1972 as a body to coordinate overlapping computer activities of members of different EPS Divisions. Most members of CPG are also members of at least one of the Divisions and, indeed, any member of a physics division can at no extra cost become a member of CPG. We encourage you to do so.

In the past, the main activity of CPG was the organization of conferences (eight were organized) on specific computational physics problems, but we realized at the last conference (on computational plasma physics) that interest in such an activity had gone down owing to computational physics having penetrated into all branches of physics. Computational x-physics is x-physics. We do not believe that such specialized conferences any longer have a future. More promising are interdisciplinary conferences and a first general interdisciplinary Computational Physics Conference is to be organized in Boston, in September 1989 by the Computational Physics Division of the American Physical Society (created last year). The idea is that computational x-physicists can mutually learn from computational y-physicists. It is proposed that the CPG-EPS organize the second such conference in autumn 1990 and that Japan takes over in 1991. This means that every third year such a general computational physics conference will be organized in Europe. We are looking for an organizer for the 1990 conference.

Another main activity of CPG was the organization of seven Summer Schools in Czechoslovakia and one in Switzerland. The themes of these courses are interdisciplinary subjects based on the use of computers for physicists. Typically, the last summer course (9-18 June 1987, Bechyne, Czechoslovakia) was on "Microcomputers in Physics", the one this year (5-9 September 1988 in Puidoux, Switzerland, see announcement on page 19) on "Numerical Methods for Parallel Vectorcomputers" and next year on "Man-Machine Communication in the Scientific Environment" (19-28 September 1989 in Skalsky dvur, CSSR). The main subjects to be treated in this last school will be logic programming, symbolic manipulations, scientific text processing and interfaces to CAD systems with applications. The manuscripts of the lectures are published in the journals Computer Physics Communications *and* Computer Physics Reports.

Besides these summer schools, the CPG Board proposes to start new activities in the following domains:
- *creation of electronic libraries (important articles on specific subjects, efficient subroutines for high speed computers)*
- *organization of an electronic newsletter*
- *electronic mail system management*
- *collect video cassettes on computer physics, courses, seminars and organize a service of distribution.*

To be able to realize these ideas we are looking for young interested physicists who will become members of the CPG Board. An election campaign is under way. We look forward to your proposals and thank you in anticipation.

R. Gruber, Chairman CPG

Personal Computing in a Physics Environment

Ian Willers, CERN

The Influence of Personal Computing

One year a student came to CERN and was given the task of writing a program that would automatically print labels that were to be placed on magnetic tapes. A suitable printer was purchased and attached to the IBM mainframe computer that runs the IBM central computing service in CERN. The student produced a program that he had laboriously proved to be correct. However, on running the program it instructed the printer to print labels with such rapidity that the printer gave up. The student then introduced some loops that merely consumed time and did no useful computing. The program worked for the first time. The program also consumed vast amounts of computing power and the student easily received that month's award as the prime user of computing power at CERN.

He visited the Systems Group who provided him with a routine that suspended the job for a given period of time during which it used no resources. Suspending the job however put him at the end of the job queue. When the IBM had only a few users the program ran perfectly. But when the number of users increased the period between the printing of labels got longer and longer. The solution was to use one of IBM's more humble products, the IBM-PC, which was happy to print labels at a constant rate with none of the complications of using the main IBM computer.

There are now many personal computers doing this type of work inside physics laboratories. However, a closer look shows the personal computer user preparing documents, doing accounts, writing small analysis programs with graphical output, controlling equipment in accelerators and so on. When the humble personal computer runs out of steam, its big brother the personal workstation takes over. The personal work-

Contents

Published by the European Physical Society: 10 issues per year. *ISSN 0531-7479*

station user is seen to be working on vast Fortran programs, once only possible on large central machines, powerful programs for Computer Aided Design, control of accelerators, very fast graphics routines, statistical analysis of experimental data and so on.

Why are people slowly but surely abandoning the security of a central service with powerful mainframes and preferring to manage and use a personal computer or workstation?

The advantages are many:

— There is power that you can rely on. The computing power of the personal machine is not shared with anyone else. If you run a program you know how long it will take; it is predictable. If you want one hour's time on your processor then you know it is yours. Many people with powerful workstations are running their programs overnight accomplishing major calculations which would take less time on the central service machines. But that user knows that the results will be there next morning whereas the probability of it running on the central machine depends on the other programs that are to be run that night. At peak hours the personal machine is as reliable in response time as it is at any time of the day or night.

— The machine is user-friendly. The friendliness of personal computers is well known. The graphical interface with windows where the screen is partitioned and menus which allow a particular operation to be chosen was commercially exploited on the Apple Macintosh and, as the users began to be accustomed to it, they liked it more and more. The personal workstations with their standard but older UNIX-type operating systems are not quite as friendly. They needed a standard multi-tasking virtual memory operating system which would support the most powerful applications. There was no choice, it had to be UNIX.

— The personal workstation excelled in the area of communications, with files being shared over very fast networks. The multi-taskings and the popularity of UNIX in the communications world gave them that advantage. The personal computer does not cost much more now than terminals did a few years ago. The physicist is able to remove his terminal and replace it with a personal computer. No extra desk space is needed and a terminal emulator enables him to continue accessing other computers. A program called Kermit (named after the frog in the TV series *The Muppets*) enables the personal computer user to transfer files slowly but surely between the central and the personal computers.

The major disadvantage to the user was the need to become his own system manager and his own operator. When the disks did not work, data was lost. When the personal computer broke down the user did not know what to do. When the printer ran out of paper, the poor user could be seen going on a raid to the central operations to steal the supplies that he lacked. In spite of all this he survived and even became proud of his new found self-sufficiency.

The Different Communities

In the general business world, the IBM personal computer and the IBM compatible personal computers are by far the most popular. In the physics community, the Apple Macintosh is the more poular with the IBM personal computer and its compatibles a close second. A third community is now growing rapidly which works with personal workstations. The Apollo offers support for large Fortran programs and was initially the most popular. VAXstations are now becoming popular because of their compatibility with the VAX machines. The SUN personal workstation also has a dedicated following in the physics community, particularly, when UNIX is an advantage.

The IBM personal computer was technically a disappointment when it was announced. However, the name of IBM caught the imagination of many people and a large cottage industry grew up to produce software in great quantity. The user is confronted by a large number of programs, in excess of 10000, for the IBM personal computer. This compares to a mere 2000 for the Macintosh. The industry quickly settled on PC-DOS (based on MicroSoft's MS-DOS) as the industry standard operating system for IBM and compatible personal computers. IBM and MicroSoft have now announced a new operating system called OS-2 which is designed to take advantage of the new processor chips in the new PS-2 range of personal computers. Windowing combined with multi-tasking will become as common place in the personal computer as it is in the personal workstation.

The Apple Macintosh pioneered user-friendly interfaces and compatibility between application programs. This level of compatibility meant that every program had a similar user interface and ability to communicate via a clipboard with other programs. Another concept which has been very successful is the concept of a single level menu system. Each menu item is either complete in itself or leads to a dialogue box which covers all further options. Within a short time after starting up a program the user is able to familiarize himself with the program and have a good idea of its range of functionality. The Macintosh also remained a closed box that should not be opened by the user. This meant that every Macintosh was identical while IBM personal computers were a mixture of boards from many vendors giving many special advantages and... problems.

The third group found the personal computers not powerful enough. However they saw the advantages of a good interface and they wanted good graphics. The added cost of a personal workstation was a deterrent but enough people took up the option that soon an enthusiastic community was built up. During 1987 the prices of personal workstations dropped dramatically. This coincided with the introduction by DEC, Digital Equipment Corporation, of their VAXstation range of personal workstations. The added competition coincided with developments in manufacturing which provided manufacturers with the ability to build powerful workstations using the same techniques as were used in the personal computer industry. A rather good personal workstation now costs a little more than a top of the range personal computer. Furthermore, the new personal computers at the high end of the range are becoming more expensive as they offer similar functionality to personal workstations.

Personal Workstations *versus* Central Computing

With the advent of cheaper personal workstations a new debate has begun — When is it best to use a personal workstation and when is it best to use the interactive central computer services?

Personal workstations are good for:

— Consistent central processor power – this is power you can rely upon;
— A good user interface and good graphics – facilities that make it easier for the physicist to concentrate on his problem;
— Good homogeneous networking – file sharing and servers giving printing or external functionality to all workstations on the local network.

Personal workstations are bad for:

— Big processing power – the central machines remain very powerful in comparison to today's personal workstations;
— Operations – the user is the operator of his machine and he would like access to central facilities such as printers and tape drives that are best managed by an operations staff;

—Peripherals – the quality and range of peripherals is still much better on the central computers.

As this argument continues there is a movement to let you have your cake and eat it too! Good networking to heterogeneous systems is at last within sight as *de facto* standards, such as NFS (Network File System for distributed file access), X-Windows as a windowing standard and Postscript as a standard output language for graphics, are being rapidly adopted by many suppliers. If all goes has planned —

File systems will be accessed using NFS. Applications will talk to windowing systems on workstations using X-Windows. All printers will accept Postscript format to print both text and graphics enabling users to print files wherever they want.

The progress in this area has been due to the various manufacturers coming to a consensus of opinion on the directions to follow. No longer do they wait for the standards organisations, they prefer to follow *de facto* but functional standards now.

Finally the debate is taking on larger proportions as the personal workstation manufacturers are announcing new machines of 10, 20, 30 and even 40 Mips of power. They are incorporating new RISC (Reduced Instruction Set Computers) processor chips with 10 Mips power each into multi-processor architectures. We are certainly looking at an exciting future.

Trends in Supercomputers

R. Gruber, Lausanne

(GASOV-EPFL)

The title of this report could also have been "Trends in Supercomputing". The reason is that powerful parallel vector computers can deliver a high computational performance only if the application software is adapted to their architectures. It becomes more and more evident that to solve the biggest problems, the software must be adapted to the computer architecture and *vice-versa*.

Present Situation

Supercomputer is a designation given to about 300 computers installed worldwide with a peak computing power of over 100 megaflops (million floating point operations per second). These machines have mainly been used for numerical experimentation in various scientific domains such as fluid dynamics, structural mechanics, seismic explorations, reservoir modelling, quantum mechanics, plasma physics, materials science. It is also believed that portfolio analysis and financial transaction business in banking will, in the near future, be executed on computers having many powerful processors and memories of the order of a few gigabytes.

Specific characteristics of supercomputers are high peak computing power achieved by very rapid clock periods (4.1 ns for a CRAY 2), parallel processors (65 536 in the highly parallel Connection Machine), pipeline architectures (instructions are divided into subinstructions which can all be in execution at the same time, so working like an assembly line) leading to up to two operations (1 add + 1 multiply) per clock period, large memories (more than 2 gigabytes for CRAY 2 and ETA 10) and very rapid connections to the outside world. The peak power and the memory space can be up to three orders of magnitude bigger than for workstations. The major flaws of these computers are that their operating systems and their mode of access do not yet match modern standards. In addition, to benefit most from the high computing power, it is necessary to formulate an application in such a way that at least 90% of all the computations are executed as vector operations. Unfortunately, this is not yet the case for most of the applications now running on supercomputers. This further implies that non-vectorizable organizatory work such as editing, interactive graphics and documentation, data handling and communication business should be taken care of by a user friendly, highly interactive personal workstation directly linked to the supercomputer via a high speed fibre optic network. At the moment, however, most of the supercomputers are accessible only through mainframes, thus reducing their attractiveness.

Scientist' Expectations

A scientist or engineer would like to be able to solve numerically the most realistic physical model described, for instance, by a set of partial differential equations (PDE). As an example, one would like to compute the time evolutionary solution of six coupled nonlinear PDE's in three dimensions. For a discrete approach (finite elements, finite differences or finite volumes) a discretization of 100 × 100 × 100 intervals is today considered necessary to obtain physically relevant results. To reduce the number of time steps, implicit methods should be used; hence the necessity of very efficient iterative solvers based upon the methods of multigrid or conjugent-gradient with preconditioning. To make the best use of the parallel architectures of future supercomputers, algorithms and programming techniques leading to the definition of codes with coarse parallel granularity should be adopted.

High parallelization at the subroutine level can be obtained by a decomposition of the geometrical domains into subdomains, each subdomain being assigned to one processor. To handle the connectivities between subdomains, fast networks built around high speed buses, global memories or direct connections of nearest processors are needed. An estimate for three-dimensional simulation programs shows the necessity for an effective computing power of 10 000 Mflops (then one run does not take more than one hour CPU time) and a memory space of eight gigabytes (all the