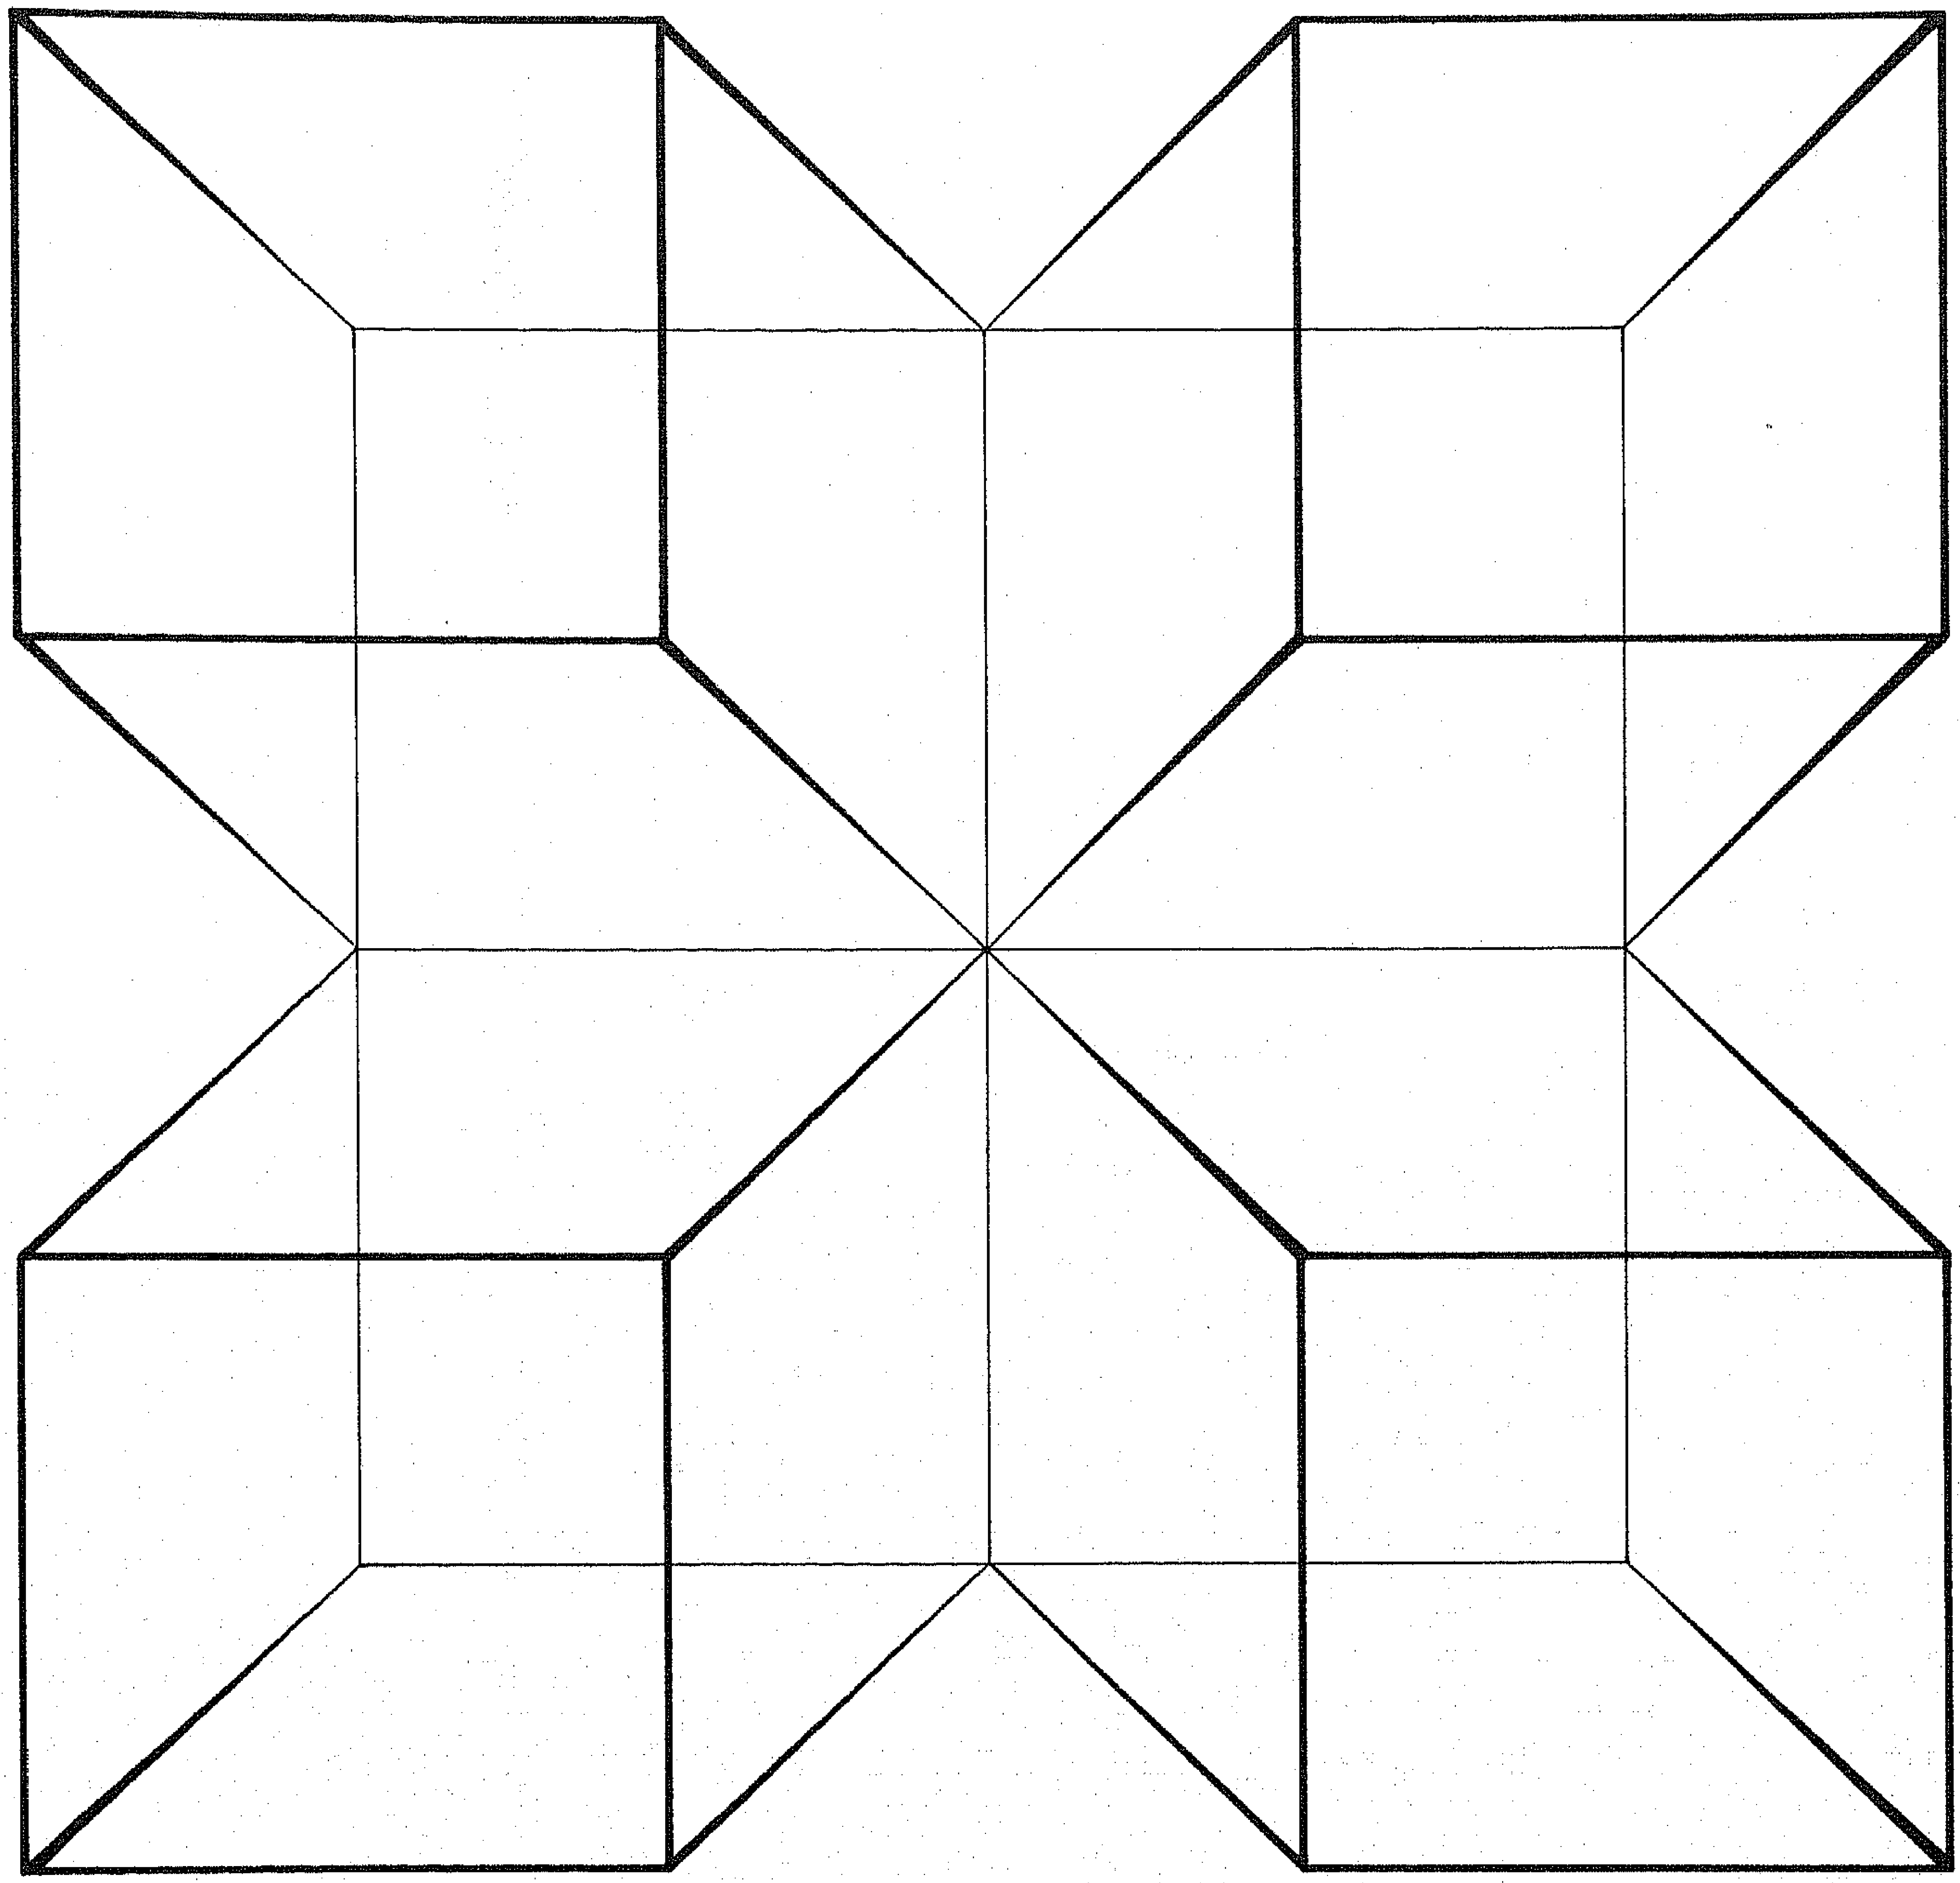

Computing Facilities Policy Document
(Beleidsplan Computervoorzieningen)
1991-1996

CWI: the Center for Mathematics and Computer Science
(het Centrum voor Wiskunde en Informatica)

Computing Facilities Policy Document
(Beleidsplan Computervoorzieningen)

1991-1996

CWI: The Center for Mathematics and Computer Science
(Het Centrum voor Wiskunde en Informatica)
Amsterdam, The Netherlands

June 1990

Table of Contents

1. Introduction 1

2. A Review of CWI's Research Equipment Acquisition Policy 2

3. Present Situation at CWI 3

4. Policy Objectives for Supporting Mathematics and Computer Science 4
4.1. Computer Architectures for Supporting Research 4
4.2. A Six-Year View of CWI's Architectures for Research Support 6

5. Projected Needs, 1991-1996 9
5.1. Basic Assumptions 9
5.2. Research Infrastructure 9
5.2.1. Research Workstations 10
5.2.2. Central Fileservers 11
5.2.3. Network Infrastructure 11
5.2.4. Central Facilities 12
5.2.5. Input/Output Facilities 12
5.2.6. Research/Management Support Facilities 13
5.3. Project-Specific Equipment 14
5.3.1. Multimedia Systems 14
5.3.2. Scientific Visualization 16
5.3.3. Computational Geometry 18
5.3.4. Image Processing 18
5.4. Computer Algebra Nederland 20

6. Budget Projections 21
6.1. Projected Total Needs 21
6.2. Financing Projections, 1991-1996 22
6.3. Yearly Specifications, 1991-1992 23

1. Introduction

This report presents the equipment acquisition and support policy for the Center for Mathematics and Computer Science (CWI: *Centrum voor Wiskunde en Informatica*) for the period 1991-1996. CWI is the research institute of the foundation Mathematical Center (SMC: *Stichting Mathematisch Centrum*), an independent organization funded primarily by the Netherlands organization for the advancement of research (NWO: *Nederlandse Organisatie voor Wetenschappelijk Onderzoek*). CWI has three primary responsibilities within the world of higher education in the Netherlands:

- to perform theoretical and applied research in the various fundamental areas of mathematics and computer science;
- to facilitate the transfer of knowledge between the research community and various educational, governmental, and private-sector institutions; and
- to serve as a national and international center where researchers in computer science and mathematics can meet, exchange ideas, and plan future work.

In order to assist it in carrying out these responsibilities, CWI has developed a computing infrastructure that allows its researchers and guests to perform their research and knowledge-transfer tasks. This infrastructure is kept up-to-date based on an internally-formulated research support policy that attempts to maximize the facilities provided to CWI's researchers while retaining a coherent environment that is adaptable to the changing needs of the research community. The equipment infrastructure is made possible via funding provided primarily by the Dutch ministry of education (O&W: *Ministerie van Onderwijs en Wetenschap*) and NWO.

Given CWI's intended international role within the mathematics and computer science research community, the provision of a first-rate, world-class equipment infrastructure is one of the basic elements of our over-all research policy. During the past two years, we have developed an equipment acquisition strategy that is based on the assumption that computing equipment has become so integral to the research process—for arriving at results, for disseminating results, and for managing the research process—that equipment access should be as common as telephone access. In general, every researcher (as well as every research manager and administrator) should have access to modern facilities that are matched to their needs, as well as access to appropriate training and support facilities to utilize this equipment.

The purpose of this document is to detail our research support policy for a period of six years. The policy that we have formulated is an extension of our existing major effort to modernize the facilities available at the Institute. In order to recover from a serious deficiency, we have had an acquisition-intensive support policy during the past two years that should be continued. At the same time, we realize that research equipment funds are not unlimited and that coordination of acquisition and placement is essential to maintain a cost-effective environment.

In the following five sections, we present our views on how the intensive requirements of our researchers can be realized in terms of projected funding subsidies of our supporting agencies. We begin by reviewing our overall assumptions on the nature of our computing infrastructure. We then describe the current equipment situation at CWI; this description gives an indication of the progress that has been made since our last broad policy statement in 1987. Next, we describe our detailed research equipment needs; this discussion is divided across general institute-wide needs and the needs of four projects that are indicative of equipment intensive research. We conclude this report with a series of budget projections (both on an annual and per-project basis).

2. A Review of CWI's Research Equipment Acquisition Policy

The past decade has seen a significant change in the way that researchers are provided with computing resources. Computing equipment has become cheaper, faster, easier to use, and more reliable. At the same time, the number of options for configuring systems has increased, as has the choice of vendors and operating systems. The most fundamental change that has occurred during the past decade, however, has been a shift in emphasis from the acquisition of large centralized facilities to the definition of an integrated environment of personal and group facilities. As a result, a need for a coordinated research support policy has arisen to ensure that the needs of individual researchers (and groups of researchers) are met while at the same time ensuring that an infrastructure is developed that allows resources to be shared effectively by all appropriate users.

In providing broad-based equipment access, we use the following principles to guide our equipment acquisition:

(a) that equipment selection should be decentralized and guided by the needs of each group of users rather than being centralized and guided by averaging the needs of several groups of users;

(b) that the basic role for a central support organization is three-fold: (1) to provide a flexible infrastructure that prevents groups from becoming isolated or requiring duplicated resources, (2) to provide consultative services to guide equipment selection, and (3) to provide a user services organization to coordinate training;

(c) that the proliferation of different types of equipment should be avoided as long as this does not come into conflict with the goals of element (a), and

(d) that equipment purchases should be done in a coordinated fashion to maximize purchasing power.

The combination of these elements have been translated into a general policy that deemphasizes the use of timesharing central computers, that phases out the acquisition of conventional personal computers, and that has concentrated on the acquisition of high-functionality individual workstations.[1]

CWI has been fortunate to have received support from the ministry O&W (in particular, via the *IAS* program) that has allowed us to make major in-roads in the development of a world-class infrastructure. It is a mixed blessing of the current computing marketplace, however, that each year brings new equipment that not only allows people to work more productively, but that also expands the type of work than can be accomplished. We feel, therefore, that it is essential that the path upon which we have embarked be continued, and that our researchers (in particular) and the Dutch research community (in general) can continue to benefit from the advancements made in the computing arena. This continued support is important for two reasons: first, the ability of a research institution to attract and maintain top-level personnel is directly related to the level of facilities support they receive; second, once hired, the ability of a particular researcher to produce and disseminate information regarding his or her research is directly related to the on-going provision of adequate research equipment.

It is for these reasons that we intend to continue to both modernize our present facilities and to expand those facilities to allow us to more productively address the needs of our researchers and visitors.

[1] Our definitions for each of these classes of equipment are given in section 4.1.

3. Present Situation at CWI

CWI presently maintains an infrastructure that supports approximately 100 local computer workstations of the class of a Sun SPARCstation, one central timesharing computer, and two parallel processors for various types of research and support use. A summary of the major research equipment in place at CWI is given in Table I.

The majority of the placement of equipment at CWI has occurred during the past two years. In this period, we have attempted to resolve fundamental problems with the quality of equipment support that was given to our researchers. We have now reached approximately the half-way point in providing all of our researchers and support staff with facilities to meet their basic computing needs. While this change has been very productive, it is important to note that in spite of intensive acquisitions, the average MIPS rating of CWI's computers dedicated to research use is less than 10; when one considers that current (1990) generation equipment has a computing speed of between 25-30 MIPS, it is clear that a significant gap exists between our current infrastructure and one that is "state-of-the-art." For this reason, we see a continual need for support needs to both expand and maintain the systems available to our researchers.

TABLE I: CWI Central Computing & Workstation Facilities (Summer 1990)

Description	*Number Installed*	*Performance (total/per unit)*	*Storage*		*Main Users*
			Memory	*Disk*	
Alliant FX/4	1	40/10 MIPS	32MB	1.5GB	Parallel Sys. Research
Harris HCX-9	1	9 MIPS	8MB	.9GB	Research Timesharing
SUN-4/490	1	21 MIPS	32MB	2.0GB	Research Fileserver
SUN-4/280	4	7 MIPS	16MB	4.0GB	Research Fileservers
DEC VAX-3500	1	3 MIPS	16MB	.3GB	Amoeba Fileserver
DEC VAX-3200	1	3 MIPS	16MB	1.0GB	Amoeba Fileserver
SPARCstation 1	40	13 MIPS	12MB	100MB	Research Workstations
SPARCstation 1+	22	15 MIPS	12MB	100MB	Research Workstations
DECstation 3100	2	16 MIPS	16MB	100MB	Amoeba Workstations
DEC VAX-2000	6	1 MIPS	4MB	80MB	Amoeba Workstations
DEC VAX-2000	10	1 MIPS	4MB	0	Amoeba Processor Pool
DEC Firefly	2	3 MIPS	64MB	0	Amoeba Processor Pool
IBM 6150	5	3 MIPS	4MB	80MB	Graphics Workstations
SUN-3/260	4	3 MIPS	8MB	300MB	Graphics Data Flow
SGI 4D/50	1	7 MIPS	8MB	170MB	Radiosity Research
SGI 4D/25G	1	17 MIPS	4MB	80MB	Networking Research
Encore 520	1	56/7 MIPS	64MB	2.5GB	Support Timesharing
DEC VAX-11/750	1	0.6 MIPS	8MB	.9GB	UNIX Development
SPARCstation SLC	33	12 MIPS	12MB	0	Support Workstations
NCD X-Windows	20	-	1.5MB	0	Support Terminals

4. Policy Objectives for Supporting Mathematics and Computer Science

In order to put our research need in a common context, we will first review our notions of architectures and present our definitions for the various classes of computing that exists. Later sections will make use of these definitions in discussing our policy directions for computing support. (Note: this section is a summary of a report *Supporting the Research Equipment Needs of Mathematics and Computer Science;* D. Bulterman, CWI Report CST-89.03.3, 16 October 1989.)

4.1. Computer Architectures for Supporting Research

In terms of classical computer architecture, a computer can be defined as a single entity that is composed of an *instruction processor*, a *multi-level memory system*, and an *input/output facility*. If instruction processors were infinitely fast, memory systems infinitely large, and input/output processors infinitely flexible, then one could select a single computer to provide support for a wide range of applications. Unfortunately, instruction processors are not infinitely fast; at best they can process only a few million instructions per second. Similarly, memory systems are not infinitely large nor are they infinitely fast; even very large systems contain fast primary memories that are typically smaller than a few tens of megabytes and slower secondary memories of a few gigabytes. Input/output systems are also not infinitely flexible: data transfer is usually done by some sort of mechanical device (and mechanical devices are slow) and the user is typically forced to enter and receive data through relatively cumbersome devices that bind the user to the machine rather than bonding a machine to a user.

As a result of the physical realities of computing technology, different computer architectures have been devised to compensate for the lack of desired performance in the three above areas. Recent advances in technology have made a wide range of computer architectures available. Some architectures use multiple instruction processors to provide a fast working environment. Other architectures provide limited facilities to a user that are specialized to a particular task. Yet other architectures try to balance the needs of many types of applications across a ''family'' of processors. In this section, we give a general taxonomy of computing types that we feel have a place in a research environment.

Figure 1 illustrates a partitioning of different types of architectures that could be used to support general and particular research needs. Three classes of machines exist at the top level of this partitioning: *multi-user systems*, *single-user systems*, and *special-purpose systems*. Multi-user systems are designed to support a number of different users, either serially or concurrently; they range in complexity from minicomputers to supercomputers. Single-user systems are designed to serve a single owner/user, although each user may run different types of applications; typical examples are personal computers and scientific workstations. Special-purpose systems are specifically built to perform a fixed task or set of tasks; these systems (which are still quite rare) are not general purpose computers of the multi-user or single-user class, but specially designed custom computers that implement a single algorithm or class of algorithms.

Multi-user systems can be divided into two sub-classes: *timesharing machines* and *dedicated-application machines*. A timesharing machine is designed to allow multiple users to use a single logical architecture at the same time. This sub-class may be implemented using a computer with a single instruction processor (such as a VAX-11/780 or an IBM-3081) or it may be implemented using an array of instruction processors (such as an Encore Multimax or a Sequent Symmetry). In this class, the user is typically not aware of how the timesharing model is implemented—that is, whether a single- or multiple-CPU model is used—except in terms of overall performance advantages. In contrast,

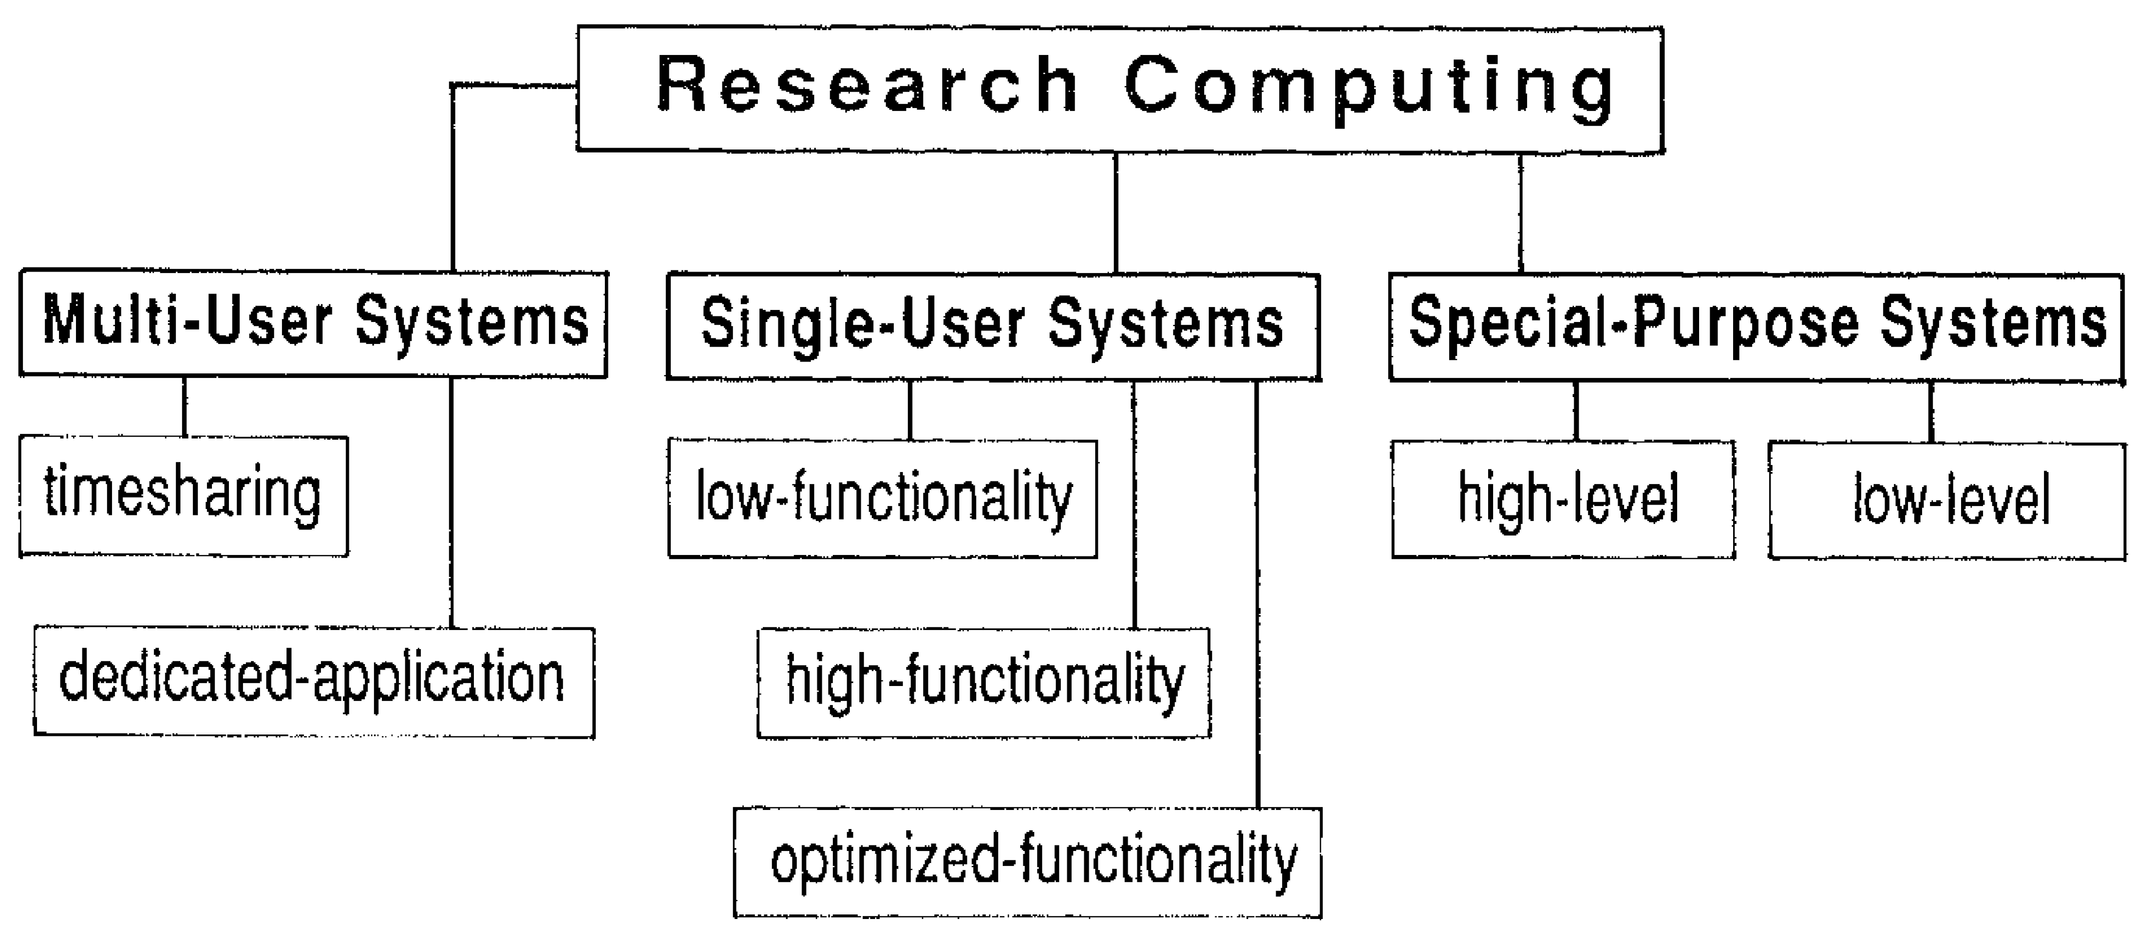

Fig. 1: A taxonomy of computer types for supporting mathematics and computer science research.

dedicated-application systems are used for only a single purpose at a time, but are built to service multiple users over time. (Implementation architectures for dedicated-application systems include supercomputers, general-purpose data flow computers, and parallel processors that are used to execute parallel algorithms.) Both timesharing and dedicated-application machines are meant to be shared resources; the differences between them are in how they are programmed and how users access the machines. Note that while a machine such as a Multimax or Symmetry can be used as a timesharing host or as a dedicated-application machine, a full machine can not be used for both simultaneously.

Single-user computers are intended to service the needs of a single individual rather than as a shared resource available to many users. Machines of this type can be divided into three sub-classes: *low functionality, high functionality,* and *optimized functionality.* Traditional personal computers such as the IBM-PC/AT and the Apple Macintosh are good examples of low functionality systems: they typically provide a working environment that allows one task to be undertaken at any one time. In contrast, personal workstations such as the Sun SPARCstation-1, the Digital DECstation-3100 or the Sony NWS-751 (to name only a few) are examples of high functionality systems: their working environment allows users to manipulate several tasks at once via sophisticated user interfaces that give access to multi-processing operating systems. The third sub-class, optimized functionality single-user systems, are systems of the type that would be classified as ''high functionality,'' except that they are further optimized to perform one or more aspects of an application's needs. This may include computer graphics transformations, high-speed numerical operations, or special-purpose software applications (such as dedicated language processors). Note that in this classification the term ''functionality'' has been chosen with care. The more common term ''performance'' gives a misleading emphasis on the speed at which operations are carried out on a particular machine rather than the range of operations that can be supported: a personal computer may be able to run a spread-sheet application more quickly than a low-performing high-functionality workstation, but the latter architecture can also do other tasks for the user simultaneously, yielding better concurrent functionality.

Special-purpose architectures are those that are designed to implement a single type of algorithm rather than general-purpose tasks. Note that this is not necessarily a class that contains all computer architecture research: a special-purpose image processing machine that implements (in ''hardware'') a single algorithm is not a study in computer architecture but a study in image processing. There is every indication that the decreasing costs of hardware fabrication and the increasing performance needs of several types of applications (such as signal processing or distributed control systems) will allow a researcher to build custom machines in the future. There are currently only a handful of these machines available in research environments, but we expect this class to have moderate (but significant) growth during the next five years.

CWI has built an infrastructure that consists primarily of high-performance workstations that are linked together via a set of local area networks. These networks provide for access to other workstations and to our central computing facilities. At present, our central facilities consist of one processor for supporting general timesharing work and two parallel processors that can be used as research tools. A number of special-purpose architectures also exist for particular projects.

4.2. A Six-Year View of CWI's Architectures for Research Support

During the next six years, most of the models of computing described in section 4.1 will need to be supported. Of these models, we expect that the personal workstation will dominate the computing environment in terms of number installed and user access, although we also see a lasting need for dedicated-application multi-user machines. On the other hand, we expect to phase out traditional timesharing services and to eliminate support for low functionality personal computers. Finally, we see the area of special-purpose systems growing substantially during the next decade. In order to support the expected growth and contraction of various types of equipment, a general policy has been developed for use at CWI. Each of our specific policy goals are given for the three classes of machine architectures defined in section 4.1 in the following paragraphs.

4.2.1. Multi-user systems

Although there will be an economic need for timesharing services for several years to come, we feel that no new large-scale timesharing computers should be acquired for a research setting. Instead, resources should be shifted to providing users with as many single-user systems as possible and with building an infrastructure that can support wide-scale application of the workstation model. This being said, it is important to realize that a transition period will exist during which it is economically and organizationally impossible to provide each timesharing user with a personal workstation. During this period, rather than acquiring or upgrading existing single-CPU, multi-user systems, timesharing services should be shifted to machines than can ultimately be used as dedicated-application multi-user systems. An example of this is the short-term use of a shared-memory parallel processing system (such as an Encore Multimax) as a timesharing host and then having that same machine serve as a basis for parallel systems research after the need for timesharing services have ended. Note that we expect to access very large facilities (such as supercomputers) at a national or international level rather than at a local level; the limited needs for the production-style use of such large central facilities do not justify their local acquisition, maintenance and support.

4.2.2. Single-user systems

It is clear that the major component of any research equipment acquisition plan should be for single-user computing systems. The relatively low acquisition cost of these machines combined with the facilities they present to the user make them indispensable in a research environment. The availability of single-user systems has made the central policy issue surrounding them not: ''should we buy them?'' but: ''what type do we buy and how many?'' We will provide our views on these two questions in this section.

In section 4.1, we defined three sub-classes of single-user systems: low functionality, high functionality and optimized functionality. As a matter of policy, we find that low functionality systems typified by the IBM-PC and the Apple Macintosh (which we here-after simply call PCs) have no long-term place in a research setting and that their use should be discontinued as soon as is practical. PCs are unattractive because they provide software environments that foster isolation of users rather than for coordinated, decentralized use. They also require a separate support structure and software base, raising support costs. Instead of continuing to support these facilities, resources should be put into developing an infrastructure to support high functionality and optimized functionality workstations. The PCs that remain should either be sold or used as replacements for home terminals.[2] (The reader is directed to the previously referenced report of a complete discussion of relative value of PC's in a research setting.)

The bulk of the investment for equipment should be for single-user high functionality systems, here-after simply called ''workstations.'' This investment will consist of single workstations and related infrastructure equipment. This latter category includes file servers, disk storage, network facilities, printing facilities and backup storage. In term of selecting facilities, two issues dominate our planning: one is the selection of vendors for single-user systems and the other is the rate at which workstations are integrated into our environment. Each of these are discussed in turn in the following paragraphs.

In selecting workstation vendors, a balance must be struck between the desires of individual research groups and the need of the institute as a whole to be able to manage equipment acquisitions. In order to provide some degree of flexibility, we allow the acquisition of several types of equipment provided that each such type presents a reasonably consistent user interface and underlying system structure. In specific terms, all workstations at CWI need to support the UNIX[3] 4.3BSD networking protocols, thin-wire ethernet interfaces, NFS network file sharing software and the window system X from MIT. We have found that the presence of this standard software is adequate to support a fairly broad range of architectures.

The second issue regarding the acquisition of workstations is the rate at which these workstations are integrated into an environment. Given adequate financial resources, we restrict our acquisition rate to no more than 20% of our total workstation requirements in any one year. We do this because of the rapid pace of development of workstation products and our own limits on usefully supporting a massive influx of computers. If one considers the changes expected in the workstation market, then if 50% of one's workstation needs are supplied at one time, one can expect that within 2 years, the bulk of one's facilities will be totally out of date. Even if stability was expected, then we have noticed that new users of workstations require a great deal of support, and that our support facilities would be spread too thin to adequately provide this assistance. (The research world is full of examples of institutions ordering large amounts of equipment, only to have that

[2] The continued use at any level of the PC's, although potentially attractive because existing systems seem ''free,'' is in reality an expensive activity: software needs to be acquired, support personnel need to be allocated, and the equipment needs to be maintained. More importantly, however, the different work environments provided by PCs and higher-functionality systems means that conversion problems will continue between, say, an integrated secretarial support environment and a stand-alone home environment.

[3] UNIX is a trademark of AT&T/Bell Laboratories.

equipment sit idle in hallways waiting to be installed.) In terms of long-term workstation commitment, we expect that the lifetime of a typical workstation will be 4-5 years if it is upgraded adequately, so that, in fact, 20-25% of an institute's total facilities need to be replaced annually once a steady-state condition has been reached.

In addition to support for general-purpose workstations, we also see the need to support optimized function single-user systems. These systems must be chosen with care, although it is clear that certain research problems or certain contractual obligations with external partners will necessitate the purchase and support of equipment that deviates from the standard systems selected. It is our policy to allow these systems upon demonstration that they are absolutely necessary and upon demonstration that they will not negatively impact the infrastructure provided for the institute as a whole. Again, the primary goal of providing equipment support should be to facilitate research in a practical way, not to alter the research to facilitate the support structure.

As a final note on single user systems, there has been a great deal of interest in a new class of computer terminal call a *windowing terminal*. The best example of this type of terminal is a graphics-based terminal that runs the X protocol. These terminals are attractive in that they extend the life of many timesharing systems by providing the user with an ''almost personal workstation'' environment while still providing the administrative and economy-of-scale advantages of a single timesharing host. (These terminals also have the advantage of being user-environment friendly, in that they are small, produce little heat, and usually produce no noise.) While we find these terminals useful during a transition period, their per-seat costs are only attractive if one neglects the costs of the timesharing host to which they are connected. We therefore discourage their acquisition except in particular environments—such as secretarial support—where there is a need to centralize data on one host and where there is little expectation of a growth in computing requirements by their users.

4.2.3. Special-purpose systems

One of the most important growth areas in hardware development will be the ability of research users to design special-purpose architectures to implement a single algorithm or a class of algorithms. Although it is tempting to think of these systems as the exception rather than the rule, we expect that research institutes will need to support this type of machine development if they wish to maintain a leading position in the research world. This has the consequence that resources must be diverted to establishing system development and modelling stations, as well as providing for the testing and evaluation of the systems produced. From a practical perspective, it will probably not be economically feasible to engage in actual system production within the next few years unless a partnership is established with a new or existing production facility.

5. Projected Needs, 1991-1996

5.1. Basic Assumptions

The research priorities at CWI are given in the policy document ''Beleidsnota 1990-1995'' and a detailed description of all of CWI current and planned research projects is given in its document ''Scientific Programme 1990, Long-Range Plan 1991-1995.'' In the sections below, we describe the needs of the Institute based on these project descriptions. This discussion follows our internal method for supporting research: while all research equipment purchased is based on ''per-project'' needs, all equipment selection is divided into *general infrastructure equipment* and *special needs equipment*; this distinction reflects the degree to which exceptions can be made to our standard facilities requirements, as illustrated in the following sections. We begin with a discussion of the needs of the general research infrastructure at CWI; this infrastructure serves the research needs of most of the project groups by providing coordinated computing facilities and a general support structure. Next, we highlight four project areas that have special equipment needs that fall outside of the general CWI infrastructure; these projects are: Multimedia Systems, Scientific Visualization, Computational Geometry, and Image Processing. Finally, this section discusses the relation of the independent foundation Computer Algebra Nederland (CAN) to the infrastructure maintained by CWI.

In discussing our research support policy, we will concentrate on the types of equipment that are necessary to support CWI's research and on the timing issues surrounding the placement of that equipment. In order to put these needs in a tangible context from a subsidy point-of-view, we also include cost estimates for the 1991-1996 period with each of the categories in sections 5.2 and 5.3. Since it is nearly impossible to project pricing and cost trends for more than two years into the future, we base our price estimates on the current value of equipment for the initial period of this plan, and then assume a gradual price reduction in later years. As a result, price information should be used as an indication in the years following 1993.

5.2. Research Infrastructure

In considering the general needs of our users, we first review our projections for staffing levels during the period 1991-1996. The numbers are important because of our need to equip all technical staff members of CWI with adequate facilities so that they can be integrated into the total Institute's infrastructure.

As an indication of our overall infrastructure needs, Table II presents our most recent staffing estimates for the year 1993; this is the mid-point of the period for which this policy document has been written.

TABLE II: **CWI Staffing Projections, 1993 (in fte)**

Classification	Projected Number
Mathematics Researchers	57
Computer Science Researchers	57
Systems and Applications Programmers	26
(Non-technical) Support Staff	50
Totals	190

As can be seen, of this group, 140 fte are technical staff members and 50 fte make up the non-technical support staff. We envision that the long term needs of this total group includes the provision of basic research computing support, the provision of computers to make use of our internal and external communications facilities (such as electronic mail, document composition and presentation, and local and remote file access), and the provision of equipment for technical support and training.

While a great deal has been done to improve CWI's central infrastructure in the past two years, we have barely been able to keep pace with developments in the technical marketplace, and foresee a need to continue the growth and refinement of our facilities through the period 1991-1996. During the next six years, we see six areas that will require major capital expenditures: general research workstations, central fileservers (and associated disk and tape facilities), local network facilities, central shared computing facilities, central input/output production facilities, and library and research management automation facilities. Each of these six areas is considered in the following paragraphs. (Note: the budgetary implications of these areas are summarized in section 6.)

5.2.1. Research Workstations

The major component of our equipment acquisitions will continue to be for high-functionality computer workstations. While we include this category as a general facility, each of the workstations that are purchased are assigned to researchers who need computational facilities for their particular projects. Since, however, the use of a workstation provides not only local research support but also a ''window'' into the entire research support facility of the Institute, we group the acquisition of typical user-level workstations into one group. Experience has shown that approximately 20% of our total workstations needs are acquired each year; given a workforce of 190 researchers and staff members, this means an annual acquisition rate of roughly 40 units. Of the 40 units, approximately one-half are purchased for the equipment-intensive projects listed in section 5.3, and as a result these are not counted here. This results in an estimated acquisition rate of 20 units per year. (Note that in initial years of a project, use is often made of general machine rather than special ones, so that a strict division cannot be made.) Once a workstation is acquired, we follow a policy of *shifted placement* rather than *rotated placement* of workstations at CWI: this means that we always give new workstations to the most computationally intensive users each year, and then shift their present equipment to other users. In this way, all users receive updated facilities at least every four-to-five years. (Exceptions to the shifted placement policy occur in years when it is necessary to upgrade a series of older equipment sooner than the annual rate mentioned above. This is usually done to reduce the costs that grow when many different versions of computers need to be supported. We expect this to occur every two or three years.)

In terms of particular acquisitions, CWI maintains a workstation evaluation committee that keeps track of recent technology developments and recommends individual workstations for purchase each year. Requests for support, based on current needs, are passed to this committee by the respective research groups, at which point a determination is made on the yearly specifications for our general purpose workstation choice. During the past year, we have received an increased number of requests for workstation-based computing for those researchers who maintain computer access off-site. While this latter group as traditionally used Apple Macintosh PC's and high-speed modems to access our central facility, we expect that, starting in 1991, there will be a growing demand for (inexpensive) workstations to replace personal computers for this work. While the costs of these facilities will be slightly higher than new personal computers, we expect that the increased functionality and the cost-saving aspects of duplicating the on-site computing environment will justify this extra expense.

In summary, we expect to continue to acquire approximately 20 workstations per year for general research use throughout the six year period of this plan, and also expect to augment this amount with an additional 10-15 lower cost workstations to satisfy remote access requirements of our researchers. We have noticed that prices for state-of-the-art workstations have been nearly constant at a price of approximately *f* 40.000, with each year bringing vastly improved performance; we expect prices of inexpensive workstations to stablize around *f* 9.000 within a year. Based on these prices, we estimate a six-year investment of *Mf* 5.30. Of this total, we expect that 20% of the newly acquired equipment will be for direct replacement of existing facilities (for reasons of support and maintenance manageability), and 80% will be to directly expand the types of equipment available to our researchers.

5.2.2. Central Fileservers

It is a requirement of current workstation technology that each new workstation be supported by adequate fileserver facilities. These facilities include the fileserver computer and associated disks. They also include adequate backup facilities to ensure the integrity of the information keep on the fileserver. It has been our policy to acquire a new fileserver each year to serve the newly acquired workstations. We also attempt to keep the amount of disk space across all file servers up-to-date in terms of disk speed and disk density. At the same time, in order to reduce maintenance costs, we also attempt to consolidate multiple fileservers into one machine when new generations of technology are available, as long as this consolidation does not cripple research when one fileserver becomes inactive because of hardware or software problems.

The supplier for fileserver services has traditionally been the same as that year's vendor of workstations. We expect that in future years this will change as emerging standards in the operating systems area make the acquisition of third-party fileservers more attractive. We expect that this will eventually reduce fileserver prices, although we do not anticipate a dramatic drop in prices for another three years.

Our budgetary forecast is for an expenditure of *f* 375.000 per year on new fileservers, additional disk space, and consolidation and replacement costs. We expect that 20% of our annual expenditures will be for replacement equipment (such as new disks and disk controllers) and that approximately 80% will be for new and expanded facilities to serve new workstations.

5.2.3. Network Infrastructure

In 1988 CWI installed a building-wide network to service all users of computing equipment. This network consists of three parallel thin-wire ethernets. The networks are essential to the operation of our computing infrastructure, since they carry all data between the workstations and their related fileservers. In order to be kept reliable, we need to invest approximately *f* 100.000 in new network controllers and in network troubleshooting equipment. This amount is sufficient to instrument one network per year. Given this level of investment, we expect that the current infrastructure will be adequate for the coming three years. In 1993, we expect to replace major portions of the network with fiber-optic network links, since new generations of workstations will need vastly improved network performance over the current 500-kilobits effectively provided by an Ethernet. We expect the transition process to be spread over a several year period, until all of our current workstations and computing equipment have been converted. We expect that the cost of this conversion will be approximately *f* 900.000. Our experience is that even after this expenditure, an annual estimate of *f* 100.000 will be required to maintain a complex infrastructure. As a result, we expect a total six-year investment for networking support of *Mf* 1.50. Of this amount, approximately 15% of all of our funding over the period will go toward replacements and the remainder to enhancements.

5.2.4. Central Facilities

At present, CWI maintains three shared-use central facilities: one general timesharing host (a Harris HCX-9), one vector-based parallel processor (an Alliant FX-4), and one shared-memory parallel processor (an Encore-520). It is clear that, in spite of the general trend toward workstation-based computing, there will be a continued need for shared computing facilities if we are to effectively support a broad range of research problems. In particular, the area of parallel computing will continue to grow as more researchers concentrate on the use of parallel algorithms. In order to support this trend, we plan to concentrate our central facilities on parallel computing by replacing, upgrading, and expanding the three machines now in use.

Of these machines, the Harris is nearing the end of it useful life; it is a relatively slow processor (by current standards) and is not fully software compatible with our workstation-based infrastructure. In line with our policy of not acquiring new single-processor central facilities, we do not expect to purchase a direct replacement for the Harris during the period 1991-1996. We will attempt to keep the Harris in service as long as possible by migrating faster disks to the architecture, which we feel will extend its life for another year or two. At the end of this period, we expect the penetration of workstations to be far enough along to allow us to discontinue the computer without need for a replacement.

We expect to upgrade the Encore to that company's new technology in 1991; this upgrade will be required to allow us to use Encore's new operating systems software, all of which is being developed around its new RISC technology. At the time of its upgrade, we expect to also expand the number of processors available from 8 to 16, so as to make the machine more attractive for research into parallel algorithms where a moderate number of processors can be used to solve particular problems in mathematics.

The Alliant will need to be upgraded or replaced in 1992; at this time, the current processor-loan agreements with the manufacturer will expire, meaning that the architecture will need to be purchased (an unattractive option, given the age of the equipment) or it will need to be replaced with new equipment from the same or a different manufacturer.

While upgraded shared-memory and parallel vector facilities will support a wide range of work, they cannot support research into massively parallel algorithms. As a result, we expect that a new central compute server will be required mid-way through the 1991-1996 period to address research into parallelism over a very large number of processors (on the order of 100s of processors); although current technology to support this research area is still prohibitively expensive, we expect that by 1993, prices will be more in line with our expected budgets. Note that in the period 1994-1996, we expect to upgrade the equipment purchased in earlier years to protect its investment and to keep pace with appropriate technology.

While it is difficult to make precise budgetary predictions for large central machines, we estimate that, on average, we will need to purchase one such machine per year (as detailed above), each with an average cost of approximately *f* 500.000. (We expect this cost to decrease slightly in later years). Of this total, approximately 35% of the investments will modernize existing facilities (such as the Encore and Alliant) and 65% will be to expand the capabilities of these machines.

5.2.5. Input/Output Facilities

While a great deal of effort has been placed into the development of a workstation-based infrastructure at CWI, we are currently quite deficient in our generalized capability to present results at lectures and meetings, and in capturing other than typed-in data. At present, CWI provides two high-quality postscript laser printers and one postscript phototypesetter as central output facilities and approximately 10 low-quality laser printers for preview use. Our input facilities (other than the keyboard and mouse) are restricted to a single low-resolution digitizer. Given our ability to produce state-of-the-art color images on display stations, we need expanded facilities to communicate these results to larger audiences.

Our priority needs with respect to output production are for a capability to produce color postscript output via a color printer, a color slide production facility, and a video projection system coupled to a video tape recording system. While some of this equipment will be purchased under auspices of the multi-media project (see section 5.3.1), most of the facilities need to be acquired for general use. For example, we expect to acquire a color printer in late 1991 or early 1992, as soon as the resolution approaches the 400-pixel (per inch) level now supported by our high-quality facilities. At the same time, we see the need for expanded printing facilities to accommodate the increased demand from our research and support users. We plan to acquire a higher-speed printer in 1991 and another in 1992, each at a cost of approximately *f* 70.000. We also plan to gradually replace the existing low-quality printers when each of them reaches the end of their life cycles. (We expect this to be at a rate of two per year during the entire period.) Finally, we expect that our present phototypesetting facilities will need to be enhanced before 1995 with a version that also supports color printing, meaning an upgrade of our publishing facilities at approximately the same time.

In terms of input systems, we feel that the major general need is for an input scanner that can be used to digitize high-quality images. This facility is discussed in section 5.3.4, although it will also be made available for Institute-wide use. We also have a need for an OCR facility to reduce our dependence on typed-in text as an input mechanism.

During the entire period, we expect to invest *f* 500.000 in input/output facilities; this includes a yearly single-item purchase of a major output device (such as a laser printer for phototypesetter) of an approximate value of *f* 70.000 and an input device set of approximately *f* 15.000. Of these amounts, approximately 20% will go toward the replacement of existing facilities and 80% to expand the facilities currently available.

5.2.6. Research/Management Support Facilities

At present, CWI uses a generally segmented research support computing structure that uses personal computers within the support departments that operate separately from the research facilities at the Institute. While this policy made sense in 1987, it is a policy that can not be justified given the general desire for improved efficiency through enhanced communication among all sectors at CWI. As a result, we expect to replace significant portions of our current PC computing infrastructure with low-cost workstations during the period 1991-1996. Most of this replacement will not require additional equipment funding, since our policy is to make older research equipment available to the support sectors when new research equipment is required for high-speed work. There are two exceptions to the shifted placement policy they will result in a direct expenditure for equipment. One exception is for a central automation facility that can support the automated interrogation of our library facilities and the maintenance of research project administration services. The second exception is for new facilities to support the financial administration sector as they make a transition from PC facilities to workstation facilities. (This latter move is required to improve integration into the Institute's infrastructure and to remove the need to support multiple classes of computing systems.)

We expect that the support sector will require an investment of less than *f* 100.000 per year during the period 1991-1996, with the exception of the acquisition of a central facility for library automation in 1991-92, which will cost a total of approximately *f* 500.000 for hardware and software during those two years. This investment will allow the CWI library to provide UNIX-based interrogation of our local collections and will also provide a link to the PICA-maintained library catalogue information maintained external to CWI. (Discussions between PICA and CWI on integration of these facilities had begun as this document was written.) Integration of the financial sector into the CWI infrastructure will depend on the availability of appropriate software; we expect this to occur before 1994. The total cost of this facility, including software, will be approximately *f* 300.000 (including workstations).

5.3. Project-Specific Equipment

In the sections below, we highlight four areas that are representative of equipment-intensive projects at CWI. As was mentioned above, all of the central workstations and shared computing facilities acquired for use at CWI are, in essence, purchased at the request of individual research groups. The projects listed below are not listed in these general totals because their equipment needs are dramatically broader than most of CWI's work.

Given that the purpose of this report is not to justify research projects but to detail a research-support policy, each of the projects discussed below will not contain a detailed technical description. To put each projects needs in context, however, we will preface each section with a short summary of the goals of each of the research areas and a cross-referencing to our long range scientific plans. We then summarize the projected needs of these projects during the next six years.

5.3.1. Multimedia Systems

Project Summary

Multimedia research at CWI is a long-term coordinated effort among several projects that began in 1990. The goal of the project is to investigate the use of multiple information forms (such as text and sound or graphic data and voice annotation) to expand the possibilities of human/computer interaction. This research can be subdivided into a number of focus areas, all of which are being coordinated to use a common technology base. The focus areas are: *distributed operating systems, computer networking, database design, user interfaces*, and *interactive systems*. Research in the area of distributed operating systems is concentrated on use of the Amoeba system developed jointly by CWI and the Vrije Universiteit; this system provides a fundamental communications and process interactions base for all of the multimedia projects. Research in the area of computer networking is concentrated on efficient implementations of protocols for local and wide-area interconnections for multimedia hosts. The database design research focuses on efficient implementation of multimedia information stores that contain (synchronized) text, image, voice, and graphic data. Work in user interfaces involves the coordination of various data streams in terms of a common user access point. Research in interactive systems involves the overall coordination of the expression of various types of concepts in terms of a flexible assortment of data types. Multimedia work is divided among the departments AA, CST, and IS, with supplemental contacts in BS and AM.

Project References

AA-2 (Distributed Systems/Amoeba), AA-3 (Computer Systems and Ergonomics/User Interfaces), AA-4 (Distributed Adaptive Information Systems/Database Design), IS-2 (User Interfaces), CST-2 (Network-based Task Submission and Performance Monitoring).

Equipment Requirements

The equipment requirements for the multi-media research projects can be divided into six categories: multi-media workstations, video capture and display equipment, audio capture and display equipment, high-speed machine interconnection equipment, performance monitoring equipment, and local disk space. Note that fileservers are not included in this list; we assume that the cost figures allocated for the workstations include a fileserver ''tax'' that can support either a local fileserver or access to shared facilities. In the paragraphs below, we will summarize our needs in these areas.

- *Multi-media workstations:* The principal requirements for a multi-media workstation is that it have extremely fast processing capabilities and that it adequately support video and audio inputs and outputs. These stations also need to be of the class that supports the Amoeba operating systems as a base for communications support. Given the start-up nature of the multi-media project, we expect initial workstations needs to be moderate, but expect that over the six year period of this report that each researcher will consume two generations of workstation equipment. In the first year, we expect to expand the range of architectures supported for production-style by providing the Amoeba with the first of the multi-media workstations acquired. (This is necessary not only to begin the project, but also to ensure that the multi-media operating system kernel is ready in time for extended use in 1992. The Amoeba workstations would replace 16 obsolete DEC VS-2000's, which have been in use for approximately four years by this group. For all group members, we expect a total of 30 workstations will need to be purchased in the six year period, with each workstation having a cost of *f* 55.000; the bulk of these workstations (65%) would replace existing equipment.
- *Video capture and display equipment:* An essential component to multi-media research is the capture, reduction, manipulation, and display of video information. Given the desire to study data compression algorithms for video signals, we see a general need to support video input for a relatively large percentage of our multi-media workstations. This input will either come through commercial video cameras or through special purpose terminals with built-in video input. In addition, we will need to integrate output signals as a composite of video and text/graphic information; this will require video adapter cards to many multi-media workstations. As a whole, we expect to spend *f* 150.000 over the six year period of this policy period.
- *Audio capture and display equipment:* Audio capture equipment is generally available on many types of workstations. (Most of these come with microphone input and some rudimentary output facilities.) While this equipment is sufficient for, say, audio annotation, it can not be used for higher quality manipulation of audio data. For this reason, we foresee the need of acquiring better audio input and output facilities to demonstrate the broader integration of sound into a work and production environment. Our expectation is that most multi-media workstations will need appropriate microphone and playback facilities within two years. The total scope of this project includes some 20 loudspeaker sets and matched high-quality input devices, and several demonstration systems available for general evaluation. The total cost over the six year period will be approximately *f* 100.000.

- *High-speed machine interconnection equipment:* While we initially expect to use existing building network services, it is clear that the transfer speed requirements of (near) real-time multi-media work will necessitate a 10-100 fold improvement over existing network facilities. Although it is impractical to rewire our building for this one set of projects, it is clear that a number of local networks will be needed to support prototype development and test transfer and compression algorithms. (Note that discussions are underway with a number of network and workstation suppliers about the integration of their proprietary networks on a research basis; we expect to be able to make a full decision in late 1991.) Since a portion of these new networks will replace existing facilities, we expect that approximately 40% of the *f* 80.000 needed for research purposes can be viewed as replacement costs and the balance as new resources.
- *Performance monitoring equipment:* One of the challenging aspects of multi-media interconnection is the ability to measure the effectiveness of the protocols developed, and to compare these results with the expect theoretical outcomes. To accomplish this, we will require monitoring equipment that is similar to that used for our general purpose networks—although it will need special adapters for the new network technologies outlined above. We expect that this equipment will be important in the initial phases, and that it will become less import in later years, when more emphasis will be placed on applications. The total cost for this equipment will be approximately *f* 100.000.
- *Local disk space:* One final aspect of multi-media equipment funding will be the requirement to store very large amounts of data for transient or permanent use. A portion of this storage will be maintainable at fileservers, although most of this space will need to reside locally for performance reasons. To this end, we expect to acquire between 10 and 151 gigabytes of storage of increasingly fast technology during the six year period. As a part of this storage will be used to replace existing central facilities (about 70% of a total allocation of *f* 350.000), we expect to reserve approximately *f* 100.000 for expanded local storage purposes.

5.3.2. Scientific Visualization

Project Summary

In several CWI projects, the visualization of scientific data plays a significant role in interpreting and understanding results. The processing power of current computing equipment (such as CWI's Alliant parallel processor) provides a means for investigating large and complicated mathematical models; scientific visualization is indispensable in presenting, explaining, evaluating, and illuminating the scientific output of this type of research. The understanding of chaos in dynamical systems has depended in an essential way on numerical simulations and the visualizations thereof. For the past several years, the need for a dynamical systems laboratory for investigating population dynamics, spread of diseases and environmental pollution, the evolvement of physical processes, etc. is felt and formulated in the CWI-policy documents. Apart from scientific support, it give the possibility for demonstrating the CWI-projects to outside partners in contract research, and for popularizing several areas of present day mathematics.

Project References

AM-3 (Nonlinear Analysis and Biomathematics), NW-1 (Evolution Problems), NW-2 (Steady Boundary-Value Problems). Supplemental contacts: AM-5 (Image Processing and Reconstruction) and BS-6 (Image Analysis).

Equipment Requirements

The equipment requirements for the scientific visualization research projects can be divided into three categories: visualization workstations, attached processors, and information recorders. In the paragraphs below, we will summarize our needs in these areas.

- *Visualization Workstations:* Like multi-media workstations, the computers for use in scientific visualization need to be more powerful than the average research workstation at CWI. Unlike multi-media research, however, scientific visualization requires primarily superior graphics processing facilities and the coupling of output to high-speed processing engines (see next paragraph). In terms of the research at CWI, we view machines of the class of a Silicon Graphics as a *minimum* requirement. During the term 1991-1996, we expect that this research will require two central visualization machines similar to an SGI 4D/280VGXB and 10 lower speed machines (such an SGI 4D/210) as developers workstations. The total costs of the each large facility will be approximately *f* 250.000, with additional workstation components costing *f* 60.000 each. (This yields a six year total of *Mf* 1.10.) We will purchase one large visualization machine and one lower speed workstation in 1991. The large facility, which has a vector draw rate of well over a million pixels per second (and thus adequate for at least the next few years), is considered to be an expanded facility, and the workstation components are considered as replacements to existing equipment. We expect that the second large central machine will be required in late 1992 or early 1993.
- *Attached processors:* At present, much of the work that is connected with the general area of scientific visualization is associated with the reduction and presentation of results from supercomputers or high-performance parallel systems. While we expect to make use of facilities already available as CWI central equipment or via supercomputers external to CWI, we also envision the need for local attached processors to accelerate the computational facilities on local machines. These attached processors, which are usually dataflow-oriented arrays of computers that plug into workstations, are required because the workstation itself is often busy with data conversions rather than data production. The processors would be used during initial algorithm investigation across a broad range of mathematical and computer graphics applications. Given the relatively low volume of productions of these units, their individual costs are high. As a result, we expect to need to reserve approximately *f* 175.000 per year for this equipment. While these costs are high, they are minor when compared with the costs of maintaining supercomputer facilities—especially if one considers that the performance difference between attached processors and supercomputers on a per-application basis are relatively small.
- *Information recorders:* As with our general support environment, there will be a need to catalogue information generated on visualization workstations on video recording devices. These devices will be either conventional video tape facilities or video disk facilities. Video disks are preferable because of their greater degree of selectivity. Unfortunately, the increased resolution of visualization equipment may require special-purpose equipment that is different from general facilities acquired for CWI. While we do not expect to need special facilities in the first years of this project, we do anticipate a total allocation of approximately *f* 300.000 over the six year period to support data recording.

5.3.3. Computational Geometry

Project Summary

A great many practical problems contain a geometric component. Examples include robotic movement and placement problems, VLSI routing problems, pattern recognition and orientation problems, and geometric placement problems within computer graphics. The collection of all of these (sub)problems is substantial, with many practical and theoretical issues as yet unresolved. The solutions of many aspects of computational geometry problems involve combinatorial as well as heuristic aspects. These problems have been addressed separately at CWI by the projects *combinatorial optimization algorithmics, computer graphics,* and *computer-aided design.* It is expected that in coming years, research equipment will be coordinated across these computationally-intensive applications to increase overall efficiency.

Project References

BS-1 (Combinatorial Optimization and Algorithmics), AA-1 (Algorithms and Complexity), IS-1 (Computer Graphics), and IS-4 (Intelligent CAD Systems).

Equipment Requirements

The equipment requirements for the computational geometry research project can be divided into two categories: research workstations and attached processors. In the paragraphs below, we will summarize our needs in these areas.

- *Research workstations:* The workstations required for computational geometry also require a substantial graphics capability. As with Visualization and Multimedia research, these workstations will be required to display complex data in spatial terms, using color to highlight individual data characteristics. While the processing power of current-generation workstations will probably be sufficient to service the needs of this group, it is clear that each workstation will need expansion capabilities for supporting special processors (see below) as well as being able to execute a common software base that is often determined by third-party research. In terms of budgetary allocations, we expect to provide initial placement of conventional workstations for this group in the initial years of the project, with a total six-year expenditure of an average of one special workstation per year; this translates to a sum of approximately *f* 700.000, with the major concentration later in the period.
- *Attached processors:* Given the parallel nature of many research approaches in computational geometry and the vast amounts of data to be considered, it is essential that sufficient local processing extensions be provided to this project. While it is difficult to predict the direction of this work, it is expected that at least four to six special attached processors will be required; of these we expect an even division between data flow approaches and parallel approaches. Current costs for a single processor (with software) approaches *f* 125.000 per workstation (above the workstation cost); given general diminishing pricing trends, we expect a six-year total cost of *Mf* 1.15 for all of the processors, with the major acquisitions coming near the end of the period.

5.3.4. Image Processing

Project Summary

CWI has, since 1985, embarked on a major and long term program of image processing research. The work done at the Institute falls into two general categories: *image reconstruction* and *image analysis.* Image reconstruction is the process of inversion or back-transformation of signal data to produce viewable images, as in computed tomography

(CT), synthetic aperture radar, and seismic exploration imaging. Image analysis is the process of extracting usable information and making decisions based on image data; it includes the recognition of image features or spatial patterns, numerical descriptions of such features, and consequent actions such as classification and clustering. The shared goals of both types of imaging research are: research on mathematical aspects of image processing algorithms, using techniques from mathematical and numerical analysis, mathematical statistics, computer-intensive statistics, and computer science; development of algorithms and associated software; pursuit of research contacts with applied scientists (medical investigators, biologists, physicists, geologists) and with industrial research laboratories.

Project References

Current imaging projects include AM-5 (Image Processing and Reconstruction) and BS-6 (Image Analysis).

Equipment Requirements

The equipment requirements for the image processing research projects can be divided into five categories: imaging workstations, compute servers, image capture stations, video disk playback facilities, and video recording facilities. In the paragraphs below, we will summarize our needs in these areas.

- *Imaging workstations:* We project a six year need for imaging workstations that each will exceed the performance requirements of the standard research workstation acquired by CWI. These workstations are characterized by both computational intensity and the ability to support high-speed pixel manipulations. Image workstations also have the need to support very ''deep'' image memories: typically, 24-32 bits deep is a requirement for reasonable color and/or texture representation. During the period 1991-1996, we expect to acquire 8 workstations for these projects, each with a projected cost of *f* 80.000. During the initial years, we will expand the types of systems available to image processing work by acquiring co-processors to existing workstations and increasing memory capacity. In later years, we expect to acquire several new workstations that will provide both faster processing and better support for full color images. We will focus on the workstation as a presentation device for these projects, with special processing requirements satisfied by the compute servers described in the following paragraph. Two of the imaging stations are expected to be purchased in 1991 as expansions to our equipment base,
- *Compute servers:* One characteristic of much of the imaging work to be done at CWI is the requirement to process large amounts of data that has an initially uncertain structure. This means that segmentation on to a standard parallel architecture is difficult, since it is often the structure of the data that determines the ultimate segmentation of the algorithm. As a result, much of the image manipulation research that can be done requires a flexible processing architecture in which the various image transformations can be scheduled on the basis of a particular image. In order to support this research, we expect to acquire two image processing compute stations during the period of this research. In addition, approximately half of the imaging workstations will require acceleration processors to support enchanced pixel manipulation facilities; these facilities will replace existing workstation equipment. In general, we anticipate allocating *f* 500.000 for processor-based extensions to the imaging computing environment, and *f* 200.000 for extensions and replacements to attached processor equipment. During the first two years of this project, we will experiment with small-scale attached processors, and acquire larger-scaled common facilities when this research expands in 1993.

- *Image capture stations:* Each of our image processing projects has a need for the scanning of data from external sources. In order to support this work, we see a need for one high-resolution color scanner, two local low-resolution scanners for color input and two local black-and-white low resolution scanners. (Note that resolution is considered in terms of an X-Y array as well as the depth—or number of bits per pixel—obtained at each location.) Low resolution scanners can typically be attached to a particular imaging workstation, while high-resolution scanners require a separate workstation to control data input. We will acquire one low-resolution scanner for color images and one for black-and-white in 1991 and, based on the experience gained by using these cameras, we will acquire additional equipment in late 1992 or early 1993. The total cost of these scanners (and control processors) will be *f* 550.000 during the six year period.
- *Video disk playback facilities:* Another characteristic of imaging work is the need to access large amounts of data rapidly from a particular workstation. Given the fact that a single image can consume approximately 3 megabytes of data, the needs of cataloguing large amounts of images can easily consume a substantial portion of fileserver and network resources. To overcome this, we want to allocate a number of digital video playback facilities for each of the imaging projects. As with all forms of storage facilities, we expect that at least two generations of equipment will be need per project in the 1991-1996 period. The images placed on disk will either come from local data banks (see the next section for a discussion of video recording facilities) or from research partners. We will allocate a total of *f* 350.000 to the presentation of high-resolution input to the imaging projects, with an initial acquisition of equipment planned for 1991.
- *Video recording facilities:* In addition to playback facilities, we will also acquire a generalized video recording system for at least one station per imaging project. Given the decreasing cost of laser disk recorders, we expect this goal to be realized in 1993. In the meantime, we will investigate the use of digital tape recorders at one station in 1991 to test its usability as a mean of archiving results obtained by our research. We have allocated a total of *f* 200.000 for the acquisition of these recorders, including medium charges, during the period.

5.4. Computer Algebra Nederland

In addition to its own equipment, SMC/CWI provides infrastructure support to a small number of external organizations on a contract basis. One of these organizations is the foundation Computer Algebra Nederland (CAN). CAN has a portion of its computing facilities placed at CWI and a portion divided across several national sites. CWI provides CAN advice on equipment selection and provides system software maintenance support. In addition, CWI provides networking expertise to allow the CAN facilities to be accessed from remote sites. Finally, the foundation SMC acts as a purchasing agent for CAN and manages CAN's disbursements. In all of these interactions, however, the foundation CAN makes the final and formal decisions on the types of equipment they acquire for their own facilities, in accordance with the wishes of their constituent members and member institutions.

6. Budget Projections

6.1. Projected Total Needs

Table III presents a summary of the total budget requirements for equipment investment at CWI for the period 1991-1996. The categories in the table correspond to the equipment needs discussions in section 5 of this report. Two types of equipment acquisitions are noted: *replacement* and *renewal*. Replacement equipment is used to directly replace an existing piece of computing equipment; it is expected that this will be done to improve performance or to increase efficiency of equipment support. Renewal equipment is used to expand the capabilities of existing equipment (either via upgrades or expansions) or it is used to provide new equipment. In all cases, renewal equipment will expand the research being performed.

TABLE III: **Equipment Investment Needs, 1991-1996** *(in Mf)*

Category	*Replacement*		*Renewal*		*Subtotal*
1. Research Infrastructure					
1.1 Research Workstations	1.00		4.30		
1.2 File Servers	0.45		1.80		
1.3 Network	0.25		1.25		
1.4 Central Machines	1.00		1.90		
1.5 Output Facilities	0.10		0.40		
1.6 Support Facilities	0.10		0.80		
Subtotal		*2.90*		*10.45*	*13.35*
2. Project Infrastructure					
2.1 Multimedia Systems					
2.1.a Workstations	1.10		0.55		
2.1.b Video Equipment	0.00		0.15		
2.1.c Audio Equipment	0.00		0.10		
2.1.d High-speed Links	0.30		0.50		
2.1.e Monitoring Systems	0.00		0.10		
2.1.f Disks	0.25		0.10		
Subtotal		*1.65*		*1.50*	*3.15*
2.2 Scientific Visualization					
2.2.a Workstations	0.40		0.70		
2.2.b Attached Processors	0.40		0.60		
2.2.c Data Recorders	0.10		0.20		
Subtotal		*0.90*		*1.50*	*2.40*
2.3 Computational Geometry					
2.3.a Workstations	0.20		0.50		
2.3.b Attached Processors	0.25		0.90		
Subtotal		*0.45*		*1.40*	*1.85*
2.4 Image Processing					
2.4.a Workstations	0.15		0.50		
2.4.b Compute Server	0.20		0.50		
2.4.c Image Capture	0.00		0.55		
2.4.d Video Disk Players	0.00		0.35		
2.4.e Video Recorders	0.00		0.20		
Subtotal		*0.35*		*2.10*	*2.45*
Totals		6.25		16.95	23.20

6.2. Financing Projections, 1991-1996

Table IV provides a summary of expected CWI funding resources for the period 1991-1996 on an annual basis. This table reflects our major subsidy providers (NWO and IAS), as well as our estimates of external funding for the period.

TABLE IV: Equipment Investment Sources, 1991-1996

(in Mf)	1991	1992	1993	1994	1995	1996
NWO Subsidy	1.30	1.40	1.40	1.40	1.50	1.50
IAS Subsidy	2.00	2.00	2.20	2.40	2.40	2.50
Other subsidies	0.20	0.20	0.20	0.20	0.20	0.20
Totals	3.50	3.60	3.80	4.00	4.10	4.20

Notes:

(1) The above-presented amounts are for equipment investments only, and are exclusive of any amounts for the expansion or maintenance of the physical buildings at CWI;

(2) The above-presented NWO subsidy contains only amounts to be used by CWI and are exclusive of any amounts intended for use of or participation in SARA;

(3) Funds from IAS subsidies are to be used only for the expansion of facilities or for the renewal of equipment with ''state-of-the-art'' facilities; normal replacements of equipment are financed through the NWO subsidy.

(4) Irrespective of the regular level of CWI funding, the above-presented funding scheme will be less that 15% of our total NWO funding request. This is the amount that is given on page 35 of the *Meerjarenplan 1991-1994 NWO, Deel I;*

(5) Any eventual reduction in the IAS subsidy will inevitably lead to an increased request through our NWO subsidy channels.

6.3. Yearly Specifications, 1991-1992

1991

Table V provides our current plans for equipment acquisition for the year 1991. This plan is based on the overall needs described in section 5. As in earlier tables, equipment categorized as replacement will be used to maintain existing facilities while equipment listed as renewal will be used to expand existing facilities, either via upgrade or via additional facilities.

TABLE V: Equipment Investment Priorities, 1991 *(in Mf)*

Category	*Replacement*		*Renewal*		*Subtotal*
1. Research Infrastructure					
1.1 Research Workstations	0.15		0.70		
1.2 File Servers	0.10		0.28		
1.3 Network	0.00		0.10		
1.4 Central Machines	0.15		0.20		
1.5 Output Facilities	0.00		0.08		
1.6 Research Support Facilities	0.10		0.15		
Subtotal		*0.50*		*1.51*	*2.01*
2. Project-based Infrastructures					
2.1 Multimedia Systems					
2.1.a Workstations	0.20		0.10		
2.1.b Video Equipment	0.00		0.02		
2.1.c Audio Equipment	0.00		0.02		
2.1.d High-speed Links	0.05		0.07		
2.1.e Monitoring Systems	0.00		0.03		
2.1.f Disks	0.02		0.03		
Subtotal		*0.27*		*0.27*	*0.54*
2.2 Scientific Visualization					
2.2.a Workstations	0.05		0.25		
2.2.b Attached Processors	0.07		0.10		
2.2.c Data Recorders	0.00		0.00		
Subtotal		*0.12*		*0.35*	*0.47*
2.3 Computational Geometry					
2.3.a Workstations	0.10		0.00		
2.3.b Attached Processors	0.00		0.00		
Subtotal		*0.10*		*0.00*	*0.10*
2.4 Image Processing					
2.4.a Workstations	0.08		0.08		
2.4.b Compute Server	0.00		0.04		
2.4.c Image Capture	0.00		0.12		
2.4.d Video Disk Players	0.00		0.02		
2.4.e Video Recorders	0.00		0.04		
Subtotal		*0.08*		*0.30*	*0.38*
Totals		1.07		2.43	3.50

1992

Table VI presents a summary of our acquisition plans for 1992. Given the uncertain nature of technology development, we provide only a summary of the needs of individual projects rather than a detailed specification.

TABLE VI: **Equipment Investment Priorities, 1992** *(in Mf)*			
Category	*Replacement*	*Renewal*	*Subtotal*
1. Research Infrastructure	0.50	1.50	2.00
2. Project-based Infrastructures			1.60
2.1 Multimedia Systems	0.20	0.30	
2.2 Scientific Visualization	0.15	0.30	
2.3 Computational Geometry	0.10	0.05	
2.4 Image Processing	0.15	0.35	
Totals	1.10	2.50	3.60