

Prof. G. Shaviv • Prof. M. Wolfshtein

# SUPER COMPUTERS in Israel

**Feasibility and Alternatives**



TECHNION — ISRAEL INSTITUTE OF TECHNOLOGY

**THE S. NEAMAN INSTITUTE FOR ADVANCED STUDIES IN SCIENCE & TECHNOLOGY**

**The S. Neaman Press**



הטכניון - מכון טכנולוגי לישראל  
TECHNION - ISRAEL INSTITUTE OF TECHNOLOGY

THE S. NEAMAN INSTITUTE  
FOR ADVANCED STUDIES IN SCIENCE & TECHNOLOGY

מוסד שמואל נאמן  
למחקר מתקדם במדע ובטכנולוגיה

---

# **SUPER COMPUTERS IN ISRAEL**

**A Study of Feasibility and Alternatives**

**Prof. G. Shaviv**  
**Prof. M. Wolfshtein**

**August 1988**

All opinions expressed and arguments employed in this publication are the responsibility of the authors and do not necessarily represent those of the S. Neaman Institute.

Copyright ©1988 by the Samuel Neaman Institute for Advanced Studies in Science and Technology. All rights reserved.

No part of this publication may be reproduced in any form without permission in writing from the publisher, except by a reviewer, who may quote brief passages in a review to be printed in a magazine or newspaper.

Printed in Israel by The S. Neaman Institute Press, September 1988

## Table of Contents

ABSTRACT	6
EXECUTIVE SUMMARY	7
1. INTRODUCTION	9
1.1. Background	9
1.2. Literature Review	10
1.3. Outline	11
2. GENERAL BACKGROUND	12
2.1. What are Super Computers?	12
2.2. Scalar, Vector and Parallel Machines	13
2.3. What Makes Super Computers so Fast?	17
2.4. Who Needs a Super Computer?	18
2.5. The Considerations in Selecting a Super Computer	20
2.5.1. Speed	20
2.5.2. Memory	21
2.5.3. Libraries and Applications	22
2.5.4. Growth in the Field	23
2.5.5. Ease of Use and Application	24
2.5.6. Code Conversion and User Migration	24
2.6. New Developments in the Field of Super Computing	25
2.7. Super Computer Performance Measurements	28

3. PRIORITY FOR INVESTMENT IN SUPER COMPUTERS	33
4. THE ISRAELI SITUATION	36
4.1. The Needs for Super Computers in Israel	36
4.1.1. Needs of the Academic Community	37
4.1.2. Needs of Government Institutions	39
4.1.3. Needs of the Private Sector	41
4.2. Classification of Potential Super Computer Users	41
4.3. Estimate of the Demand for Super Computers in Israeli Institutions	43
5. THE POSSIBLE SUPER COMPUTERS IN ISRAEL	48
5.1. Type of Machine	48
5.2. Some Notes on Vendor Selection	49
5.3. Additional Considerations	51
6. COMMUNICATIONS AND NATIONAL CENTERS	53
6.1. The Relation between Communication and Super Computing	53
6.1.1. Interactive Work	53
6.1.2. Interactive Graphics	53
6.1.3. Super Computer Networks	54
6.2. The Economics of Local and National Super Computing	55
6.3. Price of Communication in Israel	59
6.4. Can Israel have a National Super Computer Center?	60
6.5. A Little More on the Financial Side	63
6.6. The Case for a Distributed Israeli Super Computer System	65

7. USER ENVIRONMENT	66
7.1. Connectivity	66
7.2. Operating System and Utilities	67
7.3. Consulting and Instruction	69
7.4. User Migration	70
8. DISCUSSION AND CONCLUSIONS	71
8.1. The Necessity for Super Computers in Israel	71
8.2. Cooperation in a Distributed System	74
8.3. The Case for some Local Centers	76
8.4. Conclusions	76
APPENDIX A: Evaluating Computer System Performance	78
APPENDIX B: Costs of Computer Communication in Israel	81

### ABSTRACT

This document discusses the need for super computers in Israel and possible ways to install such equipment. An examination of the current situation in Israel is presented. In addition, the concept of a national central facility for super computing is compared with that of distributed local centers. The report is concerned with the case of Israel in general and the Technion in particular.

## EXECUTIVE SUMMARY

1) Super computers are among the drivers of the forefront of technology. Indeed, the lack of super computers in Israel impedes the progress of science and technology of this country.

2) Super computers should be placed at the top of the priority list of multi-user equipment required by the academia in Israel.

3) A national center for super computing with T1 communication lines is the best solution on condition that the price and size of the central facility is above a critical minimum. Our estimates, based on expected funds allocated to a super computer on the one hand, and the demand for super computer power on the other, are that this solution is too expensive.

4) The preferred solution at the present demand for super computer power is a small super computer on each campus.

5) It is our conviction that the super computer operating system must be UNIX.

6) At present, vector technology (hardware and software) is more developed than the parallel one. We are convinced, however, that the future belongs to parallel vector machines.

7) At present vector machines should be preferred to parallel ones. However, options of parallel processing in general, and in available machines in particular, should not be ignored.

8) Consulting services are extremely important and any new super computer must be supported and supplemented by proper consulting services.



9) Vector technology, parallelism and problem solving on vector and parallel machines should be included in the curriculum in computer sciences, and science and engineering departments in Israeli universities.

## 1. INTRODUCTION

### 1.1 Background

The application of parallel computing technology to science and engineering started about 15 years ago with the introduction of the first "vector" or "pipe-line" machines. However, the first machines were very expensive and their operation rather difficult. Thus, only the central advanced research centers could afford them and mastered the expertise to operate them. Slow progress followed during the late seventies and the early eighties. Yet vector computers (or super computers as they came to be known) did not spread to real industrial applications until the late 1980's. When this happened the arena changed in a grand way. Suddenly super computers became widespread and indeed essential for industrial research and development. Presently any country striving for high tech industry can not afford not to have such a machine.

To date, there are several hundred super computers in use throughout the world, most of them in the USA, but also in other western countries and Japan. With annual sales of more than 50 machines this has become an attractive product and there are now about ten manufacturers who compete in this market world wide. As usual competition brought about price cuts which became possible by improved technology and larger quantities. Thus the late 1980's can be characterized by two main trends in this field: Firstly, these machines have become a necessity for many real life applications where their use can be economically justified and secondly, their price reduces steadily. Today we find ourselves in a critical period; a nation aspiring to any kind of leadership or excellence in science and technology cannot afford to neglect super computing. Israel is no exception.

With this background in mind it should be recognized that any developed country will require super computers within the next decade. Countries which do not acquire this new technology in the near future endanger their position in the high tech marketplace and reduce their potential for manufacturing advanced products. This is a great change from the past when it was argued that small countries could benefit by neglecting the field of super computing, and many small countries actually did so.

## 1.2. Literature Review

We have used two major documented sources of information for the preparation of this report. The first is "Performance of Various Computers Using Standard Linear Equations Software in a Fortran Environment" by Jack J. Dongarra, Technical Memorandum No. 23, Argonne National Laboratory, Illinois, February 29, 1988. This report contains the widest selection of test runs of a single program on many computers, ranging from small personal ones to very large vector computers.

The second source is the report "Technion Supercomputer Benchmark" by J. Arnon and C. Weil, Taub Computer Center, Technion, June 1988. This report describes the benchmark and results prepared by the Technion Computer Center. The benchmark consisted of nine programs donated by a number of heavy Technion users. The results of this report are classified, and we and therefore have used only general conclusions reached from it, without quoting quote the numbers.

### 1.3. Outline

The report is intended for a variety of users, with varying degrees of expertise on the subject. It is therefore possible, and may even be desirable for many users to read only parts of the report. This outline should help such readers select the sections of particular interest to them.

Chapter 2 contains general introductory notes on super computing. Readers with experience in the field may skip this chapter. Chapter 3 outlines the reasons for our conviction that the purchase of super computers by Israeli institutions is justified at the present time. Local Israeli considerations are discussed in chapters 4 and 5. The problem of the impact of communication on super computers is discussed in chapter 6. In particular, section 6.5 outlines the reasons for which we believe that a distributed system of super computers is the preferred choice for Israel at the present time. Finally, in chapter 7, we outline the requirements for making a super computer easily accessible to users. This chapter does not refer to the problem of choosing a machine but to the problem of how to get the most out of it. Chapter 8 contains the final discussion and conclusions.

Readers who wish to get only the essence of the report may read only the executive summary, and chapters 3 and 8.

## 2. GENERAL BACKGROUND

### 2.1. What are Super Computers?

The term super computer is used today to loosely describe computers that can perform many calculations very quickly. Some computer scientists use the following definition: Super computers are the biggest and fastest machines available at a given time. This definition is an ambiguous one: Clearly it changes every time a new generation of big computers appears. As time progresses the development of new technologies and new logic and their embodiment in computer architecture lead to a speed-up in super computer performance (an elaboration of this point is provided later in the report). Another weakness of this definition is that it does not reflect either the tremendous impact which super computers (whatever their definition may be) had on the emergence of new computer logic, mathematics and architecture or the propagation of the influence of super computing on thinking in all other scientific disciplines. In this report we shall adopt the following definition: Super computers are machines which are at least ten times faster than typical Israeli campus mainframe computers (e.g. IBM 3081D or CDC CYBER 180/850).

Super computers have become an important part of the modern technological and scientific environment. Nearly 400 such machines have been installed worldwide in a variety of government, academic and industrial institutions; they are used for many practical applications like weather forecasting, automobile and aircraft design, simulations and graphics, stress analysis of complicated structures, bio-medical engineering, the design of new drugs in the pharmaceutical industry etc. Basic research benefits from super computers as well: The use of super computers enabled physicists and chemists to explore new

problems and widen their scientific horizon to the extent that the new fields of computational physics and chemistry are now well recognized and established disciplines. Super computers have penetrated even into exotic applications such as economics, the social sciences and applied art.

## 2.2 Scalar, Vector and Parallel Machines

To better clarify the necessity of moving to super computer technology let us elaborate for a while on computer technology and architecture. The first computers were what we call today scalar machines. A scalar machine is characterized by performing a single operation on a single operand. Thus, to multiply 100 numbers by 100 numbers a scalar machine has to perform 100 multiplications (obviously) and also 100 instructions. The instructions must be fetched from the memory and brought to the program register where they are deciphered and the operands must be brought to the arithmetic units where all the arithmetic operations are carried out. In short, the time of the Central Processing Unit (CPU) goes to collect the data and perform the multiplication on one hand and to interpret the instruction on the other.

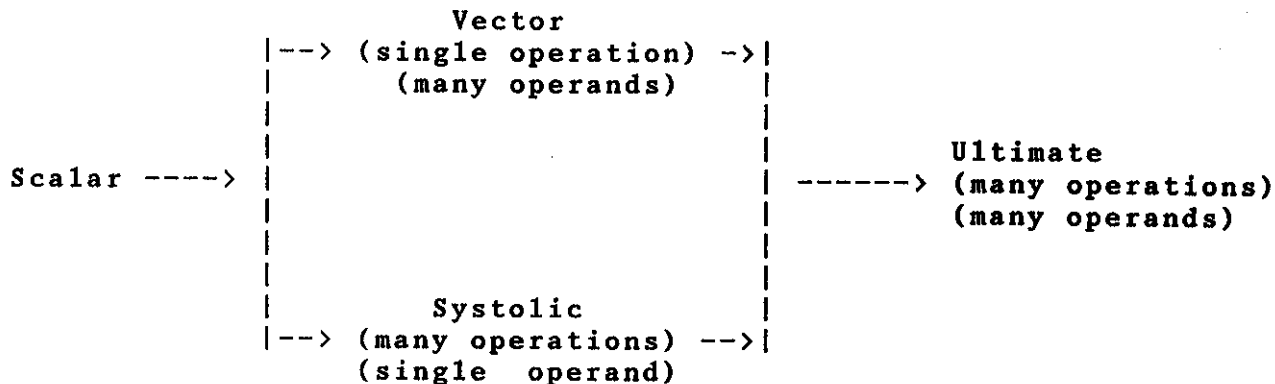
Vector machines are distinguished by being able to perform the same operation on  $N \gg 1$  operands by means of a single command - a vector command. The setup time required for the execution of a vector operation is longer than the setup time required for a simple scalar operation. However, if the vector is sufficiently long the time required for the entire calculation is reduced by a huge factor. Vector machines introduced therefore a new concept - a vector command which has many ramifications on our thinking and the way we pose problems. Today most super computers utilize vector instructions using

special hardware devices and it is one of the dominant factors in giving these machines their enormous speed.

Yet another possibility is the systolic machine in which a single instruction creates a sequence of operations on few data or operands. For example, if  $a, b$ , and  $c$  are three operands then  $a+b*c$  can be a single instruction which is performed by hardware. Thus time is saved by creating groups of the most frequently used instructions on a set of operands. Such computers are called systolic computers since they 'shortcut' the instruction. Other examples are machines with many built in functions like  $\sin$ ,  $\cos$  and the exponential functions which exist in hardware and not in software.

Not all problems can be formulated in a way amenable to vector instructions. Moreover, even programs which are highly vectorizable have a scalar part which cannot be vectorized. Finally, recursive calculations pose a serious problem to vector machines. A systolic computer is the best one in such cases.

The ultimate computer will be the one that combines the properties of the vector machines (many operands) and the systolic one (many operations) and will allow the user to perform many operations on many operands by means of a single computer instruction. We see therefore a clear pattern in computer logic as expressed in the following diagram:



So far all the computers discussed here had a single processor. What about several processors doing a job in parallel? Before discussing this approach we would like to clarify why parallel processing will be a viable solution in the future. The cost of reducing the basic cycle time of the computer (the basic switching time) increases very quickly as the cycle time decreases. On the other hand, the cost of the lower speed CPUs decreases as the technology used for their manufacturing ages. Thus N CPUs are cheaper than a single CPU working at N times the speed of each of the first CPUs. Clearly parallelism is a way to achieve super computer power cheaply. Moreover, many consider parallelism as the only solution to progress in computer speed in a few years time. Once the speed of a single CPU gets down to the range of 1-2 ns cycle time further progress will be prohibitively difficult and expensive.

Parallelism is not necessarily tied up with vectorization. Indeed we see today many machines with simple scalar instruction sets but with several CPUs which can perform calculations at a rate comparable to that of vector machines. The ALLIANT may have 8-16 and the MYRIAS may have 64-512 CPUs. The hypercube computer contains 64,000 CPUs. Still,



the combination of vectorization with parallelism is the most powerful combination today. The large ETA and CRAY machines which may have up to 4 and 8 processors are typical examples.

In view of the high expectation from vector/parallel machines the obvious question is why buy a vector machine today if we know that the future is elsewhere? First the solution will still be many vector CPUs and not scalar ones. Secondly, while the state of the art of vector computing is now growing out of infancy, parallel computing is just being born. There are many computer architectures leading to parallelism and the industry is still looking for the optimal solution. Furthermore, the mathematics, the logics and the algorithms for parallel computations are very underdeveloped.

And what will happen in the next step? Probably clusters of machines and distributed load. However, this architecture carries us too far into the future, hence we stop here.

As can be seen from this survey of computer logic and architecture, the new innovations in computer architecture lead to new ideas, solution methods and thinking in a variety of scientific disciplines.

Before going any further we have to clear the question of array processors. Typical array processors are special units which connect to a conventional computer and can perform vector operations (the CSPI machine is a good example). Such machines appear to offer financial advantages to small institutions. Yet they suffer from an inherent disadvantage as they usually require special programming commands rendering their usage cumbersome. Moreover, their financial appeal is nearly gone now, as the price performance (number of calculations per say 1\$ of cost) is not better than that of

present super computers. Consequently array processors are used today mostly for very special applications and seldom as multipurpose machines.

### 2.3 What Makes Super Computers so Fast?

Several factors combine to yield the great speed of super computers:

1) Very fast cycle time. Cycle time is the shortest period of time to exist in a computer. It is the time between two clock ticks. The new machines have a clock time in the range of 5-25ns.

2) Vector instructions. The meaning was discussed above.

3) Pipelining. The CPU can be considered as composed of several elements each performing a specific job on the operands. Once the element finished a job, the result is transferred to the next element. If a given operation requires say N elements the execution time will be N times the number of clock cycles required for each element. Before an element gets the operand and after it has done its job, it stays idle waiting for the entire operation to finish. When a long sequence of the same series of operations exists, we use pipelining. The idea is to load each element with the next operand as soon as it finishes its job on the previous one. Thus, if in a scalar machine one obtains a result every N cycles, in a piping machine there is a short setup time after which the results appear after each cycle.

4) Chaining. If several operations are to be performed on a single vector, a lot of time can be saved if the vectors are kept in the arithmetic unit and not returned to the memory to be

fetched from there. This is achieved if the output register can serve as input register to the next instruction. To a large extent chaining means the ability to perform several operations on the same operand.

5) Overlapping. In principle two independent operations can be allowed to overlap in time. Suppose you have an addition which follows an independent multiplication, then the execution of the second operation can begin one clock cycle after the first operation has started.

Clearly, vector and parallel programming differ markedly from conventional scalar programming. It is often necessary to apply completely new algorithms and to restructure the entire program and its data if full advantage of the super computer is to be achieved. Thus, the art of heavy computing and the solution of complicated problems is developing into a new discipline requiring the reformulation of problems in an appropriate form so as to be prepared for solution on these new machines. The complete implementation of these details takes learning and time. All these new features and ideas must be explained to the users by a group of qualified consultants.

#### 2.4. Who Needs a Super Computer?

Super computers are newcomers to the scientific arena. The first such machines were introduced about 15 years ago, but their price was so high and their utilization required such special skills that their spread was rather slow. Even today these machines require special numerics and programming skills in order to take full advantage of their tremendous potential.

On the other hand, super computers allow us to approach large problems which could not be dealt with by the much slower

scalar computers. Therefore they attracted the attention of scientists and engineers following their initial appearance on the market. Indeed the advent of super computers has contributed to the formation of new kinds of scientific disciplines such as 'computer physics', 'computer chemistry', 'computational fluid dynamics' and so on. Today, super computers are applied to two kinds of situations: On one hand they are used as tools for the advanced analysis of very complex phenomena for applications like aerodynamic design, crash analysis of automobiles, weather forecasting, etc. On the other hand they also serve as scientific tools that allow the investigation of situations which are not amenable to usual experimental practice like the long term behavior of the planetary system under changes of say the ozone in the atmosphere, molecular structure of complex molecules, and the structure of turbulent flows, to name only a few.

It may be stated that whenever an interaction between different processes occurs, or even when the interaction between various quantities is important, regular mainframes are too small to produce an answer with sufficient resolution and accuracy. Thus, the super computer is turning quickly into an important scientific tool.

The technological importance of the super computer is very significant as well mainly due to the new possibilities which it opens for engineering design. Aeronautical system design may serve as a good example: Typical aeronautical systems are so complicated that they cannot be analyzed in full on a regular mainframe. The advent of super computers has created for the first time the possibility to 'put the entire plane' into the computer and analyze it as a whole unit. This was never possible before and the design of such systems had to be performed in sections or regions, later to be matched to one another through trial and error. Now it is possible to design such systems as a

whole unit and to avoid the expensive and uncertain iterative trial and error process. In general we may refer to 'the computer design revolution'.

So who needs super computers? Scientists in the various fields need super computers to perform their research, and engineers need super computers in order to avail themselves of the technological applications which are quickly developing. There is no doubt that the practicing engineer in industry is soon to feel the impact of the new design tools as well as the results discovered by basic research on his daily life. Thus it becomes important to use super computers in science and engineering, and to include the art of super computing in the general scientific and engineering curricula.

## 2.5. The Considerations in Selecting a Super Computer

The following factors determine the best choice of a computer:

### 2.5.1. Speed

The speed of the computation depends on hardware and software. Typical hardware features are the computer architecture, the technology of the CPU (affecting the intrinsic cycle time), the size of memory available to a single user, disc access time etc. Typical software affecting efficiency are the operating system and the high level language employed by the users. The effect of the different factors on the actual performance of the machine depends on the particular application; different machines will perform differently on the same application. For this reason we have to determine what is the typical application of such a machine at the Technion and in Israel. This is not a trivial question because most of the usage of the machine will develop only after it had been

installed and made available to the research and development community for a long enough period. This problem is usually tackled by running benchmarks on the computers to be studied.

### 2.5.2. Memory

Super computers come with different types and sizes of memory. Fixed small, fixed large and virtual memory. Virtual memory is the ability of the machine and the operating system to address virtual, non-existing, cells by swapping data between the memory and the disc space. When the programs are very large, or the data and tables required consume a lot of space, having a virtual memory is an advantage from the point of view of the user. Let it be clear, virtual memory is an overhead on the system and under obvious conditions may stall the machine - the machine will spend its time on swapping rather than calculating.

Yet it is a necessity for the user who wants to run a large program on a machine with a small memory. If virtual memory is not available the user has to instruct the program manually how to link, when and which data to dump to the disc etc. - a tedious and frustrating job.

Moreover, since super computers are to be used for large problems the need for large memory is exacerbated. Let us give a concrete example. Consider a 3D time dependent flow problem. There are about 10 variables per grid point. A grid size of 100x100x100 will need therefore 10MW for the main variables and about twice as many for auxiliary variables. This is a very large memory which only few super computers can control. In summary, if the super computer is to be used for bigger problems it must come with a large memory. The example given above demonstrates the limitations on the size of a problem that can be run on a fixed memory machine. As for virtual memory

machines the limitations come from minimum efficiency and overhead.

Large memory is important for an additional reason. In many cases the same calculation is repeated many times (for example the equation of state of the gas in a multiphase flow). In many cases a very significant factor in time saving can be obtained by storing the table of the equation of state in memory and not repeating the calculation time and again. This solution is possible only if real memory is available.

### 2.5.3. Libraries and Applications

One can not overestimate the importance of software libraries and packages to the scientific or technological computer. Such software allows big savings in programmer time (large parts of the algorithm are already available) and improves the quality of the programming (libraries and packages are often programmed in a more efficient way than the user has the time or the desire for). The existence of such libraries is imperative and the larger the variety the better. In many research projects as well as engineering classes libraries and application packages are required. For example, the analysis of structures (in civil, mechanical or aeronautical engineering) need finite element packages like NASTRAN or ANSYS to analyze large and complicated structures. Packages like FIDAP or PHOENICS are used to solve problems in fluid mechanics. Chemists require an application called GAUSS to calculate the structure of molecules and so on. It is often the case that application packages are more readily available on machines which are more popular. The implication of this consideration is that a popular machine may be a better choice than a machine which is scarce, even if the second machine possesses some better features.

#### 2.5.4. Growth in the Field

The size of the computer and the money invested in the first machine should be adequate for the needs of the community at the time of its purchase. However, a natural growth in demand for super computer time may be safely assumed. Thus it is important to allow field upgrading of the computer so as to provide the best services when required. Professionally this process is called 'growth in the field' capability. We note that the rate of growth in demand for computer resources is greater than even the rate at which new computer technologies emerge and as a consequence an upgrade may be required before a new computer generation becomes available (we remark parenthetically that the rate of growth at the Technion was about 50% per year over the past years). The growth in the field capability is also important from an economic point of view. When the computer starts operating it will not be used at full capacity. The larger the machine the longer it will take to completely keep it busy. Thus it makes good sense if the computer is adjustable to present needs, with a possibility to upgrade only when the needs justify it.

Past history of computer development suggests that one or at most two steps of upgrading is the most that is required before new hardware technology appears. When this happens both initial cost and equipment maintenance become much cheaper, and replacement of the old equipment often becomes more attractive than a field upgrade of the old machine.



### 2.5.5. Ease of Use and Application

The application of a super computer to any given problem is far from trivial and the ease of application is a crucial point for consideration. The same applies to utilization of existing equipment like terminals, plotters, network etc., with the new machine. Scientists often resent spending time fighting new operating systems, or learning how to read and write files from discs or how to get through the network from a terminal to the destination machine. Consequently, they will refrain from using a machine with cumbersome and unfriendly operating systems or utility software.

### 2.5.6. Code Conversion and User Migration

One of the crucial questions is user migration to the new facility: how quickly will the machine be used at full load and whether only young people will use the machine? Needless to say, the changeover is worthwhile only if vectorization is implemented because the scalar speed of the machine will not be sufficient to produce a breakthrough in performance. We therefore classify future users into two categories: 1) Those who MUST use a super computer in order to reach new frontiers and 2) those who will continue to run their present programs but may want to do so more efficiently or may want to extend the scope of their present programs. Needless to stress that the prime reason for investing in a super computer is the first group but we cannot ignore the existence of the second.

It is very difficult to estimate the speed of user migration. It is clear that smaller users, those who use few seconds of CPU time can stay on the old machine since they are not going to benefit from the move to the super computer. It is the heavy users with whom we are concerned. In general we expect

these heavy users to be aware of the running time. Also, we expect them to be well informed and knowledgeable of computer architecture in order to rapidly modify the program without massive help. There are certainly exceptions. Some scientists may consider the effort and time required for code conversion to be prohibitive and will prefer not to do it. An examination of other universities confirms this suspicion. Yet, the results of the Technion benchmark described in section 7.4 below suggest that the problem can be negotiated if sufficient programming support is available.

A similar effort is required in some programs whenever the operating system is changed. For this reason the scientific world converges towards UNIX as a universally accepted operating system with painless migration between machines and vendors. Therefore UNIX is the preferred choice for an operating system for super computers.

## 2.6. New developments in the Field of Super Computing

Until recently the situation in the field of super computing was as follows: At the top of the line were machines like the CRAY XMP and CDC CYBER 205. These were very fast machines starting at a price of M\$7-10 for the smallest configuration. Such machines are not only beyond the financial means of any single Israeli university, but also beyond the means of most foreign universities (in the USA and elsewhere). Moreover, the capacity of such machines is so much higher than the demand of the typical university campus that it is not likely that most universities will keep such a machine busy from the very beginning. Against this background the simple answer was to create big computing centers of national (or at least regional) super computers. Typical examples are the USA and Britain. Obviously, such an arrangement requires a reliable and fast

communication system between the super computer centers and the users. As the price of fast communication is high, compromises were often made and the communication system was not fast enough to allow the remote user to utilize the machine for interactive applications. This has had severe repercussions on post processing of large result files. Thus the remote user was put at a major disadvantage in comparison to the in-house user. This mode of operation is far from ideal, but it was the only path opened to most academic researchers using the super computers.

The above situation left most number crunching users at a major disadvantage. The usual mainframe machines were not fast enough, while super computers were too expensive to install locally and very unfriendly to use in a remote site. This gap was filled in several ways.

Firstly, several small manufacturers developed highly efficient add-ons usually referred to as array processors (e.g. the FPS or CSPI machines). These units often used advanced ideas like the pipe line concept and could therefore run a single job very quickly. As these were basically single user CPUs with some memory only, and all the other facilities were supplied by the host computer their price was much cheaper than that of a real super computer, yet they offered significant improvement in performance in many cases. However these were generally single user machines, and they required special programming skills.

Meanwhile mainframe manufacturers developed faster scalar machines and following the success of array processors some of them developed an add on vector facility for their scalar machine (e.g., the IBM 3090 and the CDC CYBER 990). These machines are programmed with standard languages, and the add

on vector facility is cheap. Thus they fill a gap, but they do not really resolve the problem, as they used to be much slower than super computers. Later IBM extended the power of the 3090 series by using parallel architecture. The large 3090 machines are now super computers but they are not cheap anymore.

Simultaneously with the above efforts, some small manufacturers started the development of the air cooled mini super computers (e.g. CONVEX or ALLIANT). Typically the price of such machines is below M\$1, and their maintenance and operation were very simple. With these qualities they were within the reach of universities, and even large departments. Such machines are also very user friendly and therefore they were attractive to many users. These machines are not necessarily more cost effective than the bigger machines but they offer a solution to a certain class of users.

Recently some important developments have taken place in this segment of the market. Convex announced its second model, the C2, which is about 3 times faster than the C1. More over, the machine can run up to 4 CPUs in a true parallel fashion, thus reaching the low end of the real super computer range, while still selling for about M\$1.5. The Control Data Corporation announced recently a new series of super computers replacing its unpopular CYBER 205 by the new ETA series. While the large ETA 10G8 is designed to compete with the largest CRAY machines, the small air cooled ETA 10P1 sells for about M\$1.5, and as such it is designed to compete in the mini super computer market, although it has a stronger CPU power. Thus CDC offers now a line of machines spanning the entire range of machines from the Convex C2 to the new CRAY YMP. Finally CRAY Research announced its so called baby CRAY, the X/MP14se, which sells for about M\$3. While this machine is still more expensive than

the air cooled machines, it offers the reliability and good properties of the CRAY X/MP series, and is a reasonable alternative for a campus computer as well.

The advent of machines like the ETA line, the larger Convex machines and the baby CRAY on one hand and the IBM 3090 on the other hand may be expected to bring about large changes in the super computer arena. Many universities will be able to purchase the cheaper machines and indeed some have already announced their intention to do so. Thus the accessibility of super computing will increase immensely. Under these new developments the problem of national, versus distributed super computing is very different. As super computing is now within the means of individual universities the idea of national centers loses much of its previous attractiveness.

## 2.7 Super Computer Performance Measurement

Performance comparisons between different computers are not easy. Firstly because computers are highly modular and no two computers are exactly identical. The same nominal machine may have a different number of CPUs, different memory size, different communication channels with the disks, and different disks. On top of this the quality and type of system and user software have a great influence on the performance. Last but not least are the units of performance. As different computers are performing the same operations in entirely different ways even the definition of performance units may have a large effect on the results.

The performance of super computers is frequently measured in MFLOPS (Millions of Floating Point Operations per Second). MFLOPS rates include all types of operations while the time required by the computer for the different operations varies.

Thus MFLOPS rates may be used only to compare the performance of the various machines on the same program. We must also distinguish between the peak MFLOPS rate according to the manufacturer's specification and the real MFLOPS rate measured in actual computations. (Some people define the peak MFLOPS as the speed guaranteed by the manufacturer that the machine will never surpass).

An important factor is the size of the matrices which are handled. Both vector and parallel computations can be performed only on matrices (yet this does not imply that any matrix operation can be vectorized or parallelized). If the matrices are too small the machine efficiency deteriorates very sharply. And often the machine works better on matrices of certain sizes (e.g. powers of 2). Another factor is the quality of the software. The compilers on super computers can vectorize and parallelize automatically but the efficiency of the compiled code depends on the compiler. Moreover, certain algorithms and programming styles may be vectorized or parallelized better than others. Therefore, the MFLOPS rate depends not only on the machine used but also on the problem, the level of programming skills and the compiler used. Thus it is not surprising that efficient utilization of the CPU power enforces such tight requirements on the programmer, that the theoretical peak performance can never be achieved.

The above mentioned difficulty in performance measurements really calls for local benchmarking for the comparison of super computers. Yet this task is often misleading as the programs available for benchmarking are existing programs which were not written for vector or parallel processing. Consequently they can not utilize the machines at high efficiency, and do not represent real future benefits to be gained when the level of programming skills is adjusted to

the requirements of super computing. A partial solution to the problem is using a general benchmark. Some widely used packages have been developed for this purpose and serve as general yardsticks for the comparison of computers. The most well known ones are the Dongarra LINPAK loops performing various operations in linear algebra on large dense matrices, and the Livermore loops in which some mix of typical research programs is incorporated.

Typically the speed of the common university mainframes like those found in the Israeli universities is of the order of 1 MFLOPS. The largest super computers are designed to reach 1000 to 10,000 MFLOPS. However, they seldom reach speeds of more than 10-100 MFLOPS in computation of real problems. Table 1 provides some examples from Dongarra's Argonne National Laboratory Technical Report No. 23 on the "performance of various computers using standard linear equations software in a FORTRAN environment", 1987. These runs are for various 64 bit operations in linear algebra on a matrix of 100x100 and without any programmer intervention in the original programs. The results are displayed in groups, namely: Super computers; mini super computers which are small vector machines; Array processors which are powerful additions to small computers but not usually capable of running alone; Mainframes including super mini computers and work stations. The table represents the state of the art at the time this report was written.

TABLE 1

Computer	Proc	ns	Operating System/Compiler	MFLOPS
<b>Super computers</b>				
ETA 10-E	1	10.5	ETA V/FTN200(Rolled BLAS)	62
CRAY X-MP/4	1	8.5	CFT77 2.1(Rolled BLAS)	56
CRAY X-MP/14se			CFT77 2.0(Rolled BLAS)	31
ETA 10-P	1	24	ETAV/FTN200(Rolled BLAS)	27
CRAY 2S	1		CFT77 2.0(Rolled BLAS)	23
CDC CYBER 205	2 pipes		FTN(Rolled BLAS)	17
IBM3090/180E VF			VS 2.1.1 OPT=3(Rolled BLAS)	13
CONVEX C-210	1		FORTRAN 4.0(Rolled BLAS)	10
<b>Mini super computers</b>				
ALLIANT FX/8(8CEs)			FX V2.0.19(Rolled BLAS)	7.6
CONVEX C-130			FORTRAN 4.0(Rolled BLAS)	7.3
<b>Array Processors</b>				
FPS-264 (M64/60)			F02 APFTN64 OPT=4(Rolled BLAS)	5.9
CSPI MAP-6430			FORTRAN 1.5.35	1.2



Mainframes

IBM3081D (present Technion mainframe)	VS opt=3	1.7
IBM4381-12 (Technion administration machine)	VS 1.4.0 opt=3	0.95
CDC CYBER 170-835 (Technion CAD machine)	FTN 5 opt=2	0.47
VAX 785 FPA (Technion typical VAX)	VMS V4.5	0.21

---

Workstations

SUN-4/260	f77-0 sysy4-beta2	1.1
SUN-3/260+FPA	3.2 f77-0-ffpa	0.46
microVAX 3200/3500/3600	VMS V4.6	0.41
SUN-3/160+FPA	3.2 f77-0-ffpa	0.40
IRIS 2400 turbo/FPA	f77	0.24
MICROVAX II	VMS V4.5	0.13
SUN-3/50 16.7MHz 68881	3.2 f77-0-f68881	0.087
Apple Mac/Levco Prodigy	ABSOFTE Mac Fort 020	0.076
IBM PC-AT/370	VS FORTRAN opt=3	0.033
IBM PC-XT/370	H opt=3	0.031

---

This table is given only as a guideline for orientation. Actual performance depends on the particular application, and is discussed in greater detail later on.

### 3. PRIORITY FOR INVESTMENT IN SUPER COMPUTING

Even after the recent improvement in price performance of super computers and the availability of cheaper super computers the machines are not cheap. Even a small super computer would cost M\$1-2, and its annual running and maintenance bill may well exceed K\$300-500. Therefore universities should ask themselves whether investment in super computers is the best one they can make in terms of quantity and quality of research and teaching.

The question is not easy to answer, mainly because the quantification of research and teaching does not lend itself easily to objective comparative evaluation. Still, there are some facts which should be taken into consideration when looking into this problem as follows:

1. Wide applicability: Computer equipment in general is required for the work of almost any academic unit. Therefore investment in computers is going to improve the working conditions of a very large group of members of the academic community.

2. Balanced loading of computers: It is usually the case that most of the computer users require small chunks of the computer time, and a relatively small number of heavy users swallow the lion's share of the time, thus interfering with the rest of the community. By giving these researchers a super computer option they are going to leave the general mainframe, freeing much desired CPU time for general non CPU bound usage.

3. New frontiers: The acquisition of a super computer is more than just adding more computing power. Rather it opens a

door to an entirely different type of research which can not be done without these machines. For instance certain three dimensional phenomena in fluid mechanics or physics, studies of molecular structure, and so on. The acquisition of a super computer is not only a quantitative leap but also a qualitative one, and academic institutions without super computers are likely to find themselves forced to abandon these newly developing disciplines.

4. Vector and parallel processing are themselves new disciplines: They are likely to be very important not only in super computers but also in smaller machines. While research into these disciplines can be performed using smaller machines, the availability of super computers may boost these fields as well.

5. Super computers are finding more and more practical applications: Industry has to face this challenge in order to survive in a very competitive international market place. The role of universities in this context is to give technological leadership by research and teaching in super computers. By neglecting this field the universities are abandoning their moral responsibility to the community.

In view of the above factors we reach the conclusion that investment in super computing is not only an obligation which universities should answer but is also one of the most attractive investments available to the academic community nowadays in terms of academic throughput.

These arguments are valid throughout the developed nations. Super computers are probably more abundant in Germany and Japan than anywhere else. In the USA the NSF established five very large super computer centers, and a fast (T1) communication

network between them. A large number of super computers were installed in Western Europe too. Regional organizations of super computer users have formed in all these areas, and international joint research programs have been defined. Altogether it appears that the international academic community has opted for super computers.

As for Israel, it is the opinion of the authors of this document, that we too have to act now. The desire to do research along the frontiers of science, to support the local industry by training young graduates in state of the art methods, and to allocate funds so as to maximize their effect all lead to the conclusion that Israeli universities need super computers NOW.

#### 4. THE ISRAELI SITUATION

##### 4.1. The Need for Super Computers in Israel

Over the last 2-3 years super computers have become one of the most important tools in many branches of science, technology and the arts. The impact of these machines on science and technology is tremendous.

Most Israeli universities do not have as yet machines in the super computer class, although the CDC CYBER 990 and the IBM 3090 computers installed at Tel Aviv University last year are approaching the super computer range.

This is a serious deficiency: It is necessary for any community aspiring for advanced R&D to employ super computers in order to keep pace with the scientific and technological developments in the rest of the world. Thus, if the state of Israel is to have a place among the developed nations of the world it should acquire super computer technology. We believe that a change in computer speed (and technology) in Israel by a factor of 10-100 is necessary in order to bring the Israeli R&D community to a level of computational technology comparable to that of the industrial nations. In such conditions Israeli researchers will be able to compete with their colleagues abroad. Even if we cannot purchase the best computer and must be content with a less powerful machine - it is a must to move to the new technology - or else we remain way behind.

We have to stress that this is a very particular point in time. Universal spreading of super computers has just started, and we can still bridge the existing gap in super computing between Israel and the advanced nations of the world. The boat of super computers is now leaving the harbour. We can still

embark on it. However, if we delay our decision we shall stay on shore and hear how super computers were used in faraway lands to advance knowledge in both technology and science without our participation or involvement.

Let it be understood and said that in science it is the first who is quoted and remembered. Being the second to discover does not count. And in technology being second often means never.

When considering future uses of super computers it is convenient to divide the users of super computers into three categories: academic, government and private.

#### 4.1.1. Needs of the Academic Community

The academic community has to answer the needs of basic and applied research as well as those of teaching. Research needs are of two kinds: 1) The new disciplines of computational sciences often require super computers for regular work. This is the case also for the new subjects of parallel and vector computations. 2) The utilization of super computers together with available software packages for studies of problems using simulation of experimental apparatus. In both areas Israel is in poor shape and much time has already been lost. The spread of super computers in the US began about 15 years ago. In Japan and Western Europe it started about 5-10 years ago. If the Israeli scientists do not have access to super computers soon, important disciplines will die out in Israel for lack of a capability to perform advanced work.

The following table provides interesting comparisons between computer resources in Israel and in two countries in the western world. We can of course compare Israel to underdeveloped countries but the question is if this is desirable. Should we

not strive to reach the top? The table provides information relating to the number of universities and institutes, the number of students and the installed super computer power available to them evaluated in GFLOPS .

Table 2

Comparison in Super Computer Power

	Japan	USA	Israel
Universities	460	3000	7
Number of students	2.0M	6.0M	0.06M
Installed super computer in GFLOPS	17.6	12.2	0.
GFLOPS/Student	8.8	2.0	0.8*

\* If a 50 MFLOPS machine had been installed in Israel

As can be seen, if a 50 MFLOPS machine costing about M\$4 is installed as a national facility, Israel would be equipped approximately as 1/10 of Japan or 1/2 of the USA. Alternatively, if two or three universities buy smaller machines for the total sum of \$5-6M the situation remains the same. The retardation in supercomputer availability in Israel is very obvious.

Of course, the table provides statistics only and does not compare quality of research, innovation, creativity or new breakthroughs in science. Notwithstanding, it shows a well

defined trend of the most technologically advanced countries.

Yet another important aspect is the need to attract high quality young scientists to the academic community. As the discipline of super computing is rapidly becoming a major, fast developing one, the availability of a super computer is bound to have an effect on the attractiveness of a university to acquire new recruits to the faculty. The significance of this point should not be underestimated.

Of all sections of the Israeli research and development community the academic sector is the most aware of the need, and probably best positioned to alter its priorities so as to accommodate the financial requirements for super computers. Yet, it is unlikely that this sector will be able to afford the most powerful machines now available, and will, thus, probably have to live with second line machines.

#### 4.1.2. Needs of Government Institutions

Government institutions can be divided into two groups: The first involved manufacturing companies owned by the government, usually in the high-tech area or government users of sophisticated technology. The second group are a number of government run research and development institutions, usually involved in applied research. These institutions need super computers to survive. If they do not utilize the immense power of super computers they will not be able to maintain a high standard, and compete with advanced institutions in the international arena. In the case of government institutions the situation is not as grave as it is with academic institutions, as the spread of super computers to research and development and to high tech industry in general is still in the initial stages abroad. Nevertheless, there is very little time to lose now,



and if super computers are not made available to this segment of the research and development community in the very near future, decline and regression are bound to result.

In principle this section of the user community is sufficiently strong to support institutional super computers. Yet it suffers a slow-down in the present time, possibly retarding its advance in the field of super computing. On a national scale it is important to allow this section to utilize super computing as soon as possible. If this becomes necessary, cooperation with academic institutions should be encouraged until government institutions can support their super computing needs without such cooperation. In fact, this sector may be expected to install (when the need arises) very large super computers.

A major problem for the above section two groups is the security (both commercial and defense) imposed on it. This makes cooperation in the use of super computers between this sector and the academic one very difficult. Indeed, even cooperation within this sector is not simple. Long communication lines, simultaneous use of the machine by alien users or even their access to the machine impose great security risks. This demand appears to be orthogonal to the natural requirements of the academic environment for easy and free access to large numbers of relatively unmonitored users.

#### 4.1.3. Needs of the Private Sector

The high-tech section of the Israeli industry is a potential user of super computing in the same way as the high tech government owned industry. Competitiveness in this sector is bound to compel it to use super computers. Yet this sector is often smaller, driven by short-term economical considerations and is less sophisticated than government institutions. Thus, it may be expected that this sector will be the last to install super computers, and will have to satisfy its needs by renting time in the academical or government institutions.

#### 4.2 Classification of Potential Super Computer Users

It is quite clear that many users are not and will not be interested in using super computers. This may happen because their jobs are too small, or that they require special features not available on the super computer. Still there remain many potential users for super computers in Israeli institutions.

The first group to be considered is that of current users. There are several pioneers in the use of super computers in every university in Israel. These researchers consider the use of super computers essential to their work and go through great difficulties in order to use super computers abroad. Often, this is done via colleagues by international collaboration since in most cases Israelis cannot get financial support for super computing either locally or from foreign foundations. Such work is performed via the international communication network. The machines vary but are mostly CRAYs.

To this group we should add those who work on super computers during visits to cooperating laboratories abroad or by close cooperation with a foreign group who perform the actual runs.

Incidentally, this has a negative impact on the form of export of ideas. An Israeli scientist with an original idea contacts a colleague abroad who can supply the computer means. They carry out the research and publication jointly. Unfortunately, this is one way in which Israeli ideas are channeled abroad freely thus reducing the benefits to the Israeli community.

A second group consists of researchers who are not involved in super computer work at the present time, but may benefit from such possibilities once they become available locally. These researchers are the majority of potential users. They do not apply super computers to their current work due to the present difficulties obtaining access to the machines and the complications in their use from remote stations.

Last but not least are teachers. Teaching in this area has hardly begun due to a lack of possibilities. Scientific or engineering education cannot be practiced through the use of a remote machine unless it is extremely fast, reliable and convenient communication to this machine is made available. The price of current fast communication links prohibits such possibilities. It is not anticipated that basic programming courses will use the super computer, at least not in the near future. However, teaching in advanced subjects must be related to super computing and the creation of new syllabi reflecting the influence of super computing in various fields is very desirable.

In closing this section we wish to comment on the use of super computers at sites where a super computer will not be installed in the near future. Users at such sites are bound to be placed at a disadvantage simply because they will be far from the machines. Still we expect some researchers at distant sites to find their way to super computers at other Israeli

locations, and they may then constitute a part of the total load. To some extent the demands of this group will depend on the central policy adopted by the Planning and Grants Committee and its subcommittee on computing.

#### 4.3. Estimate of the Demand for Super Computers in Israeli Institutions

It is very difficult to estimate the demand for future super computer time. Yet, by discussing the demands of the three user groups listed in the previous section, some estimates may be obtained. The results are very approximate due to objective difficulties estimating future usage patterns associated with new technology. Still, they can serve, with caution and reservations, as guide lines.

We could not arrive at a definite assessment of the number of computer hours used in super computers outside of Israel, by Israeli researchers. However, we believe that the number is significant. We know personally some such users. For instance, at the Technion we have at least five such users. As much of the work is done using various modes of cooperation we do not know how many super computer CPU hours these five researchers use between them. However, an estimate of about 250 equivalent IBM 3081D appears sufficiently conservative to be regarded as lower bound. To get an estimate of the whole Israeli academic community we assume that the demand in three of the other six Israeli universities is similar. Thus, the lower bound for the entire state is about 1000 equivalent IBM 3081D CPU hours. The actual number is probably bigger, and may reach 2000-4000.

To get an estimate of the potential usage by researchers who are not currently using super computers we have carried out a survey of the structure of the present needs for computer time in

all university computing centers in Israel. Table 3 presents the information collected.

Table 3

Total CPU Hours per Year in University Computing Centers in Israel

Job	Very heavy	Heavy	Medium	Small
Time per job	> 1 hour	30-60 min	10-30 min	<10 min
Year	%   hours	%   hours	%   hours	%   hours
1987	34   10550	30   9308	12   3723	24   7447
1986	28   9701	31   10740	17   5889	24   8313

The information is based on data collected from the university computing centers. All times are in CPU hours per year. The CPU hours are normalized to the performance of the IBM 3081D.

About 30% of the CPU time at all computing centers (except for Haifa University and Ben Gurion University) is used by heavy users. Many of these users must be having difficulties with their work, and can not run all the cases which they want. We also believe that most of this work could be run more economically on super computers. Thus, this group which consumes about 10000 CPU hours per year, is a natural candidate to migrate to super computers. The next group, consisting of users with large jobs, probably manages better than the first group on the present mainframes. Still, some members of this group, with 30% of the total CPU time and 10000 CPU hours per year, may benefit from migration to a super computer as

well. The other user groups of 10-30 min per job and less are not likely to contain many candidates for migration to super computers. In particular we estimate that about 20-30% of the CPU used is for I/O intensive work like word processing which can not be helped at all if carried out on a super computer. Indeed, it is far less effective on a big machine than on a small micro computer, which is much more efficient for such applications and offers a much wider range of options than big machines.

The apparent conclusion we draw from these numbers is that about 15000-20000 CPU hours (normalized to IBM3081D) in all universities could benefit by migration to a super computer. However, we do realize that not all jobs can utilize the benefits of super computing. Therefore we offer a conservative estimate: Let us assume that about 50% of the present load (about 10000 CPU hours) will migrate from mainframes to super computers.

Another important aspect of this migration is the number of users who will use the super computer. In order to examine this question we took the situation at the Technion as an example. At the Technion there are today about one hundred researchers who are heavy CPU users. The members of this group use between them about half of the CPU power of the two Technion mainframe IBM 3081D computers. If this number represents the demand at other universities too we may expect a community of about 300-500 users of super computers in Israeli universities. This is a large number which justifies major investments and some cooperation on a national scale.

Finally, we have to consider teaching requirements. If we assume that about 500 students will be involved in courses with super computing content in the whole country, and that each

student will utilize only 10 CPU hours (IBM 3081D equivalent) per term, this adds up to about a 5000 equivalent CPU hours per year.

The three kinds of super computer users described above may thus be expected to utilize between them about 17000 equivalent CPU hours per year. Assuming that the super computer will initially run about 5 times faster than the IBM 3081D this means that the currently existing load may need about 3400 CPU hours per year. On top of it, room must be left for some expansion and growth. Allowing for a 50% growth per year this means 5000 hours after the first year, 7500 hours after the second year and 11000 hours after the third year. Assuming 500 working hours per month (excluding maintenance and system time) we must conclude that there is definite room for at least two super computers in Israel. However, the migration from existing mainframes to the super computers is going to be a relatively slow business, in particular in the first year. Therefore the installation of super computers may be gradual. Meanwhile universities who will have a super computer will be able to provide CPU hours to users from other universities. The actual rate at which the machines will be blocked depends on the number and the size of super computers to be installed and on the quality of the communication network between the super computers and their users.

A rule of thumb for purchasing a new computer is that the present load will constitute about 40% of the power of the new machine. The expected immediate demand is perhaps a bit lower, but bearing in mind the fact that super computing is a new discipline, and the expectations from universities to pioneer new technologies whenever necessary, it appears that the purchase of super computers is already desired by Israeli universities.

On the basis of these figures it seems that super computers may answer some very real needs of the community as it is now. However, this does not take into account the new problem areas which are to be developed in Israel once super computers are locally available. Typically no serious three dimensional modeling or analysis of any scientific or technological issue can be performed in Israel at present, and the availability of super computers will undoubtedly open such problems and give a new dimension to Israeli research and development work.



## 5. THE POSSIBLE SUPER COMPUTERS IN ISRAEL

In this section we shall provide some information on vector computers within the Israeli context. The section is rather limited in scope, since it does not deal with array processors and parallel machines. The information given should be regarded as an introduction to the subject, reflecting the situation at the beginning of 1988.

The machines considered are low end machines for CRAY, ETA and the IBM3090, and high end machines for CONVEX.

### 5.1 Type of Machine

When we claim that a super computer must be brought to the Israeli academia we mean predominantly that academia here must be exposed and have access to the new revolution in the logic of posing and solving problems. By working and experimenting with the new disciplines, progress will be made.

In light of the previous sections it appears that at the moment it would be premature to bring a super computer to Israel in the form of a big multipurpose parallel super computer. That does not mean that small parallel computers should not be introduced in certain departments so that the Israeli researcher can start learning and do research in this developing field. Indeed the academic community should be encouraged to look into parallelism and its various aspects - hardware, software and algorithms, in view of future prospects. Moreover, the major drive should still be towards vector or vector/parallel machines.

Both CRAY and ETA already offer vector/parallel machines with a small number of CPU's where each CPU is a vector machine.

Convex and IBM offer machines with slower vector CPUs but with a higher degree of parallelism. Our examination of several computer installations shows that frequently they do not allow a single user to employ more than one CPU. The reason is the lack of software and poor efficiency. In this respect the system managers of the machines are inclined to value more the total throughput of the system rather than the execution speed. This problem will be resolved when efficient automatic parallelizing compilers are more readily available.

In view of the present state of art it may be argued that a vector machine with a single CPU is still a better choice for a super computer. After a certain period another CPU should be introduced. This will not only double the throughput, but will also allow certain users to experiment with parallelism.

## 5.2. Some Notes on Vendor Selection

Good maintenance of super computers is imperative to provide efficient service in order that users will avail themselves of the machine. It is therefore necessary that the manufacturer of the super computers supply maintenance service in Israel and not overseas. The following manufacturers of vector computers provide maintenance service in Israel: IBM, CDC and Convex. CRAY research is in the process of establishing a service center in Israel. (It is worth while mentioning here that the two CRAY machines installed in Arab countries are maintained by the CRAY London office). Several examples of machines which we have reviewed are given in Table 4 below.

Table 4

	MFLOPS	Memory	CPUs	Clock	Comments
	A 1a			rate	
	Dongarra	Mwords		ns	
CRAY X-MP14se	31	4	1	8.5	
CRAY X-MP1x		8	1	9.5	(used)
ETA 10P1	27	12	1	24	
ETA 10P2		20	2	24	
CONVEX C-210	10	256	1	40	
CONVEX C-220		256	2	40	
CONVEX C-240		256	4	40	
IBM 3090/180 VF		13	8	1	17.2

Prices are not quoted here because all manufacturers offer special deals to universities, and their list prices can serve only as very rough estimates. The machines listed above are these presented to the authors at the time of interviewing. More machines and vendors may be available to the reader. Thus the list is only an illustration of the situation at a given point of time, and does not necessarily represent all possible solutions to the problem. For instance, ALLIANT entered the Israeli market recently but so far we have no information about the services they provide to users, nor their level.

### 5.3. Additional Considerations

The properties of the different machines within the price range of M\$1-3 are quite varied, and require some attention. A CRAY machine within this price range is a mature machine, one which is very popular around the world; it is likely to be a used one of a discontinued line. The other machines of similar computer power are the ETA 10P and the Convex C-220 or C-240. These are new and untried machines. The Convex C-240 is the top of its line and upgrading of such a system depends on future development of new models. The ETA machine is the bottom of its line and offers field upgrading up to the ETA 10Q, which is 27 times as fast, and sells for about \$20M. The 10P is an air cooled machine and field upgrading to big machines (models E or G) requires cryogenic cooling, needing expensive additions. Moreover, cryogenic cooling is a very new and challenging technology which has not yet been fully proven in the field.

The availability of used CRAY machines is a consequence of its popularity, and the large number of CRAY machines sold in the past. Therefore there is plenty of experience with these machines, and abundant software is available for them. The Convex C-2xx series is a direct continuation of the Convex C1 which has an extensive library of software.

The final possibility is the IBM 3090 with vector facility. This machine is really a very fast scalar machine with a vector CPU as an optional add-on. Consequently, the architecture of the machine is not that of other super computers. The basic strength of this machine is in its scalar processor, large memory, parallel architecture and software, the wide availability of software, and the frictionless migration from existing VM machines. One of the largest of its kind (3090-600E) is installed at Cornell University. According to

inquiry at Cornell the machine provided a smooth transition from an IBM mainframe they had operated on campus before, and allowed the utilization of the well established familiar software. The same argument holds also for the Ecole Polytechnique at Montreal.

The operating systems are important as well. Cray offers its traditional stable and proven COS O/S, and its new UNICOS O/S (UNIX). ETA offers the new (and not yet stable) EOS O/S, and expects to be able to deliver a UNIX O/S in 1989. CONVEX offers a mature UNIX O/S. IBM offers its mature VM and MVS systems and declared its new AIX (UNIX) O/S, expected to be available for the 3090 early next year.

Finally, another point should be raised here: An interesting way to reduce the financial burden of the academic system for the installation of super computers is by cooperative programs with the manufacturers. Such programs may include not only subjects related to hardware and general purpose software (like operating systems and general utilities) but also to application packages and user programs in the scientific and technological disciplines typical to super computing. Long-term cooperation is to be preferred to an ad-hoc single project cooperation, and the effort to establish such joint projects should be made by both the universities and the manufacturers. However, it should be noted that such connections can be made between a single university and a single company and are unrealistic on a national scale.

## 6. COMMUNICATIONS AND NATIONAL CENTERS

### 6.1. The Relation between Communication and Super Computing

#### 6.1.1. Interactive Work

Most jobs run on the super computer in batch mode. However, quite frequently a job must be run interactively. For example, debugging a very complex problem in which strange phenomena appear in large quantities of data. Another example is a problem in which the user may want to change a parameter during the calculation. Thus, the communication lines must be able to provide and support interactive work.

#### 6.1.2. Interactive Graphics

Super computers are used to solve large problems. The amount of data processed by super computers is enormous, in particular when three dimensional problems are considered. Therefore graphic pre- and post- processing is often required. Consequently, the normal mode of operation is to use the super computers mainly for number crunching while graphic pre and post processing are done mostly on work stations. The stations, which are much more suitable for graphic and interactive processing can perform some of the calculations required for the graphic processing, but number crunching for the graphic processing is often delegated to the super computer as well, and the two tasks on the two machines communicate with one another. The proportion of graphic processing delegated to the super computer depends on the relative strength of the two machines.

As an illustration let us consider a simple problem in three dimensional fluid dynamics. Such a problem may require

about 1M mesh points and about 3M words storage for the dependent variables. These are about 200 Mbit. For post-processing it is required to transfer all these data to the work station. A simple calculation shows that communication lines band width should be measured in MBits/s rather than KBits/s. Thus T1 communication lines are the minimum requirements for communication between remote sites and the central facility, and ETHERNET lines are the minimal requirement for local communications.

### 6.1.3 Super Computer Network

When a national center with communication lines to remote sites is considered, a dedicated "STAR" communication system is the preferred configuration. It should be noted that most networks serve sites which are equal in importance, and use therefore a "NET" configuration which is less attractive for communication to a central super computer. Moreover, if the communication is performed on a general purpose "NET" system the amount of data from the super computer is bound to slow down the transfer rate of general data to unacceptable levels. The practical significance of the above discussion is that the price of a fast "NET" communication system must be added to the price of a super computer. The communication is cheaper when the remote users are not far (say within a large campus), but it may be quite high for geographically distant sites.

The same problem is encountered elsewhere: The NSF in the USE has established not long ago several national centers of super computers with a 1Mbit/s network. Some super computing networking exists also in Europe, in particular in the UK and in Germany. While these networks do serve users who do not have local super computer capability it is quite apparent that the users wish to have local capabilities even if they are connected

to a strong network, and they prefer to use the local facilities beyond their capability before they resort to a remote machine. On top of the problems of transfer of huge files other difficulties may arise as well. For instance the quality of consultation available locally is much better than advice through the communication network, and delays in the lines disturb the continuous work of the users.

A crucial factor is the quality of the line. We have heard from users in the USA that the 56KBits/s line seldom stays on for more than 15 minutes without interruption. If this is the situation in the USA where the telephone service is a good one what will happen in Israel?

## 6.2. The Economics of Local and National Super Computing

We have seen that the choice between a national center and local ones depends not only on the cost of a central facility versus the cost of a local one, but also on the cost of communication. Before going into a detailed analysis of the two possibilities we have to point out that since the CPU power/dollar invested rises with the cost of the machine, a central facility is more powerful than the sum of local ones purchased for the same total cost. Moreover, a large central facility can, in principle, run bigger problems than the smaller local facilities. On the other hand, the number of users on a local facility is smaller, and therefore their access to the machine is easier.

From the previous discussion it should be realized that the economics of network super computing is not simple. The communication system alone is rather important, but in addition



some other factors should be considered as well:

a) A central system may require a double storage system (for the same files, locally and at the super computer), so as to minimize the amount of traffic on the communication network. Thus for the solution of a national facility the total storage will be larger than the total storage required if only local super computers are installed. It should be noted that the price of disk storage constitutes a large part of the total price.

b) A central site often requires more workers for its regular operation.

c) The communication system for a central site connected to remote users is more expensive than the system required for the same number of peer sites with local capabilities.

d) Even when universities are connected to a central super computer they need some local vector capabilities for teaching and training as well as local test processing and preparation of work. The computing power of the local capability is inversely proportional to the band width of the communication network.

Most of these factors are difficult to quantify. Fortunately it so turns out that communication, which may be the major item for a small system, is also the easiest to estimate. It is therefore reasonable to examine the ratio between the annual cost of a central facility to that of communication. In principle the annual cost should include depreciation as well as running expenses. However, in view of the relative magnitude of these quantities it may be sufficient to consider the depreciation only. Obviously such a measure gives only a rough

guide, and should not be used for limiting cases. Still, we believe that the following model may serve as a guideline for most cases:

Consider a super computer with a life span of  $Y$  years which serves  $N$  sites, and costs  $P$  versus  $N$  smaller machines. Let us assume that the price of the super computers rises as the square root of their strength. We also assume that the purchase of local super computers for the sites can be justified only if the total installed capacity is twice as big as that of the central site. Under these assumptions the break even point between the two options is given by

$$P + Y * C = N * P * (2/N)**0.5$$

or

$$r = (P/C)_{\text{limit}} = Y / [ (2 * N)**0.5 - 1 ]$$

For  $Y = 5$  and  $N = 5$  the result is  $r = 2.3$

For  $Y = 5$  and  $N = 4$  the result is  $r = 2.7$

For  $Y = 5$  and  $N = 3$  the result is  $r = 3.4$

For  $Y = 5$  and  $N = 2$  the result is  $r = 5.0$

The meaning of the above example is that for a system with 5 remote sites the total purchase price must be at least 2.3 times more than the annual communication bill to make the system attractive. For a 3 site system the corresponding ratio is 3.4. The results for a number of cases are shown in the table below.

TABLE 5

Break even point for central versus distributed super computing

P            Price of central machine  
C            Annual communication price  
r            P/C  
N            Number of users of distributed system  
Y            Number of years of utilization

$$r = Y / [ (2*N)**0.5 - 1 ]$$

Y	2	3	4	5	6
N					
1	4.83	7.24	9.66	12.07	14.49
2	2	3	4	5	6
3	1.38	2.07	2.76	3.45	4.14
4	1.09	1.64	2.19	2.73	3.28
5	0.92	1.39	1.85	2.31	2.77
6	0.81	1.22	1.62	2.03	2.43
7	0.73	1.09	1.46	1.82	2.19

However, in view of the crudeness of the model it would be wiser to state the following rule of thumb:

When

$$r > 2 * Y / [ (2 * N)**0.5 - 1 ]$$

a central super computer is preferable.

When

$$r < 0.5 * Y / [ (2 * N)**0.5 - 1 ]$$

a distributed system is preferable.

In the above example of a 3 site system and for an annual communication cost of M\$2 -- a central super computer is attractive if it costs more than M\$6.8. If it costs less than M\$1.7 the distributed system is more attractive. In intermediate cases it would be better to perform a more accurate economical study, taking into account all other expenses and factors.

### 6.3. Price of Communication in Israel

Long distance computer communication lines are usually dedicated lines hired from public communication vendors. The prices of communications to remote sites is determined by these vendors. In Israel the vendor is Bezek. Current Bezek prices are given in appendix B together with estimates of quantities and prices for a national super computer communication network. We have included two options: A national super computer center in Haifa, and one in Tel Aviv.

The annual price of the cheaper option of a center in Tel Aviv is about \$2M. This price may be justified (perhaps) if the central site is large, with a \$20M machine, but not with a \$5M (or less) machine. Thus, a T1 national communication system appears too expensive to be justified at the present time. A significant saving may be obtained by degrading the band width of the line to 64 KBAUD, at least temporarily. In this case the relevant annual costs are \$240,000/year for a center in Tel Aviv and \$480,000/year for a center in Haifa . In other words about a factor of 10 less. Could the national network operate on such lines? In the USA, Germany and Britain it was found that

1MB lines are a better choice because if slower lines are used the network becomes overloaded and the response time becomes painfully long. In such circumstances the use of the national center is seriously disturbed and users resort to local facilities albeit the higher cost to the system.

Another possibility is to seek cheaper rates from Bezek. Obviously this problem can not be resolved inside the academic community. Yet the option ought to be considered carefully. However, comparison with prices abroad does not suggest price reduction of the order required. While Bezek is not cheap, it is still within the price range of some other countries. West Germany is a typical country with high cost communications (the German example is also similar to ours because the German universities are financed directly by the state administrations). Still, the high cost of communication have contributed to the fact that the state computers are not sufficiently in use and the universities operate their local machines. In the US we know about a local DOD institution which has two separate centers and used to operate two 1MB lines between them. In view of the high monthly rent they canceled one line and installed many local machines in one of the two centers.

#### 6.4. Can Israel have a National Super Computer Center?

Israel's size, as well as its small number of universities appear to suggest that the best way to satisfy the super computer needs of the country is by forming a large national super computer center. Such a solution may be able to attract more support, and thus allow the installation of a machine larger than any single university could afford, and may also form a nucleus for a center of excellence in super computing.

The major factors working against such a solution are of two kinds: Firstly, such a solution is viable only if the price of communications does not exceed a certain portion of the price of the computer. Thus a relatively small machine may be used for a central site only if a cheap communication method is selected. However, a cheap communication network limits the possibilities of research on the machine to problems requiring small volume of data transfer only. As much of the work done on super computers is the solution of large three dimensional systems, the graphic post-processing of the results is an essential part of the work. However, graphic post processing is of necessity a task requiring the fast transfer of large amounts of data. Thus, a cheap and slow communication system will eliminate a large and important class of problem areas from the machine.

The financial benefits should not be ignored too: If a national center is established the rate of investment for the super computers and the network must be higher to ensure that all participants have access to the machine within a reasonably short time. On the other hand, if every campus decides on its own way to super computing the timing will be more flexible. Not all universities will decide to install a super computer at the same time (in particular those which have not expressed so far a strong interest in it (Haifa, Ben Gurion and Bar Ilan Universities). In this case, the total national expenditure will be spread over a longer period of time. Moreover, if the Plannig and Grants Committee decides to establish a national center it will have to provide most of the financial support for it. On the other hand, individual universities may find contributors for the super computers among their "Friends of the Universities" societies.

Another problem area is the incorporation of super computers in the educational process and that of local prestige. While super computers are by no means a tool for large classes and preliminary courses, the educational possibilities and demands imposed by them should not be overlooked. The development of educational programs for graduate and senior under-graduate students should form an important part of the implementation of super computing in Israel (together with research and development). However, teaching is much more difficult to perform when the machine is not on campus.

The problem of university prestige cannot be overlooked either. Having a super computer has a distinct influence on this factor, and it is not likely that campuses with a strong interest in science and technology can afford to stay without some local super computing capabilities for a long time. Thus a national super computer center will undoubtedly be accompanied by the mushrooming of smaller local centers whose existence it was supposed to make unnecessary. This trend is possible nowadays, when the price of small super computers is not higher than that of large mainframes.

Last but not least is the problem of advisory services. If super computers are to be widely used a system of consultation and instruction must be provided. From the observations of the structure of super computer centers we infer that a staff of about 10-20 people are needed to supply the demand (the Technion has for example a staff of 12 programmers dedicated to consultations on its machines and the applications run on these machines). If the national center is added to an existing university computing center it will not be able to provide the support and it is bound to fail on grounds of poor service. Thus, the budget for the payroll must be added to the cost of a national center. On the other hand, if the super computer

is to service one campus only the demand for consultation will be smaller and the ability of the institution to dedicate the required personnel either by new staff or from its present computing center staff is significantly larger.

The points mentioned above stressing the importance of the installation of super computers in Israel carry to the local university level as well. This is an important question which should be considered in the light of other changes occurring in the university computing environment. Practically all universities are running today some mainframes and super mini computers. These are used for various central tasks, but their role is declining as powerful work stations become more readily available. With the prices of small super computers decreasing as they do now, the following question arises: To what extent should universities consider small super computers as a part of their future central campus system?

In view of what has been said above it becomes clear that different institutions will strive to possess super computers. If cooperation is to exist all universities must have some vector computing capabilities so that they can contribute in some way to the development of the new discipline.

#### 6.5. A Little more on the Financial Side

Let us consider the case of four universities pooling all their resources to buy a super computer jointly. The total amount of money they can raise under the most favorable conditions is somehow related to the total cost of the equipment installed presently in a typical university campus. We estimate that the total value of the equipment is of the order of M\$3-4. On this basis let us assume that each university will be able and willing to allocate the sum of M\$1



for this purpose. Thus they can raise together M\$4. On the basis of our previous assumptions M\$8 can buy a 50 MFLOPS machine. The maintenance cost will be 10% a year or \$800,000/year. On the other hand, if they decide to buy four \$1M machines they will get about 40 MFLOPS with about the same maintenance cost. The price of communications for the central site is estimated at \$2.1M/year. This calculation does not take into account the fact that the central facility must have storage of its own on top of the disc storage area in each university. Clearly a national center is more expensive than the sum of all super computers considered.

Typically, US installations are much larger. Let us look again to the situation in the USA. The installation at the University of Illinois costs about \$75M and the operational budget for 5 years is \$75M (see the University of Illinois at Urbana Champaign new bureau Feb 25, 1985). The John von Neuman Center for Scientific Computing (JVNC) at Princeton will cost \$123M over five years. The NSF contributed \$200M over 5 years for the establishment of 4 university centers of super computers and the centers are expected to raise the rest of the money from donors and industry. Indeed IBM donated to Cornell its 3090-600E machine (as well as the previous machine) and CDC donated to the JVNC its CYBER 205. Under such conditions the relative cost of the communication is significantly smaller.

The total number of MFLOPS is still larger in the central facility and in principle the jobs run on the central facility may be heavier than those run on any of the smaller distributed super computers. However, this statement must be viewed with caution. The longest job on a typical mainframe is about 3 CPU hours. When problems demand longer CPU time, the job is stopped, the data dumped on a disc and later continued from the point where it stopped. This is not convenient but

many users find it necessary. The same can be done (and eventually will be done) on the smaller super computers. It is only for the very big jobs the same technique will be used on the larger machine, as it may not run on the smaller machines.

The prices mentioned above assume that the deal with the company is only for the super computers. However, some of the vendors of super computers have other equipment on various campuses and the deal may include other pieces of equipment. The simple economical calculation becomes meaningless in this case. Thus, it should be borne in mind that the accuracy of the above numbers does not take into account additional deals.

#### 6.6. The Case for a Distributed Israeli Super Computer System

From all the above it becomes clear that the central solution is not a viable one in the Israeli context. If the country could afford to pay about M\$30 a central solution might have been a feasible one. Such a solution would have included a big national super computer staffed with about 30 system, advisory and operating personnel, a fast communication system, and small vector machines distributed in all campuses. However, with the current financial restrictions imposed on the academic system the total amount of money available appears to be a smaller order of magnitude, and therefore the central solution cannot be recommended now.

## 7. USER ENVIRONMENT

### 7.1. Connectivity

One of the most important factors affecting user migration to a new machine is the ease of connection. Most scientific users are interested mainly in their scientific work. They may be proficient in numerical methods and programming languages as well. However, they only seldom take the effort needed to master operating systems and communication software. Consequently, they are very reluctant and slow to accept changes in these elements. Unfortunately, system and communication software have the tendency to change often, and in particular the introduction of a new machine requires retraining of the entire user community. This is not a simple problem, and it has been proven once and again in the past that user migration to new machines is inversely proportional to the effort involved in mastering the new system software.

One of the possible answers to this trend is to allow the users to use the new machine from their previous environment. This may be achieved either by adding user friendly communication software to the old machine or by creating an artificial environment (shell) on the new machine which attempts to look as much as possible like the old system. Once such software is available user migration becomes a much simpler task, and the users are much more agreeable to migrate to the new machine. Ideally, one would like to have all machines use the same operating system and communication software, a move which may eventually materialise with UNIX and TCP/IP. However, at the present time a very big portion of the existing machines use specific vendor software, which is often much more user friendly and often better geared to the needs of the users.

It is therefore important that a new super computer be equipped with user friendly and easy to use communication software which allows the operation of the machine from other environments like IBM MVS or VM, DEC VMS, or CDC NOS and NOS/VE. The ideal communication software allows the user to work from his familiar environment without any new alien commands. A very good example of such software is the Macintosh desk top communication to NOS/VE machines. CRAY supplies the STATION software for this purpose. It is advantageous to acquire such software together with the super computers, and preferably from the same vendor.

We have to stress this point. The investment in a super computing system is immense. It includes not only the price of the machine and its maintenance, but also a lot of work. Much of the work is performed by system programmers employed by the computer centers. However, the lion's share of this work is performed by users on the migration path. Helping the users saves them a lot of time and money, and allows a much more speedy utilisation of the machine. Thus, the investment in the machine is utilized and justified much faster. Moreover, the introduction of appropriate communication software allows faster and unimpeded progress in the field of super computing.

## 7.2. Operating System and Utilities

In the past each computer manufacturer used to supply its own operating system (or systems). This trend was somewhat changed when manufacturers of micro computers chose almost unanimously the DOS operating system and manufacturers of work stations opted for the UNIX system. While DOS is still a single user small machine system, UNIX has developed into a very popular and widely spread system, and its application to

super minis, mainframes and super computers is advancing very quickly. The advent of a single operating system for practically all machines has many important implications. One of the most important factors is the ease in software portability from one machine to the other. Today CRAY and CONVEX are using the UNIX operating system, while ETA and IBM are developing such a system. Even if the UNIX operating system is not available, a UNIX shell may supply most users with all the UNIX commands and environment necessary for daily use of the machine. It is therefore important to ensure that whichever machine is used it will have the UNIX operating system, or at least a UNIX shell. Let us stress, it is not that the machines cannot be used without UNIX, but the existence of UNIX on workstations and super computers should facilitate the exchange of machines and portability in general.

Other important standard features include compilers, analyzers and debuggers. Debuggers are important in the development of large computer programs. At the first stage of the super computer on campus when programs are to be converted, debuggers are essential. In particular it is important that users should be able to understand the messages which very frequently are cryptic, ambiguous or even incorrect. Sometimes, and in particular with jobs for a super computer, the program is so complicated that only interactive work can help. Thus, an interactive debugger is required.

Three types of programs essential for super computing are: vectorizers, parallelizers and analyzers. Vectorizers and parallelizers usually constitute a part of the compilers. For the time being automatic parallelizers exist only for the CONVEX and the ALLIANT and experience with them is still limited. Manual parallelizers, as well as automatic and manual vectorizers are available for all the machines considered in

this report. Analyzers are used to identify parts of the program which can benefit from vectorization, and are extremely important in the development of new programs, or the migration to super computers.

To sum up: the preferred operating system for super computers is UNIX. The relative quality of vectorizers and parallelizers can be established only by benchmarking. Analyzers are very important to efficient vectorization.

### 7.3. Consulting and Instruction

Vector programming is quite different from scalar programming. The changes relate not only to programming techniques but also to the algorithms used. Users unfamiliar with vector programming will tend to use scalar programs, and lose the basic advantage of a vector machine. Thus, the large investment in hardware is not utilized unless some modest amount of money is invested in advisory and consultation services, including class instruction. The need for such services may be demonstrated in the following figures:

The ratio in running time between scalar programs and vector programs is about 3-5. The improvement of an automatically vectorized code by manual vectorization is often by a factor of 3-4. Poor programming may mean that automatic vectorization cannot succeed at all leading the user to apply the wrong algorithm for the machine.

The road to good instruction and advice goes through services of expert advisors to individual users, ad hoc courses on programming tools, and incorporation of the subject in classroom teaching of the usual courses on programming and numerical analysis. To achieve this it is important to have

significant research programs in related topics. It is also necessary that the advisers have continuous experience with vector (and parallel) programming, through active participation in ongoing projects.

#### 7.4. User Migration

One of the most important issues in the introduction of super computers in Israel is that of user migration from mainframe computers to super computers. The two barriers a user should overcome in the migration process are: 1) the effort needed to master a new and unknown operating system which may be unfriendly, and 2) the need to know how to write vectorizable programs. If the system and utilities are difficult and unfriendly many users will shy away from the super computer; if the user does not manage to vectorize his program efficiently his motivation to migrate will diminish.

The first problem is tackled in many cases by using a well-known and sufficiently friendly operating system. The one preferred by most manufacturers is UNIX. Its popularity and availability as well as its relative friendliness make it almost the only possible choice. Indeed CRAY and CONVEX have already moved to UNIX and ETA, and IBM (3090) announced their intention to use it.

As for the effort required for vectorization we should recall what was found in the Technion benchmark. The benchmark included 9 programs involving heavy Technion users. All programs were found to run very efficiently after being compiled by an optimizing compiler. The programs were given to the different companies for benchmarking. It took on the average 25 man-days of a very experienced system programmer to convert the codes and reach a high percentage of vectorization. Note that the

programmers had absolutely no knowledge of the scientific content of the programs and they did not alter the algorithms. Altogether there were about 17,000 lines of code. The level of improvement varied between companies and programs. On the average an improvement of 5-10 times that of the "as is" code was found in many of the codes, in particular those which did not vectorize very well in the "as is" condition.

To reach a similar improvement by the users themselves two conditions must be met: Firstly a course on vectorization is required. Such a course may take about one week. Secondly good and abundant support and advice is required. If these two conditions are met it may be expected that most faculty members should be able to modify their codes so that they vectorize reasonably well within a period of 2-4 weeks, depending on previous exposure to vectorization, type of program and the quality and quantity of advisory services. Obviously this goal can be achieved only with active interaction of the user with the computer and its staff. The next step, of finding new algorithms and implementing them in the code is another story which may take a very long time.



## 8. DISCUSSION AND CONCLUSIONS

### 8.1. The Necessity for Super Computers in Israel

The obvious question to ask is: Why establish a super computer center in Israel if the researchers can in principle use such facilities abroad in spite of the difficulties? Indeed the use of super computers in Europe and North America is a possibility which can not be ignored when considering the needs of scientists in Israel. Such a possibility relieves the Israeli system from the financial burden and finds the solution to the few scientists who are "obsessed" with super computers without depleting the heavily loaded Planning and Grants Committee resources which appears to be permanently short in face of the needs of Israeli universities.

The first factor to be considered here is that of the communication system. Communication with North America, or even with Europe is even more expensive than local communication, and if such a solution is preferred no Israeli university would have the privilege of being sufficiently close to the machine. Thus, the present situation in which most scientists simply decline projects in which super computing is important, will remain, with all the severe consequences of neglecting the new discipline, and not having the research and development community avail themselves of the benefits of this state of the art system.

Apart from the problem of the communication lines there is a second factor affecting the decision of installing super computers in Israel. The existence of super computers on campus provides a tremendous push to the development of research in super computer applications and the relevant disciplines. Indeed, local facilities would have the chance to become centers

of knowledge and research and hence should be here and not abroad. Although it is impossible to quantify this factor it should not be ignored. For example, the national super computer centers in the USA gave strong support to the formation of centers for theoretical studies which host many researchers in various disciplines. This may be seen in the NSF sites in the USA, e.g., Cornell and Princeton, and in the collaborative effort of NASA-AMES and Stanford.

The third point to consider is the prestige of any university in which a super computer is to be installed. This prestige translates into various practical benefits, like the willingness of potential donors to contribute to the university, or that of agencies to support research, the capability to participate in joint international projects, or the desire of young academics to seek employment in the university. Indeed, if Israeli universities strive for excellence and wish to be listed among leading international institution in engineering and the sciences, they cannot afford to neglect the field of super computing, which has become so crucial in most of these branches.

The fourth point, super computers have already grown out of their infancy, and they are recognized and used by industry. It is therefore essential to include the subject in the research and teaching programs in the universities, so as to ensure that the infrastructure for super computing in industry exists in Israel. This point should not be neglected: The impact of a local center tuned to the particular needs of the Israeli community is incomparable to that of a remote facility.

The final argument is an economical one: The price of medium super computers today is not much different from the price of mainframes a few years ago. Typically the machines in

question cost today between \$2-4M. The Technion's first IBM 3081D bought about three years ago was priced about the same in actual US dollars. Moreover, a very short calculation shows that the price performance of the new machines is better than that of the old mainframes by more than an order of magnitude. Indeed, the computing power of the 3081D can be purchased in a super computer for about K\$20-100.

## 8.2. Cooperation in a Distributed System

This committee was not asked to explore the possible cooperation between the Israeli universities in computer resource purchases etc. Yet, once the question of a national center is raised the problem of cooperation floats. Two extreme opinions are voiced. The first is to minimize cooperation and let the universities compete among themselves. Competition on prestige can become a good leverage to improve conditions for research on the campuses -- with the hope that better conditions will attract the best scientists and result in better research. The other extreme opinion says that the universities are too poor for internal competition and they should join forces so as to exercise greater pressure on the companies to reduce prices, prevent duplicating the effort in software development etc. We do not intend to express our opinion on this matter. We will however mention briefly those topics in which collaboration can be useful to a point.

The universities as a unified body are larger than a single institution and hence are expected to be able to get superior deals from the different companies by buying larger quantities of hardware and software. In the past this was not the case because:

a) The companies did not make concessions give in to the Machba more than it did to the single university. Moreover, the various deals with the single university contained various pieces of equipment which are not always required by another university at the same time or even at all.

b) There is always the question of academic freedom. Should we force the researcher to buy a compromised equipment rather than the particular one which best fits his requirements?

Even if the universities decide to compete among themselves there is plenty of room for cooperation in software. Most university computing centers today are multi-vendor sites and hence there is a significant identity in equipment and overlap in required software. As a result, cooperation in software, operating systems, compilers, editors and some applications are highly desirable. Consistency in software will ease the collaboration between researchers since they will enable researchers from one institute to use the equipment in another institute.

Last but not least, cooperation in courses and instructions are highly desirable.

### 8.3. The Case for some Local Centers

We have discussed above the future of super computing and expressed our prediction that the future belongs to vector parallel machines. Hence it is our recommendation that small parallel machines be installed in the academia and be made available to researchers in computer sciences as well as science and engineering. Typical machines are the ALLIANT and the bigger CONVEXES. They are not very expensive (in terms of the total equipment installed in any single university computing center) and will provide a proper avenue to an early start in a newly emerging technology.

### 8.4 Conclusions

1) Super computers are among the drivers of the forefront of technology. The lack of super computers in Israel impedes the progress of science and technology in the country.

2) Super computers should be placed at the top of the priority list of multi-user equipment required by the academia in Israel.

3) A national center for super computing with T1 communication lines is the best solution if the price and size of the central facility is above a critical minimum. Our estimates based on expected funds allocated to a super computer on one hand and the demand for super computer power on the other, are that this solution is too expensive.

4) The preferred solution at the present demand for super computer power is a small super computer on each campus.

5) The super computer operating system must be UNIX.

6) At present, vector technology (hardware and software) is more developed than the parallel one. Yet, we are convinced that the future belongs to parallel vector machines.

7) In view of the above it is clear that at present vector machines should be preferred to parallel ones. However, options of parallel processing in general, and in available machines in particular, should not be ignored.

8) Consulting services are extremely important and any new super computer must be supplemented by proper consulting services.

9) Vector technology and parallelism should be included in the curriculum.

## APPENDIX A: Evaluating Computer System Performance

Following are some definitions of terms frequently used (and seldom defined) in the discussions of super computers.

**CLOCK SPEED OR CLOCK FREQUENCY:** this is the number of cycles per second provided to the computer by a timing device such as a crystal. These clock cycles are the basic timing reference signals for the computer operations. Clock speed varies from 5-20 MHz in micro computers to about 50-200 MHz in super computers. The actual performance depends not only on clock speed but also on the architecture and software used. Therefore the clock speed is more a measure of the potential power of a given single CPU than of the performance of a machine.

**COMPUTER ARCHITECTURE:** classical machines have the following traditional elements: central processing unit, memory, discs, tapes, and input/output channels. These elements can be arranged to interact with one another in different ways so as to achieve better performance under various conditions. The arrangement of the elements is loosely called computer architecture.

**MIPS:** this is an acronym for Millions of Instructions per Second and it refers to the basic machine operations such as: move an element from one location to another, store an element, move a register content one step forward etc. While the number of MIPS provides information on how fast the machine carries out simple elementary operations, it does not relate directly to the speed at which numerical calculations are performed.

**MFLOPS:** An acronym for Millions of Floating point Operations Per Second. This is the standard measure for the comparisons of machine performance on scientific jobs. It represents the

number of operations (irrespective of their type addition or divisions etc) required to complete a certain job and divided by the time it takes the machine to complete the job. It is well known that this number is strongly dependent on the machine architecture, on the quality of programming, and on the algorithm used to solve the problem.

**BENCHMARKS:** A comparison between two or more machines based on the actual running of the same program on the machines. Clearly, the benchmark is the summary of all the features of the machine as they come into effect in 'real life'. The major problem is however, that different types of applications will yield different relative results on different machines because of the problem formulation and the machine architecture. In spite of this difficulty benchmarks still are the best tool for comparison of machines. Some standard benchmarks are well known in the international computer community: They are important in relating the results of any given benchmark to those carried out by other institutions.

**LIVERMORE LOOPS:** As Livermore laboratories were among the first sites to operate super computers they developed a kernel of frequently run routines to compare the various machines. The routines are all in FORTRAN.

**LINPACK or DONGARRA:** LINPACK is a mathematical software library out of which a specific problem of matrix inversion is chosen. The standard matrix is 100x100 or 300x300. The requirements are that the code is not modified so as to take advantage of the unique features of every machine in the test. It should be stressed that frequently small manual changes in the program can improve the performance immensely.



WHETSTONE: The WHETSTONE is a synthetic mixture of computer instructions based on statistics of scientific calculations in FORTRAN. The program runs are 'statistical average'.

DHRYSTONE: Similar to Whetstone but in the C language.

It should be stressed that scientific calculations may sometimes be I/O bound and should this be the case, none of the above benchmarks is relevant.

VECTORIZATION AND PARALLELIZATION: Vectorization means that vector properties of the CPU are used to run the loops as vector loops. For instance the scalar product of two vectors in a super computer consists of  $n$  multiplications and a sum of  $n$  elements. In a vector machine all the products are carried out simultaneously (for  $n$  less than the vector size of the machine). In scalar machines the  $n$  products are carried out one after the other. Parallelism means that two independent tasks are calculated at the same time.

## APPENDIX B: Costs of Computer Communication in Israel

The following are the basic costs of communications as provided to us by the Bezek company.

### Fixed costs:

For a local section of up to 3km long: NS8,200;  
for the endpoint equipment: NS16,400 for the line.

For an out of town line above 3km long: NS19,700;  
for the endpoint equipment NS39,400 for the line.

### Rental costs:

NS970 for up to 3km long line and NS1780 for any line longer than 3km.

We consider two examples of national centers, one located in Tel Aviv which is close to the geographical center of the academia and one located at Technion.

### 1. Costs for a National Center in Tel Aviv

The distances between the universities in Israel are:

Tel Aviv - Haifa	90km	1 endpoint
Tel Aviv - Rehovoth	30km	1 endpoint
Tel Aviv - Jerusalem	50km	1 endpoint
Tel Aviv - Beersheva	110km	1 endpoint

Altogether 280km, 5x2 local endpoints of up to 3km, and 4 out of town lines.

The installation cost is NS82,000. The basic topology of the network is a star. The monthly rent is:

10 local lines NS9,700

4 long distance lines (280km) NS271,000

Total cost NS3,750,000/year or about \$2.11M/year.

## 2. Costs for a National Center at the Technion

The distances between the universities are:

Haifa - Tel Aviv	90km	2 endpoints
Haifa - Tel Aviv	110km	1 endpoint
Haifa - Jerusalem	130km	1 endpoint
Haifa - Beersheva	200km	1 endpoint

The basic topology is two lines to Tel Aviv and a smaller star.

Altogether:

5x2 end points,

5x2 local networks of up to 3km long,

5 long distance lines.

The installation cost is NS82,000.

The monthly rent is:

10 local lines: NS9,700

5 long distance lines (620km): NS601,000

Total cost: NS7,330,000/year or about \$4.6M/year.