N° d'ordre :

THESE

présentée

A L'UNIVERSITE LOUIS PASTEUR DE STRASBOURG

pour obtenir le grade de

DOCTEUR DE L'UNIVERSITE LOUIS PASTEUR

SECTION : SCIENCES

par

Frank BECK

Membre Permanent du CERN

CONCEPTION ET RÉALISATION DU CENTRE DE CONTRÔLE POUR L'ACCÉLÉRATEUR SPS DU CERN

Soutenue le 17 Juin 1976 devant la Commission d'examen :

MM. A. GALLMANN	Président
G. SUTTER	Membres
A. FRIDMAN	
G. MONSONEGO	
D. LORD	

THE DESIGN AND CONSTRUCTION OF A CONTROL CENTRE FOR THE CERN SPS ACCELERATOR

F. Beck

"and what is the use of a book" thought Alice,
"without pictures or conversations?"

Lewis Carroll
Alice in Wonderland (1865)

April 1976

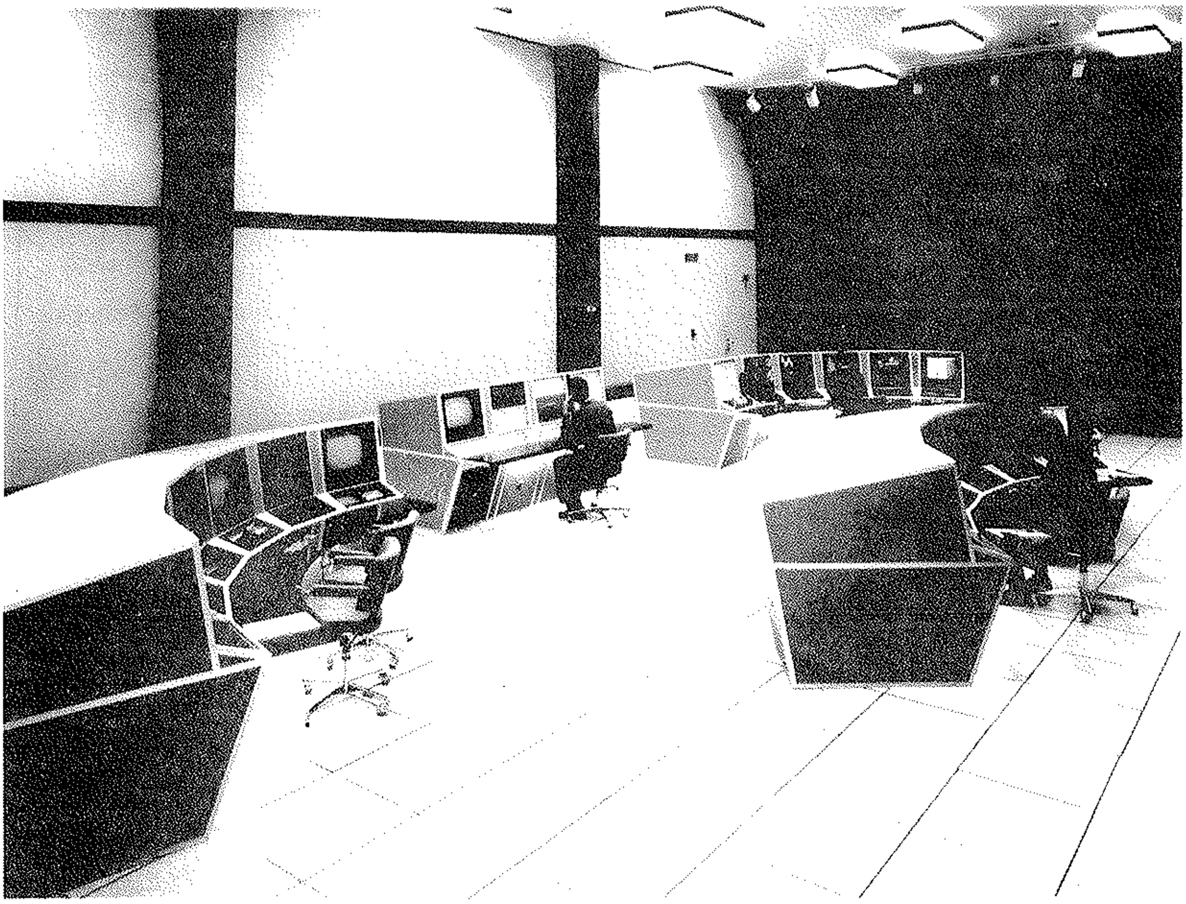

CONTENTS

INTRODUCTION

The following pages are intended to describe the conception and design of a control centre for the CERN Super Proton Synchrotron (SPS) accelerator[1]. This control centre, whilst based on ideas from a number of previous ones, embodies a new over-all concept in control centres, as well as numerous design details which have been conceived and developed especially for the purpose.

The control centre is now substantially complete and is currently being used for the running-in process of the accelerator; but the usage of its facilities is still developing, and, in the light of experience, details of its design will undoubtedly change. Indeed, the system already differs in small particulars from what is reported here: there is some difficulty in describing a complex which is now being adapted for day-to-day use by a variety of engineers and scientists. There is a strong temptation to abandon attempts at documentation in a fluid situation such as this, but what is described is substantially correct.

The rest of this introduction gives briefly the historical background of the new control-room concept, and then goes on to state what are the original elements of the new system.

Chapter I is a very compact account of those parts of the accelerator control system which were external to the work subsequently described, but whose many original features made necessary -- and possible -- the kind of central control philosophy described later. This chapter has been entitled "The environment of the problem", and it has been made as brief as is compatible with the requirement of guiding the reader quickly from his general knowledge of the state of the art into a specialized understanding of this particular control system.

Chapters II to V describe in turn the considerations which led up to design of the individual control-room components (consoles, graphics, and interactive devices), discussing in each case the alternatives offered by the availability of commercial components, the physical constraints, and the practicability of making needed devices specially for the project. This last is an important engineering constraint. The difficulties of making many commonly available devices (such as cathode-ray tubes and integrated circuits) to certain specifications are too great even for a large organization such as CERN, and in many cases a good idea must be abandoned on the grounds of unpracticability unless it can be implemented as a combination of commercially available components.

Chapter VI is an attempt to show how the complex mentioned in the previous chapters is applied by software to make the sum of the individual elements, which have been explained throughout the document, into a control centre. This is really what the whole thing is about: it could be considered conversely that the previous chapters describe elements developed to make the contents of Chapter VI possible.

The seventh chapter is a brief excursion into the future of control rooms of this kind. It suggests some lines of development which could be pursued by anyone following the author's ideas.

Control rooms for accelerators

The need for a control room for an accelerator has not always been obvious. When accelerators were smaller, it was quite feasible to walk round the machine switching on and adjusting components, and remote instrumentation was only needed when there was a radiation risk for the operator. Personnel access and television monitoring were the chief purposes of a control desk -- if one was supplied at all.

As machines became larger there arose a tendency for all the equipment to be brought into one room, except that actually involved in the acceleration of the particle beam. This enabled much of the equipment, for example the control units for power supplies or accelerating cavities, to be controlled in a central situation by an operator who could at the same time watch the results of his manipulations on oscilloscopes connected to beam monitors. Such central complexes, whilst they were called control rooms, were actually equipment rooms, with the controls for each piece of equipment mounted on the front panel.

Soon there arose an organizational need for something more elaborate than the direct control of individual power supplies. Stepping motors enabled various units to be controlled remotely, and groups of correcting elements could be ganged for simple control operations.

Introduction of computers

Initially, there was some reluctance to use computers. Early computers were expensive, on-line interfacing was tedious, and, in the absence of modern interrupt structures and other instruction refinements, slow. The unreliability of the earlier computers was acceptable when they were being used for off-line batch calculation. In a real-time environment it was out of the question.

Gradually, with the arrival of multiprogramming computers with real-time facilities and reasonable reliability, computers began to be seen near the equipment rooms. Initially they were used only to monitor information, not to perform control. Logs were printed, phase-space diagrams plotted, and emittance diagrams drawn on screens. Eventually the temptation became too great as the simplicity of applying Fourier correction components to the magnet coils directly from the computed values was recognized. Computer control had arrived and the computer, although it was a weak link in the chain, was accepted as being part of the accelerator. With the introduction of integrated circuits and the trend to lower power consumption, reliability has risen to the point where non-computer equipment may be the biggest potential cause of faults.

Control languages

The ability to perform control from the computer was limited in its usefulness by the tedium of writing programs. Unlike a chemical plant or a power station, an accelerator is not subject to a fixed regimen which can be hand-wired or programmed once for all time. Any change in the layout or connection of the control instrumentation gives rise to a requirement to change the computer program. Since an accelerator is usually under development even when it has reached the stage of being useful for experiments in physics, most of the large accelerators currently installed have a team of programmers permanently keeping the software up to date with the hardware.

To minimize this difficulty, interpretive languages can be used. A successful application of this technique was exhibited at the Rutherford Laboratory in England in 1971 in the control of a proton beam-line from the "Nimrod" accelerator[2]. An interpretive computer language differs from a compiled or assembled one in that is is decoded by an interpreter, which obeys each statement of the source text immediately and not as a subsequent process.

Statements can initiate computing operations, or control and acquisition as a subsequent process. Interpreters are easy to write, can be made very powerful and sophisticated for a small expenditure of computer memory, and their one drawback, lack of execution speed, is not relevant in an application where the slowest link in the chain is the reaction time of a human operator.

Central Control Philosophy

Control languages are very flexible, but they can be rather tedious to use in the every-day control of a large machine. The main ring of the National Accelerator Laboratory's proton synchrotron at Batavia, Illinois, was controlled by a system of this kind for the first part of its life, and the lack of conventional control facilities finally became evident. It is easier to push a button than to type in a request. It is easier to turn a knob than to type in a series of requests. An accelerator seems to require conventional controls in the form of a knob, buttons, and switches, as well as the flexibility of a control language. It is this double requirement which has prompted the approach used for the SPS control room. Conventional (and unconventional) control devices are used, but these are not dedicated to particular pieces of accelerator hardware. The links between the hardware controlling an operation and that being controlled are entirely constructed out of interpretive software, and the exact method of control being used can be changed at any time without the need for hardware changes, either to the accelerator or to the equipment on the control desk.

This new approach to the design of a control centre turns out to have certain very important advantages. The fact that the operator is connected to a computer, and not directly to the hardware to be controlled, means that the controls can be modified to suit him physically, the software connection being modified to present an unchanged interface to the accelerator. This in turn means that the same control devices can be used on different occasions by the re-configuration of existing material according to a new philosophy. While the principles of good engineering dictated the use of reliable and tried components wherever possible, new hardware or software has been devised whenever a lack existed.

The following features of the control system are the main ones which are either original, or were developed in the central controls section especially for this project:

1. A centralized graphics system, with pooled resources for all users.
2. The use of a two-level graphics system, giving high definition or colour character information, or both, at a reasonable price.
3. The integration of closed-circuit TV pictures with a graphics system on the same screens.
4. The building of a graphics language into the corpus of an interpretive control language.
5. The use of a "word-store" or macro facility, as a control device in conjunction with an interpretive control language.
6. The development of a purely electronic touch-screen device and its use as a hierarchical selection device or as a set of computer-controlled buttons.
7. The concept and implementation of a computer-controlled multipurpose knob as a control device.
8. The integration of these and a number of more conventional devices into a general-purpose control complex for simulating various kinds of control panels.

Some of these ideas have been partially realized elsewhere[3-10], but the effect of installing and combining them here has been to make a control room whose major features are already being recognized as desirable by designers with similar problems elsewhere.

ACKNOWLEDGEMENTS

This document is intended to describe the steps taken by the author during the process of deciding on the design of a control centre for the CERN SPS accelerator, and the way in which the components, once chosen, were combined to form part of a control system being designed by a team including several other engineers.

Few of the decisions, however, were made without extensive discussion with other people. The author's supervisor, M. Crowley-Milling (leader of the team mentioned above) had rather definite opinions on what was of practical use in a working accelerator. At the beginning of the project he supplied a short paper in which he had written his original ideas, and subsequently followed this up with much valuable advice. The author is grateful to him for keeping his (the author's) feet firmly on the ground, as well as for many useful ideas. A colleague, Dr. J.T. Hyman, supervised the team which designed the computer system and its interfaces. This team included Dr. G. Shering, writer of the interpreter for the NODAL language, and Dr. J. Altaber, writer of the system used with the integrated computer network and supervisor of the contract for the message-transfer system.

The actual work of realizing the equipment required much detailed engineering design and computer programming. This was done by a team of CERN staff members under the supervision of the author, and many original ideas were contributed by each of them in the implementation. The members of the team with their major responsibilities were:

Mr. B. Stumpe, Technical Assistant
(touch-buttons, knob, ball and scan converters)

Mr. S. Summerhill, Technical Assistant
(TV monitors, display devices)

Mr. M. Wenger, Technician
(plasma panels, electronic thumbwheels)

Mr. P. Anderssen, Programmer
(graphics, word store, video switching)

Mr. P. Farley, Programmer
(systems and interfacing, hard-copy)

Mr. R. Koremann, CERN Fellow
(alarm system)

Mr. P. Harris worked with the team as an agency technician throughout the project.

The waveform switching program described in Chapter VI was written by Dr. G. von Holtey.

In addition, many members of the controls group and of the workshops and other services of CERN Lab II are to be thanked for help and advice. There are too many people in this category to mention them individually, but one particular contribution is worth noting: Mr. A. Merlino put a great deal of ingenuity into the mechanical design of the computer-controlled knob.

CHAPTER I

THE ENVIRONMENT OF THE PROBLEM[11]

1. THE COMPUTER USED

Computers were needed for various parts of the SPS control system. One of the basic decisions made was that all computers should be software-compatible, so that one software system could work in all computers and a maximum economy of programming effort be achieved. Preferably all computers should be identical. In any case, large and small computer configurations would be necessary.

Another important decision was to use an interpretive language for the control programs. It was thought that much of the arithmetic parsing process could be taken out of an existing interpreter if there was one for the computer chosen.

Programming for minicomputers can be very tedious if only the real-time monitor is available under which the program will finally run. It is a great help if the computer chosen, particularly in the larger configurations, will support a variety of editors, assemblers, compilers, and a system generation program.

The computer had clearly to be fast, in order to perform some of the rather elaborate calculations needed for computing waveform corrections and other functions. In addition, there is a need for extreme speed to make up for the generally slow inherent speed of an interpreter.

The computers needed were as follows:

i) Six general-purpose computers, with mass storage, for control apparatus in six geographical areas of the accelerator.

ii) A large computer with mass storage to control the power supplies for the magnets.

iii) A computer with mass storage to control the radio-frequency acceleration system.

iv) Five computers to control various aspects of the beam transport process (injection, extraction, and beam-line control for two experimental regions).

v) Two computers for experimental area beam control and one for monitoring the injector accelerator (the CERN PS).

vi) A computer for handling files and archives, a large computer for services of all kinds, including lengthy computation and the handling of magnetic tape, a computer for collecting and displaying alarm messages, and a computer for creating displays.

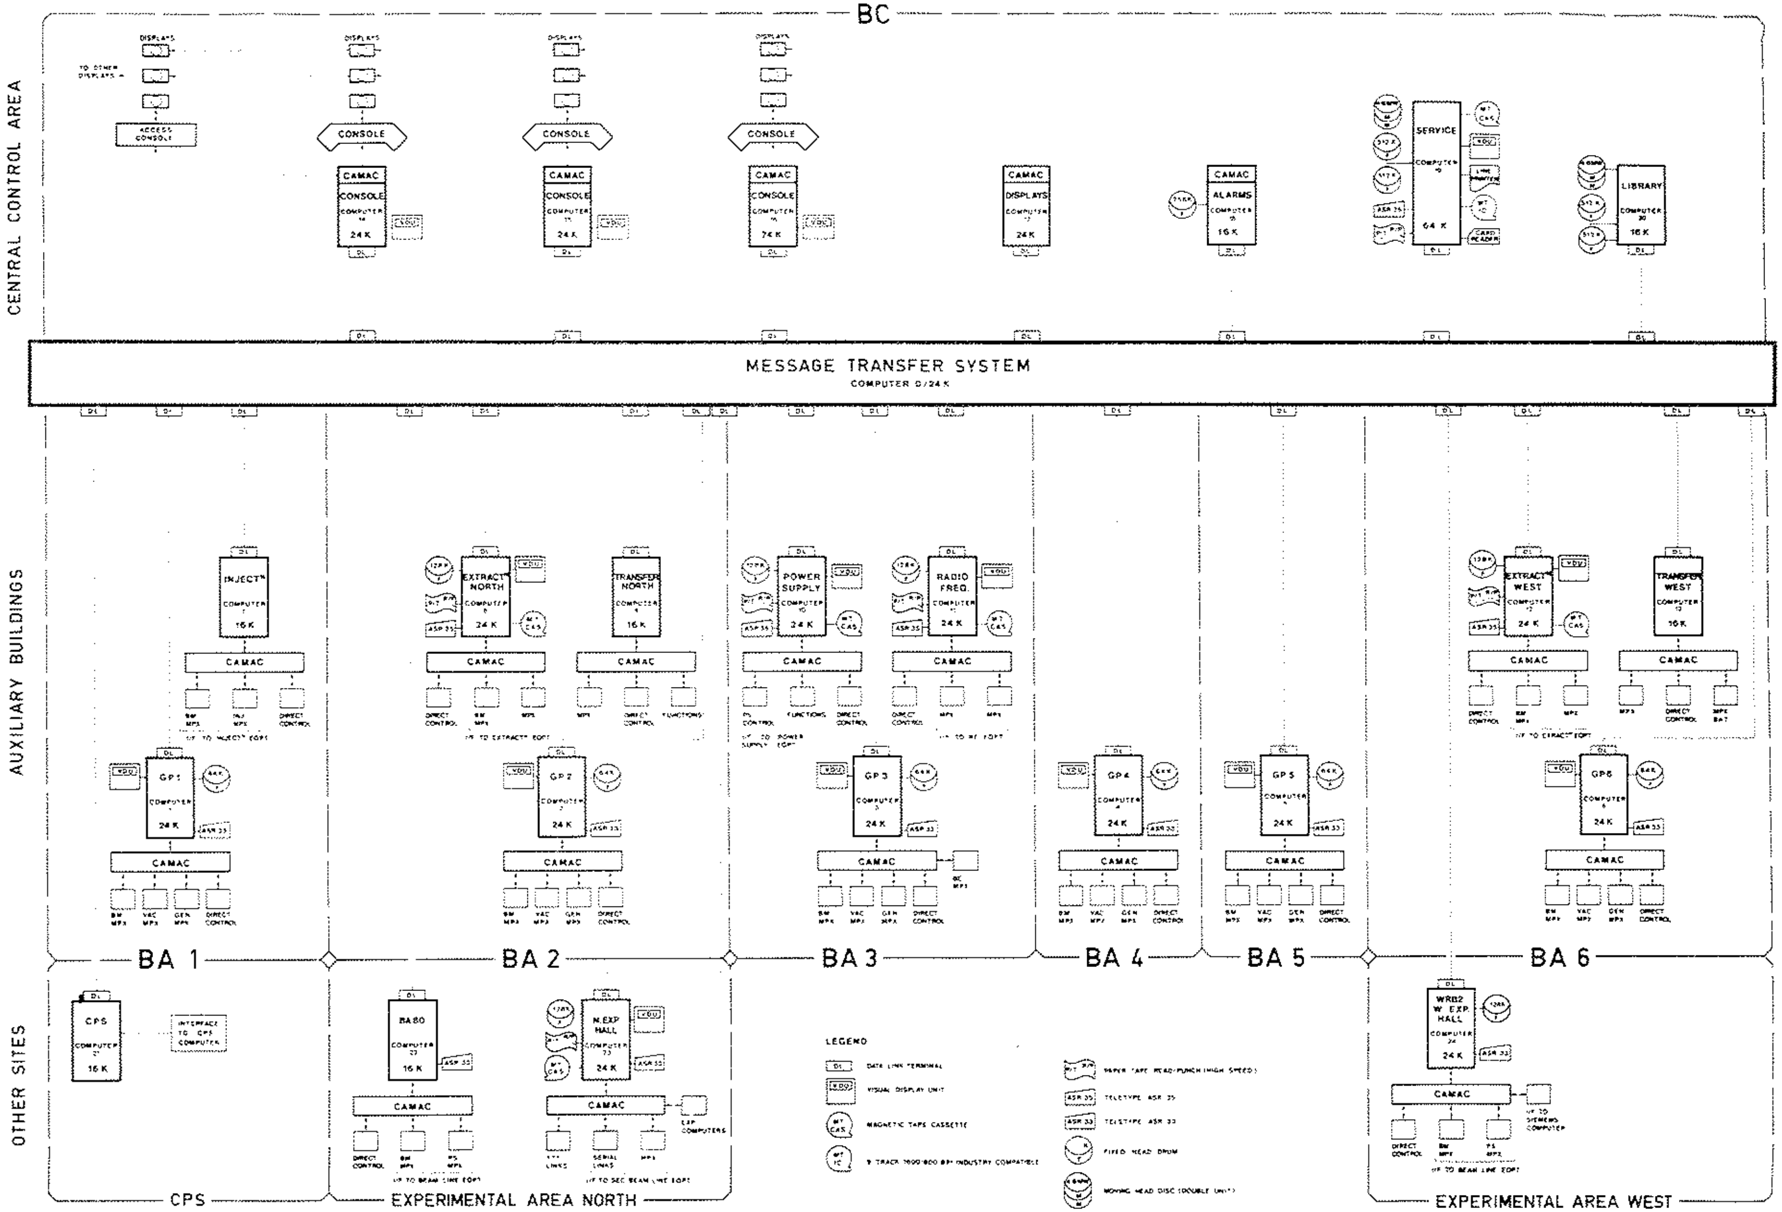

Fig. I.1 Computer system for the SPS

vii) One computer for each console. As will be explained in Chapter II, it was decided to have three of these. A computer dedicated to each console allows an operator to write on-line interpretive language programs of arbitrary power without disturbing, or being disturbed by, other operators.

Twenty-four computers were obtained for these purposes by means of a competitive bidding process which found the cheapest of those computers that satisfied nearly all the requirements. The computer chosen was the NORD-10 (Ref. C1) which is a 16-bit machine having a microprogrammed processor with 300 nsec basic instruction time, a complete set of registers for each of 16 hardware interrupt levels, 256 software interrupts, floating-point hardware in the standard version, a memory capability from 16K words in the smallest version to 256K in the largest, and a very flexible and easily interfaced I/O system involving two data buses, each of them allowing fully interleaved direct memory access from a number of sources at up to 1 million words per second.

Twenty-three of these computers were connected as shown in the diagram (Fig. I.1), and a twenty-fourth was used as a message-switching system, or telephone exchange, at the centre of a radial net.

2. THE MESSAGE-TRANSFER SYSTEM

The fact that the NORD-10 computer has such an excellent I/O system makes it very suitable for use in a store- and forward-message-switching system. The computer simply does the management of messages and storage buffers, and the processor is occupied with the messages themselves merely to the extent of starting autonomous transfers of information in to and out of its memory. The message-transfer system was made as an attachment to a standard NORD-10 computer and programmed under contract by a Paris-based computer house (Ref. C2). This company also made serial data links of high speed and reliability for transmitting the signals between computers over the large distances involved, which in some cases can be several kilometres.

A short message (one buffer of 64 words) can be transmitted by this system from a buffer in one computer to a buffer in another within 5 msec, even when other messages are being transmitted between other pairs of computers.

3. HARDWARE INTERFACES

It was decided quite early on in the project that the CAMAC system -- an international standard -- should be used as the primary interface to the computer. A crate controller was ordered under contract from the computer manufacturer, which enables the external bus of the computer to be brought past each CAMAC crate and to give full computer communication to the controller of each crate.

CAMAC is used directly for all high-speed functions and for specialized functions (such as display device control) for which CAMAC plug-ins are available directly.

For certain requirements, very numerous in the control of our accelerator, CAMAC is too expensive and not sufficiently robust or industrialized. It is also, in general, not sufficiently isolated to be used in equipment rooms with heavy switchgear and large currents. A multiplexer was therefore built, driven from a group of CAMAC plug-ins, and allowing equipment to be interfaced by balanced, isolated lines to robust, cheap plug-ins in remote crates.

Ninety percent of all the equipment is interfaced through this multiplexer, rather than directly to CAMAC.

In the case of the central control room, all interfacing was done either through this multiplexer or directly to CAMAC. No interface plug-ins were designed for the central control system.

4. SOFTWARE AVAILABLE

The computer specification called for a real-time operating system capable of supporting several tasks at different interrupt levels, and such a system was supplied with the computer. This system, called SINTRAN, was modified to allow interfacing with the data-links and the message-transfer system, as well as to improve the scheduling algorithm. In this form the system is called SYNTRON, and it allows the usual monitor functions of handling peripherals, loading programs, scheduling tasks, and communicating with the operator. The data links are effectively high-speed peripherals, and communication with other computers is therefore in all respects akin to writing on a chosen drum or magnetic tape.

An interpreter was available soon after the start of the project. This interpreter[12], written at CERN by G. Shering, was based less on the manufacturer's standard software than was at first hoped. (The lack of a BASIC interpreter was one of the few defects in the definition of the computer. There was a compiler for BASIC.) The following are some basic facilities of the language, which was named NODAL (NOrD Accelerator Language):

i) all standard arithmetic commands;

ii) internal subroutine definition and a large range of control commands, including looping and transfer of control;

iii) all input-output, including CAMAC;

iv) real-time scheduling facilities;

v) data-link control facilities ("perform paragraph 3 in computer 7");

vi) string-handling in the larger computers;

vii) full access to the system control commands at interpretation time;

viii) on-line editing and diagnostic facilities.

In addition, the same language (NODAL), as well as the other system software, is available in fully compatible form on a large time-sharing NORD-10 (later to be used in the network as the service computer). This facility allows many programs, not involving hardware interfaces, to be developed off-line, and the developed programs to be transferred to the actual control system on paper tape. The time-sharing system is also used for system generation.

A final and most important section of the software is the writing of the specialized input/output subroutines controlling individual groups of accelerator hardware. These routines, called DATA MODULES, are called from NODAL and their existence reduces the control of hardware to a simple software function. The data modules are written in most cases by the builders of accelerator subsystems, and they contain the tables needed for hardware control, as well as many checking mechanisms. Thus, switching on a pump or testing the status of a power supply is a simple matter of calling a data module. The fact that the data-base is distributed over data modules in the various computers and not kept in a library except for back-up, means that data-link traffic is kept lower than if all information were kept up to date centrally.

5. THE COMPUTER BUILDING AND THE CONTROL ROOM

The design of this building was not part of the task of building a control centre. Two large rooms (8 × 12 m each) were provided, double-height and air-conditioned. The control room is separated from the computer room by a fairly thin partition which can be cut for mounting displays, etc., but is nevertheless sound-proofed. Visitors can look in, without disturbing the operators, through specially provided windows.

The air-conditioning operates by excess pressure, and the plant is outside the building to avoid noise in the control room. Both rooms have standard suspended computer floors, allowing the concealment of cable ducts.

Illumination was provided to specification. It was decided to have bright lights ($\sim$ 200 lux) for cleaning, installation, and maintenance, and a very low level for working, so that the displays would be clearly visible. It was subsequently discovered that no lighting is low enough to avoid reflections on the console screens, and a complete overhaul of the lighting had to be done after installation. Lighting has to be kept below table level if reflections are to be avoided, and such lighting is now installed.

The rest of the computer building is not air-conditioned. It provides a few offices, a conference room, a kitchen for the accelerator operators, and some workshops and a workroom for visitors.

CHAPTER II

BASIC DESIGN DECISIONS

1. CONTROL DESKS, THEIR NUMBER AND SHAPE

An early conclusion was that if the control desks could be identical, then they would be much more proof against failure than if reliance were placed upon the excellence of the equipment. A number of identical desks have the property that if one of them fails, the user can go to another. Differences in operation for control of different subsystems of the machine can always be obtained by making the consoles behave differently by means of software.

1.1 Computers

Since an operator may want to monitor several functions simultaneously, he must have access to programs at several interrupt levels. For instance, he may need to have a graph plotted every accelerator cycle, and while this graph is repeatedly appearing on one screen, he will want to pursue some other investigation manually on another. He may also have launched a third program to keep a watch on some other parameters in the meanwhile. It is clear that considerable advantage can be obtained if each console has an entire computer at its disposal, and indeed this was allowed for in the original design.

1.2 Consoles

Given the need for adequate back-up and the expense of a computer, interfacing, and a console full of equipment, the number of consoles was severely restricted. The number three was decided on as the minimum which would give adequate back-up in the event of a failure. Throughout the project specious arguments have been used to introduce special equipment on one of the desks, but the attempt to keep them identical has been more or less successful.

The desks themselves are dictated in shape and size by the equipment which is to be fitted into them and the persons who are to operate them. The operator of a console will usually be sitting in front of it, though standing operation should not be precluded.

With the operator sitting down, he should be able to put his knees under the table top and to see over the consoles in order to see and speak to other operators and to view wall-displays on the front wall of the room. At the same time, with only a limited amount of stretching, he should be able to reach all the panels before him.

1.3 Plan view of consoles

Operators often work in twos, and a number of interested persons should be able to stand behind and watch or advise. The advisability of making the electronics fit into internationally standard 19-inch racks makes the maximum number of bays six (five would have been closer to the desired shape from this point of view). Appendix II shows the plan view which is imposed by these considerations.

1.4 Side view of consoles

The side view is also severely restricted. The distance between the knees of a seated person and his eye-level gives the maximum usable console surface. Given a table top, a control surface, and a viewing surface, as in Fig. II.1, the only choice available is that of the depth of each of these three surfaces. Again considerations of what is commercially practicable took over. Television monitors which are available on the market were only obtainable in a high enough quality if they were at least five international units of 44 mm height. Since the plan, as explained later, involved a lot of television monitors, it was thought advisable to make the viewing surface ten units high, and the control surface five units high.

1.5 Dimensions

For the precise values of the remaining parameters (height of table top, angles of surfaces, etc.), a literature search was made for guidance on the precise anthropometric proportions needed. Most of the information available on this subject is published by the U.S. Government in connection with their military program. A short list of references is given[13-17].

The rest of the console design is involved with the mounting of interface equipment with easy access, with ventilation, cabling, and other details. Appendix II gives the actual specification which the author sent to manufacturers in a successful effort to obtain a suitable console at the lowest price.

Figure II.2 shows the installations on the console as they were at some time near the end of 1975.

2. REQUIREMENT FOR SPECIALIZED DEVICES

In the last section we discussed the geometry of the operator interface without in any way specifying what the panels were to be used for. We finished up with six "viewing" panels 44 cm high and six "controlling" panels 22 cm high, disposed in a curve around the operator. This gives us a total of two square metres of interface, and it is clear that only the middle half of this is prime space. Therefore each square centimetre of the surface must be made to count.

Two approaches are possible: ultra-miniaturization, in order to pack as much material in as possible, or a small collection of devices so powerful and flexible that whatever operation the user wants to perform these few devices will do his work for him. An attempt at the first approach is a losing battle. Even if such miniaturization can be done successfully, each new control requirement needs more devices. The two square metres will eventually run out. The second approach is unavoidable.

Thus we are faced with the requirement to mount in our consoles only such devices as can be used for many purposes. The alphabetic keyboard, the television monitor, and the cathode-ray oscilloscope are three such devices. Others were developed and are described in later chapters. Even in the case of television monitors it is not tolerable to dedicate a particular screen to a single camera. Switching and sharing are needed to make each screen go a long way. As will be seen later, this principle was even extended to allow the

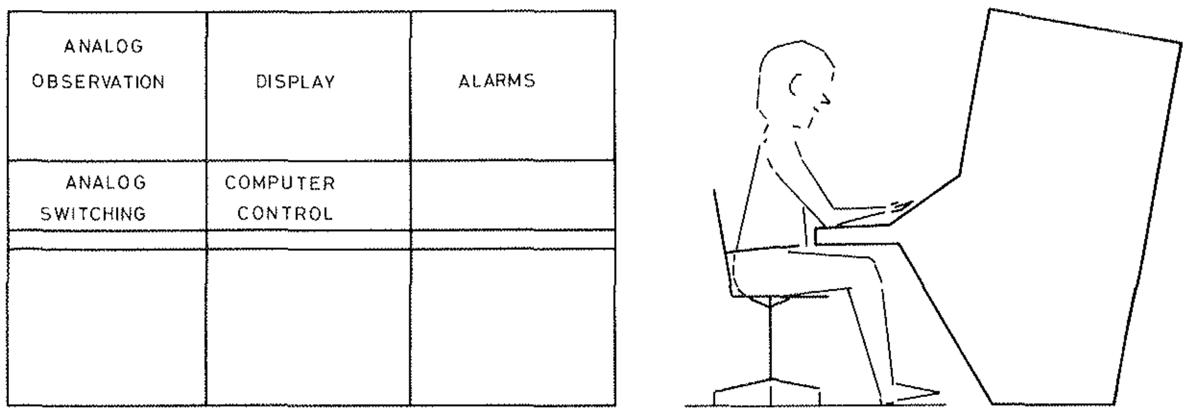

Fig. II.1 The console layout

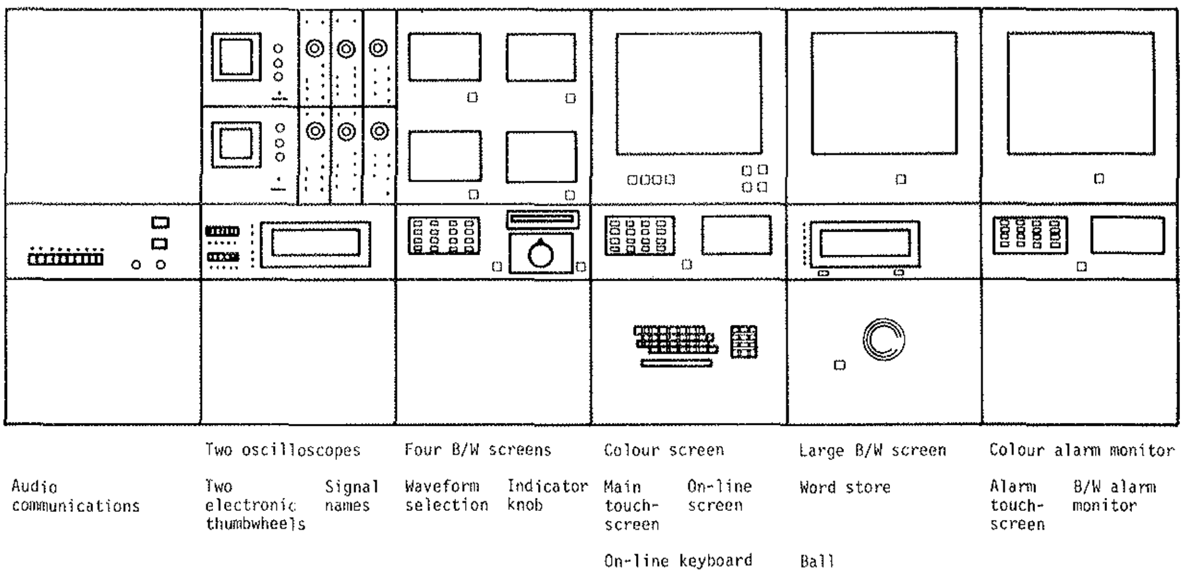

Fig. II.2 The console installations

same screen to show graphics. However, at this stage it will suffice to point out that this requirement for multipurpose devices has excluded many useful and beautiful devices from consideration for use in the control room, because they were not suitable for use in a large variety of applications.

3. PLAN OF THE GRAPHICS SYSTEM

A full graphics system at the present state of the art would enable the operator to draw graphs on a screen in a variety of colours, and move these graphs around the screen at persistence-of-vision speeds by changing some manual control, such as a ball or knob.

To realize such a system, even with one screen, is very expensive. The information flow-rate required to re-draw an elaborate picture 10 times a second is very large, especially when the picture is in three colours. The information transfer rates are such that the graphs would have to be drawn by the console computer itself, precluding the use of a graphics computer and requiring full-scale back-up of the graphics system on each console for security.

It happens that accelerator control engineers have never had such comprehensive graphics facilities, and that they have usually performed their control and monitoring operations very satisfactorily with graphs on a storage oscilloscope. (The most popular storage oscilloscope allows the drawing of graphs on a 1024 × 1024 dot mesh and gives facilities only for a slow full-screen erase, with the additional option of a "write-through" mode allowing a marker to be moved over the screen without storage.)

3.1 Information involved in displays of various kinds

A study of the colour representation problem revealed the little-understood truth that while multicoloured graphs are very expensive to produce, alphabetic information may be represented in colour very easily. Typical facts are summarized below:

Display type	Dimensions	Information (bits)	Additional information for a three-colour picture
Graph display	512 × 512 dots	250,000	500,000
Character display	24 × 64 characters	9,000	2,000

Thus purely from the point of view of information storage, a colour character display costs only 30% more than a black-and-white one, while coloured graphs are a luxury, with an information overhead of 200%. Luckily the most important use of colour in a control application is for pointing out anomalies in tables of figures. Furthermore, character displays can be extended to be of "semi-graphic" type, in which each of the character positions can be made to contain a limited number of shapes in addition to the letters and numbers, and these shapes may be made to produce diagrams and histograms.

Thus a system was arrived at which would enable the drawing of black-and-white graphs on a 512 × 512, or preferably 1024 × 1024 mesh, and also provide a separate facility for coloured diagrams involving some 24 rows of about 64-80 elements. By supplying both types of graphics in the form of a TV signal, both facilities may appear on the same screen simultaneously, if necessary. In addition, a TV signal may be added in the form of a cursor, and this signal is the only one which is continuously driven from the console computer whenever the position of the cursor changes. The character and graph pictures are renewed when necessary, and preferably not more than once per accelerator cycle, from a specially dedicated display computer.

3.2 Components of the display system

We thus arrive at a system based on synchronized TV raster pictures (synchronization is needed for addition of pictures) which are supplied from picture sources consisting of:

Character generators
Graph generators
Cursor generators
Synchronized CCTV pictures.

The pictures can all be switched to a variety of black-and-white or colour monitors, and summing amplifiers are supplied on the inputs of these monitors. The switch is in the form of a rectangular mesh or crossbar, and enables a cursor to be superimposed on a graph, or a character list to be superimposed on a TV picture, or even allows a graph picture to be switched to a colour input of a colour screen, giving the very expensive colour facility mentioned above if it is needed. Of course three full graph facilities must be allocated temporarily, but it is possible.

3.3 Advantages gained by adopting a switched system

Advantages of this crossbar-switching system are that the picture-producing facilities can be used by anyone in the system with suitable TV monitors. Pictures can be viewed on all or any desks, exported from the control room, or, as mentioned later, diverted to a printer. The advantages of the TV raster form for pictures are many: each picture-generator has the responsibility of maintaining the refresh of its own picture at 50 complete frames per second, thus relieving the driving computer of this responsibility. Monitors for TV are cheap and simple, and the advantage of working in a field where advances are most likely is that the system will be able to make use of such future improvements as flat screens and flyback-time information systems, TV signal processors and other devices, when or if these become available. Already several devices have been added to the system which were not known when the system was first conceived. The TV norm which is used is the standard European one: 625 lines, 50 interlaced frames per second. Figures II.3 and II.4 show respectively the system layout and the actual facilities provided on the video-switch. The switch itself is formed of modular plug-ins, containing mechanical relays, and driven from the display computer by software which will be described later on.

3.4 Disadvantages suffered because of the adoption of the switched system

The disadvantages of using a TV norm have become evident with the progress of time. The resolution limit imposed by 625 horizontal lines is often a little coarse. The speed at which information can change along a scan-line is as high as 6 MHz, making the electronic handling of TV pictures rather tricky (though not usually impossible). The high switching

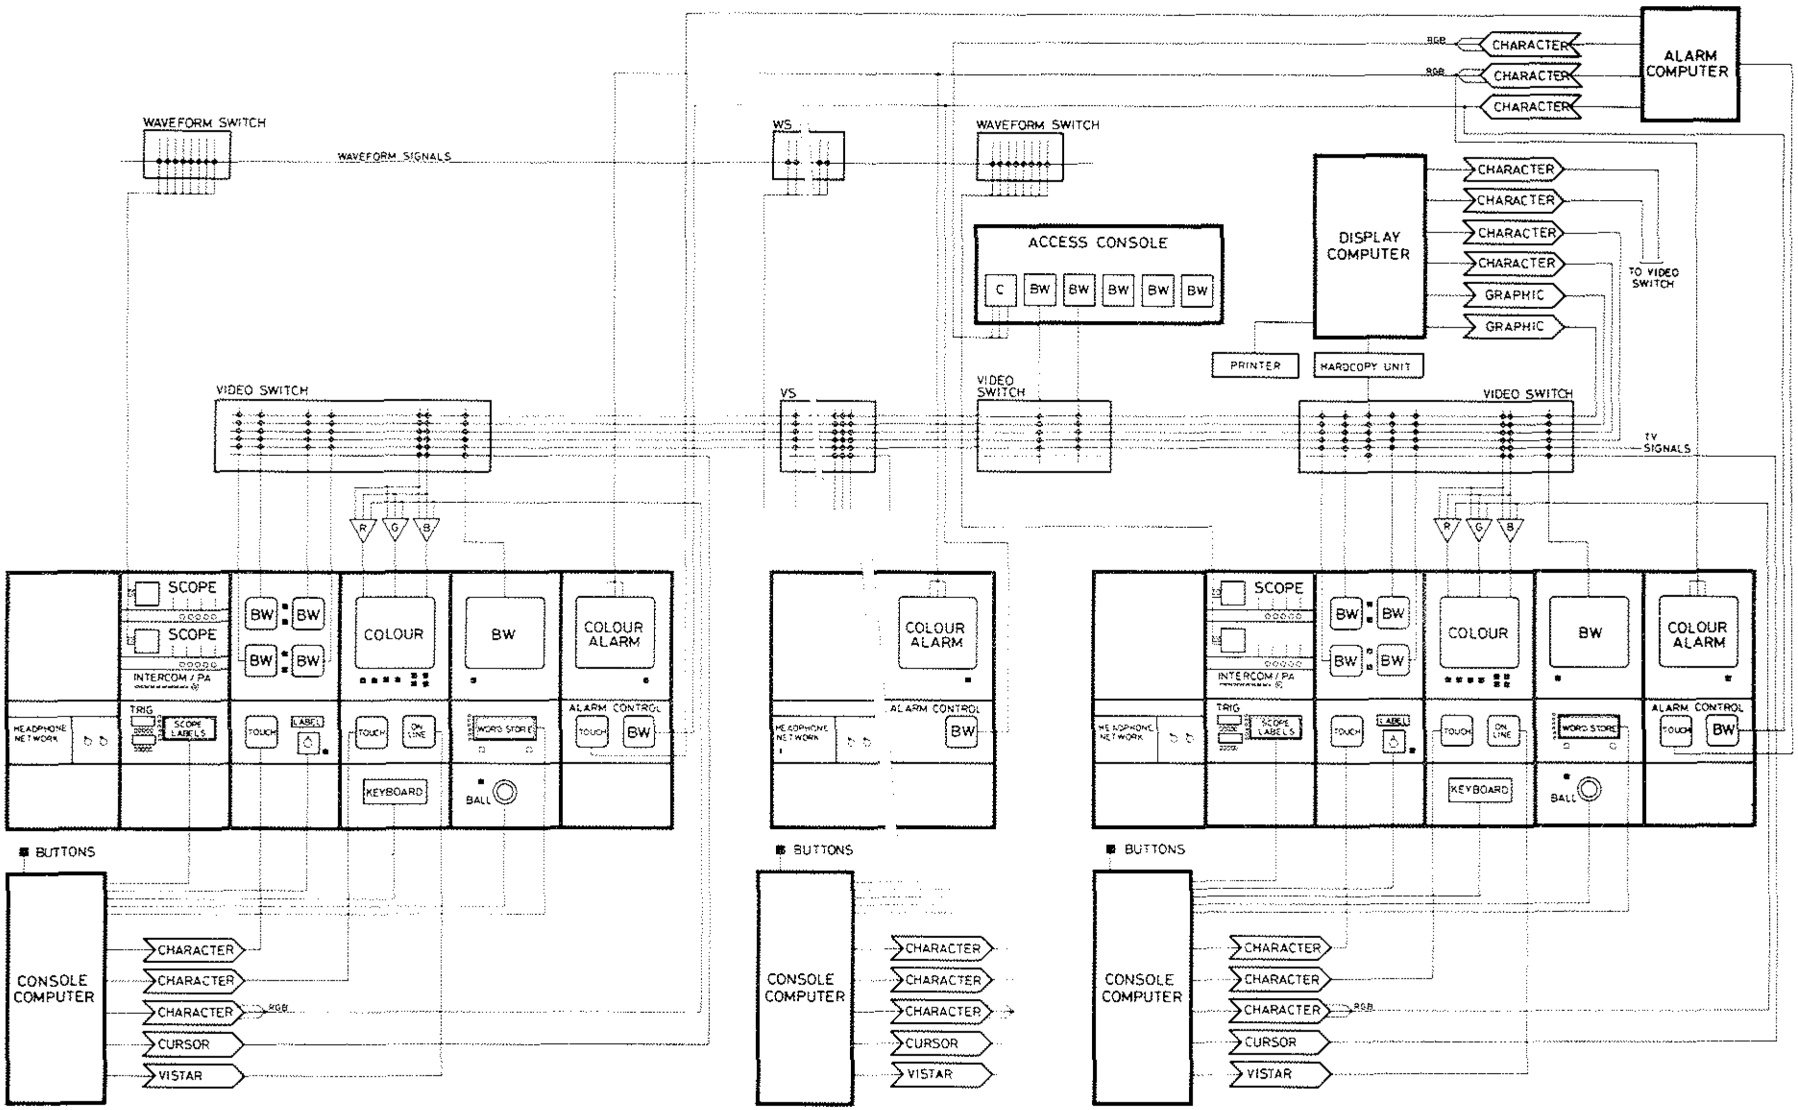

Fig. II.3 The central display system

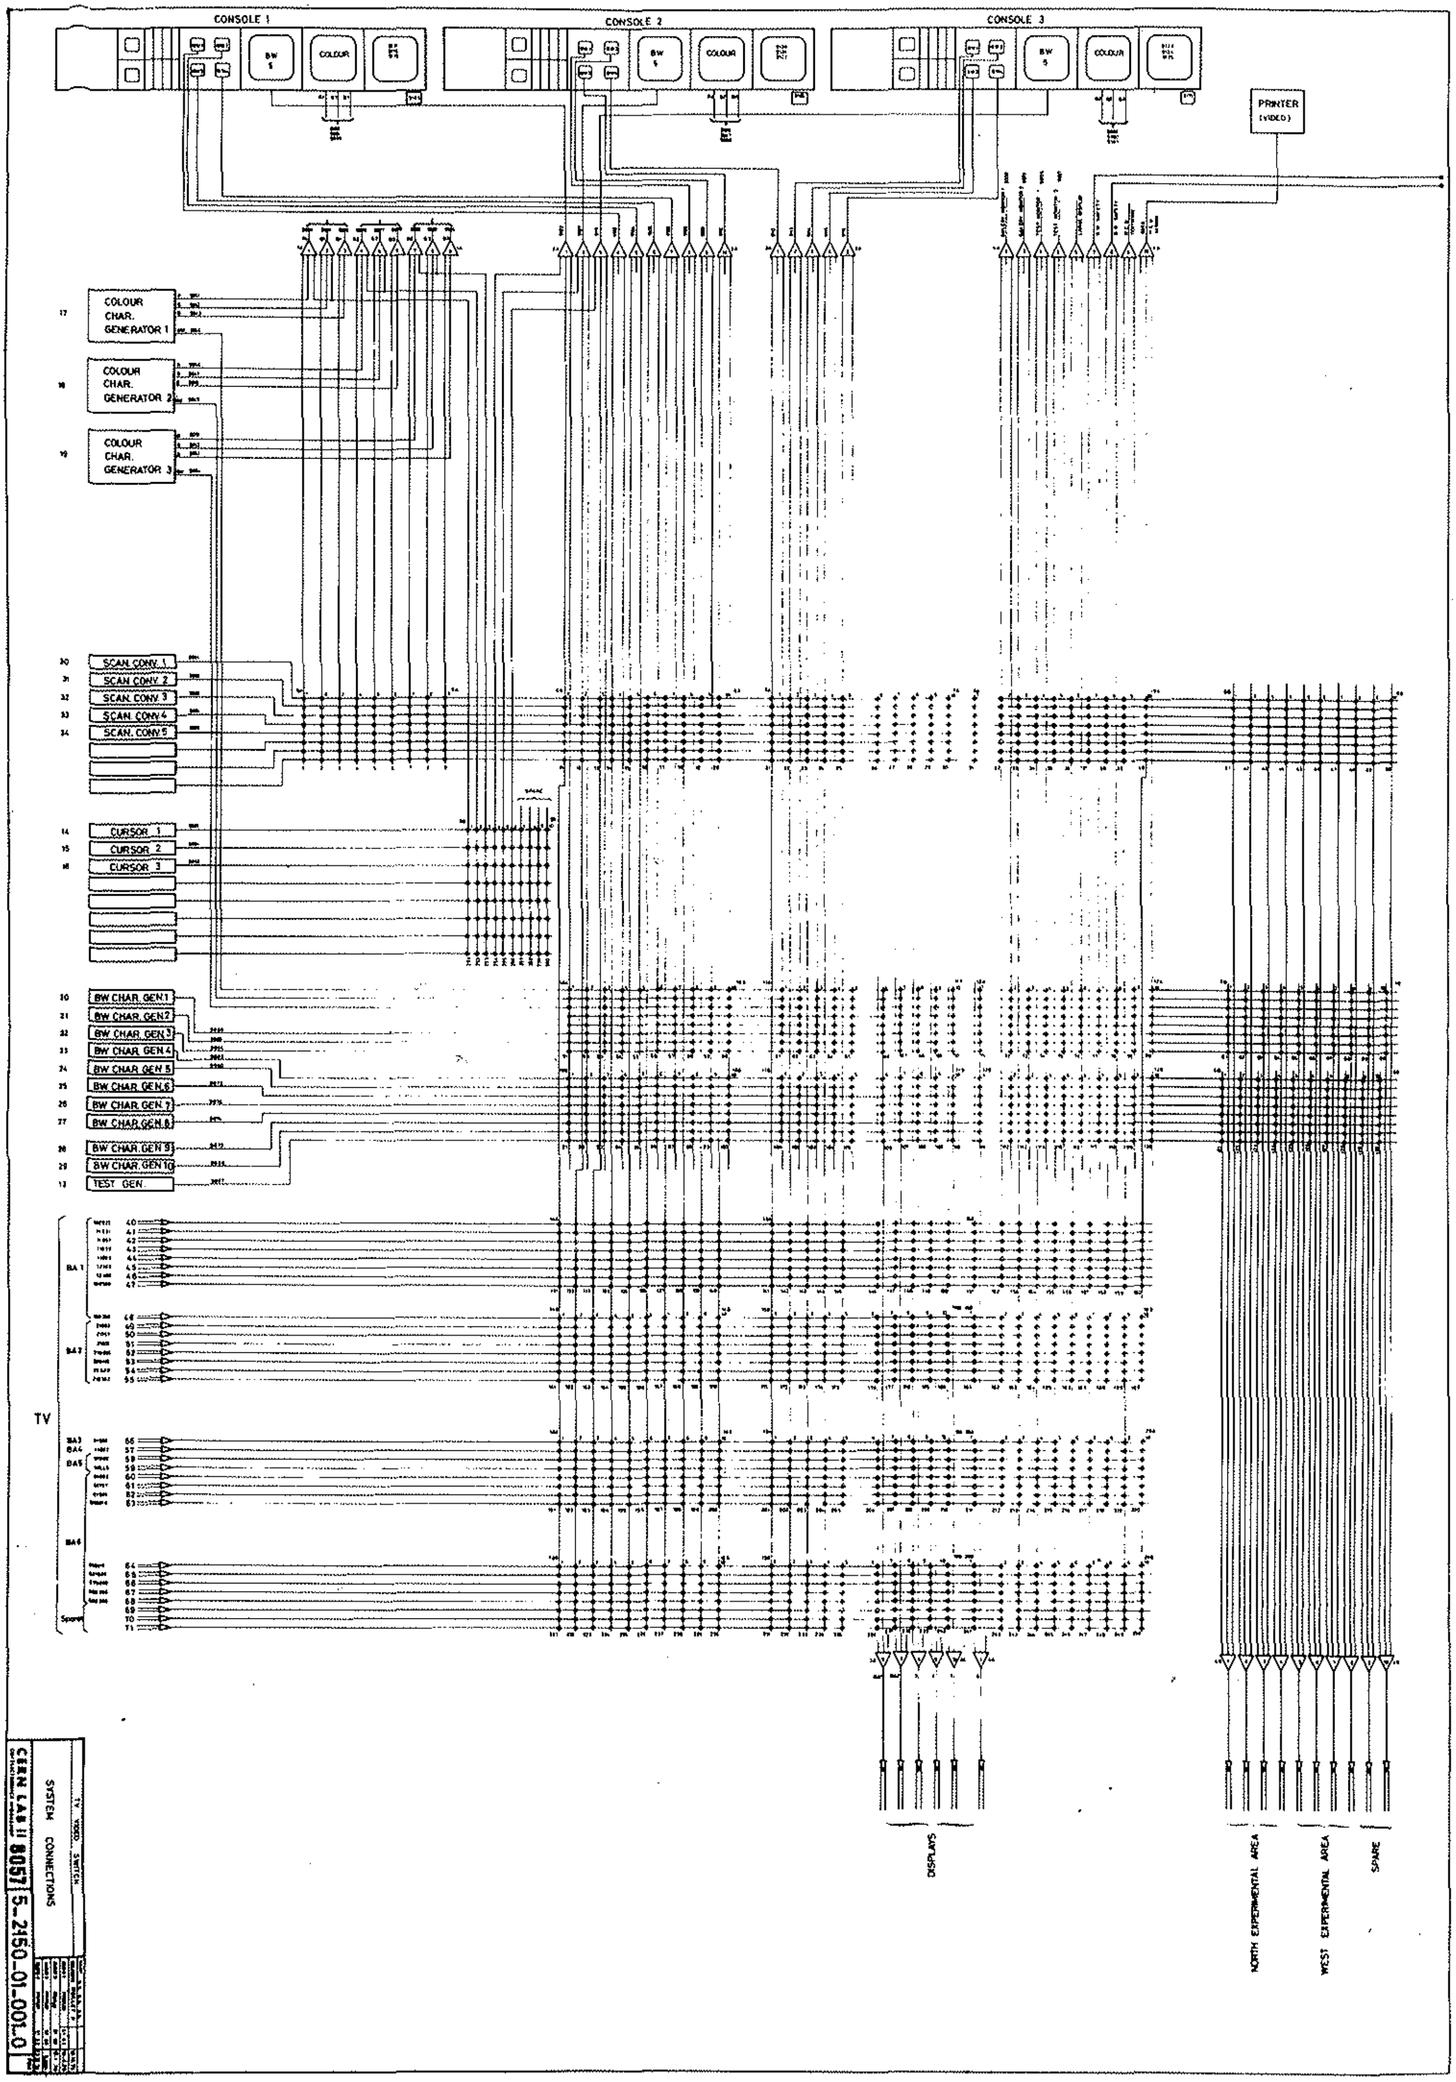
CONSOLE 1
CONSOLE 2
CONSOLE 3
BW
COLOUR
PRINTER (VIDEO)
COLOUR CHAR. GENERATOR 1
COLOUR CHAR. GENERATOR 2
COLOUR CHAR. GENERATOR 3
SCAN. CONV. 1
SCAN. CONV. 2
SCAN. CONV. 3
SCAN. CONV. 4
SCAN. CONV. 5
CURSOR 1
CURSOR 2
CURSOR 3
BW CHAR. GEN. 1
BW CHAR. GEN. 2
BW CHAR. GEN. 3
BW CHAR. GEN. 4
BW CHAR. GEN. 5
BW CHAR. GEN. 6
BW CHAR. GEN. 7
BW CHAR. GEN. 8
BW CHAR. GEN. 9
BW CHAR. GEN. 10
TEST GEN.
BA 1
BA 2
BA 3
BA 4
BA 5
BA 6
Spare
TV
DISPLAYS
NORTH EXPERIMENTAL AREA
WEST EXPERIMENTAL AREA
SPARE
TV VIDEO SWITCH
SYSTEM CONNECTIONS
CERN LAB II 8057 5_2150-01-001_0

Fig. II.4a

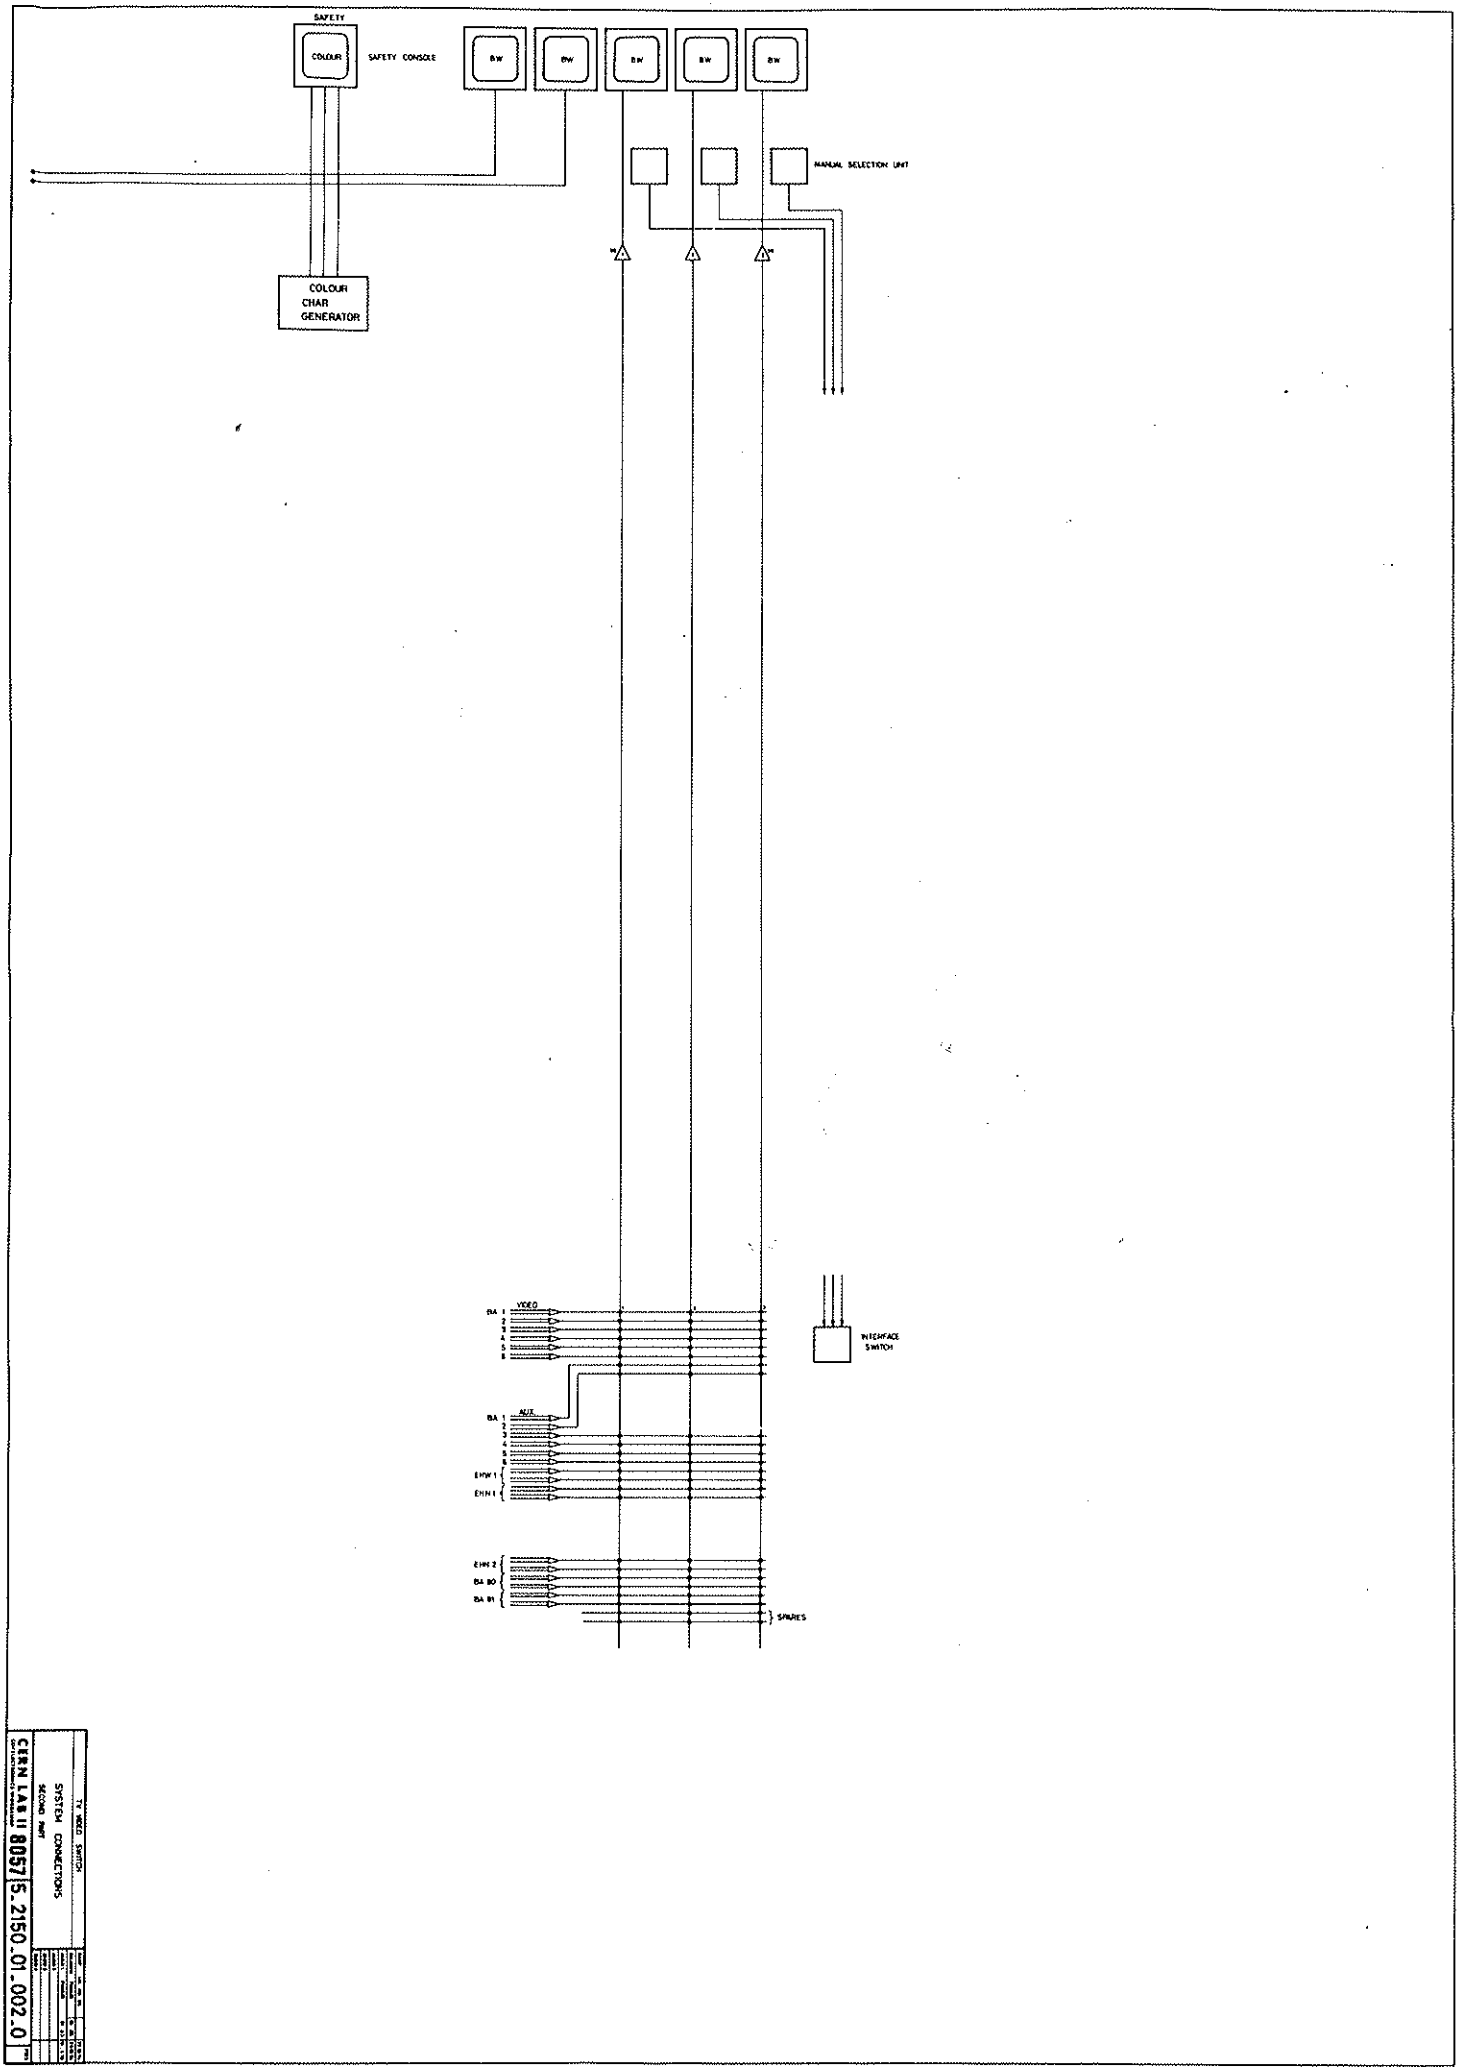
SAFETY
COLOUR
SAFETY CONSOLE
BW
BW
BW
BW
BW
MANUAL SELECTION UNIT
COLOUR
CHAR
GENERATOR
VIDEO
BA 1
2
3
4
5
6
INTERFACE
SWITCH
AUX
BA 1
2
3
4
5
6
EHW 1
EHN 1
EHN 2
BA 80
BA 81
SPARES
TV VIDEO SWITCH
SYSTEM CONNECTIONS
SECOND PART
CERN LAB II 8057 5.2150.01.002.0

Fig. II.4b

speed makes it impossible to use a number of very attractive display systems which would otherwise be available, and the 50-frame per second picture speed, while moderately free of flicker, cannot be slowed down further, as would be necessary for some colour schemes involving sequential presentation of different colour fields. Finally, there are a number of very nice systems for graphics presentation which simply cannot be adapted to TV raster form. Such systems cannot be integrated into our crossbar. On balance, however, the advantages of standardization on a 625-line 50-frame per second TV picture are so great that they are worth making sacrifices for. Special oscilloscopes, hard-copy devices, microfiche readers, and many other devices can be bought straight off the shelf and installed without alteration. All pictures can be synchronized and added without conversion, and the prices of new components are kept low by a very competitive market.

4. SOFTWARE FOR GRAPHICS

As has been mentioned earlier, the accelerator control is done through a language called NODAL. Facilities are available in this language to send both programs and data to other computers, and to ask another computer to perform a program and remit its results to the originating machine. It is quite clear from the start that these facilities will be needed for the production of graphics, especially when those graphics are to be requested at the console computer and actually produced at the display computer.

4.1 Display files

One way of building displays is for the interpreter in the console computer to be used to create a display file, consisting of codes to be obeyed by the display computer. Such codes would consist of shorthand instructions for the drawing of vectors, dots, and character strings, and also for usual combinations of these. Textbooks on computer graphics abound with such codes.

4.2 Use of the interpreter

Another approach is made possible if the interpreter is also present in the display computer. This is to add simple graphics statements to the language. Thus the operation "draw vector", "draw circle", and so on become elements of the language, and use may be made of the fact that the language already contains most of the facilities needed in a display language, such as looping, conditional flow, and arithmetic and trigonometric operations. The mechanism by which this language was implemented is discussed in detail in Chapter IV. It will suffice to mention here that by avoiding a special graphics language, and making the graphics part of the control language, much flexibility is gained. Programs for drawing pictures can be originated anywhere in the system, even if most of the display facilities are only available on the display computer.

The same thing is true of the video-switch. While the picture-switching hardware is all on the display computer, switching operations can be initiated from any computer in the system.

CHAPTER III

THE GRAPHICS SYSTEM (HARDWARE)

As discussed in the previous chapter, the problems of drawing graphics and those of tabulating alphanumeric data are distinct and in general incompatible. Alphanumeric data contain a reasonably manageable amount of information and can be represented on a screen with high quality and even in colour, as will be shown later in this chapter. Many devices are on the market which try to apply the simplicity of such displays by reducing the complexity of the picture to be drawn. Simple electrical diagrams or logical drawings of pipework can often be made out of elements which form an extension to a character set. Thus the eleven supplementary characters:

─ │ ┼ ┘ └ ┌ ┐ ┴ ├ ┬ ┤

could suffice to draw simple graphs of this kind:

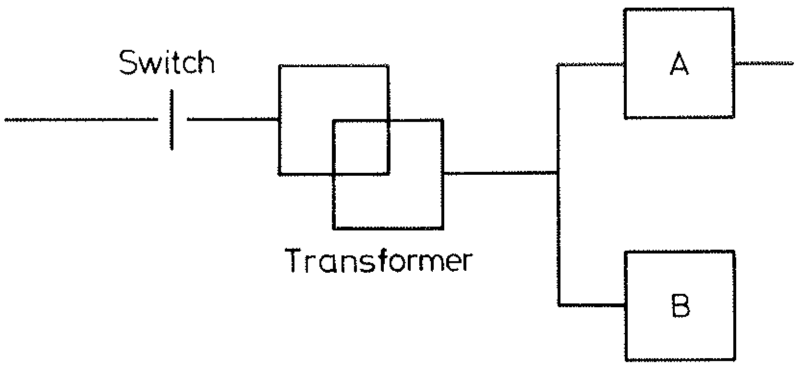

and for some applications such graphics are sufficient. The addition of a character comprising a solid rectangle enables histograms to be drawn also.

Another way to keep the information content of a graphical display low is to represent elements in memory by their parameters. Thus lines may be represented by the coordinates of their end-points, and circles by their centres and radii. Again, such a representation makes certain simple graphs like

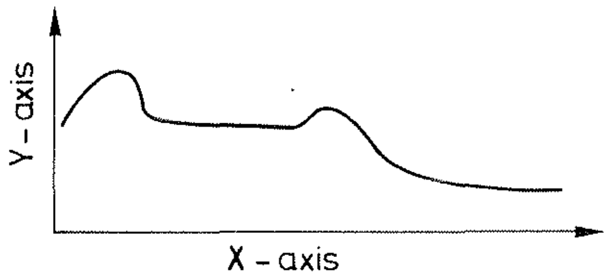

very easy, provided the curve is not required to be so smooth that the information content gets out of hand. The requirement for a family of curves means that such representations become impracticable, and the representation of the picture by holding the black-or-white status of every identifiable dot becomes mandatory.

All in all, then, it was decided to include a character display with limited graphical capability and to use it whenever possible, and to allow genuine, generalized graph-plotting by the use of specialized hardware whenever this should become necessary.

1. CHARACTER GENERATORS AND COLOUR DISPLAYS

Character generators which produce a black-and-white television picture from an input stream in ASCII or some other character representation are now commonly available commercially. Since such generators are at the base of much of the work which we have done, a typical unit of this kind will be described in some detail.

1.1 The memory

The centre of the unit is a memory, which is randomly addressable and which consists of some 64 × 24 = 1536 positions, each having sufficient binary digits to define the character to appear at a particular position on the display screen. We thus have a relationship between what is stored in the memory and what appears on the screen:

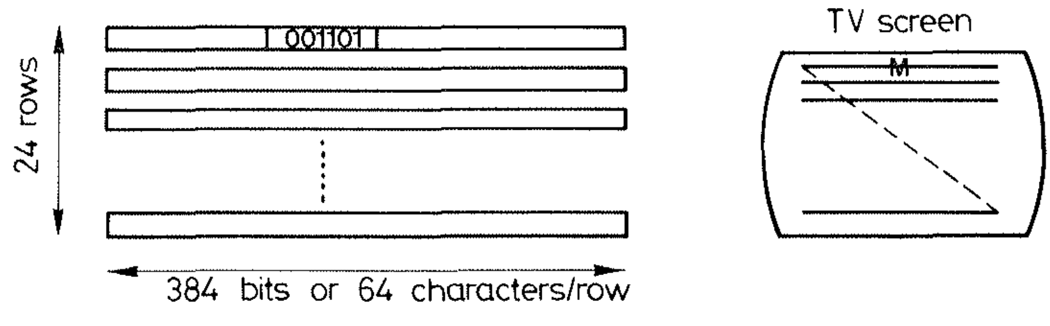

The input to the device, which is typically a teletype interface, delivers groups of binary characters to the memory for storage. Two registers hold the X- and Y-positions of the memory into which the next character will be written.

In the normal mode, whenever a character is written into memory, the X-register is incremented by one, with carry to the Y-register when appropriate, so that characters are written from left to right in sequence, a new line beginning after the end of an old one. There is also a continuation from the bottom right to the top left of the memory page.

1.2 Control characters

In addition to the sequential mode of writing, most such devices reserve special "characters" from the input stream for command operations. These characters are not written into the memory and do not step on the place at which the next character will go into the memory. Instead they cause various operations such as a reloading of the X- or Y-registers, a zeroing or clearing of the memory, or an incrementation or decrementation of the X- or Y-registers by unity, so as to cause subsequent writing to begin elsewhere.

In this way, an exact image of what is to appear on the screen may be built up by sending suitable ASCII and non-standard bit combinations to the input interface.

1.3 Producing the output

To produce a television picture from the contents of this memory, four principal digital clocks are used:

i) One clock gives the page repetition frequency of 20 msec, or 50 frames per second.

ii) One clock gives the line repetition frequency of 20/312.5 msec, or 312.5 lines per frame. This corresponds to a frequency of 15.625 kHz, re-synchronized each time the first clock signals a new page.

iii) One clock gives the character repetition frequency of 64 per line or 1 MHz.

iv) One clock gives the dot repetition frequency of about six per character, and therefore runs at 6 MHz.

A television picture is produced by a line sweeping from left to right every 64 μsec, and producing a complete picture every $\frac{1}{50}$ second. A process called interlacing drops each second frame by half the distance between the horizontal scan-lines. In real television the even interlace frame can contain different information from that in the odd interlace frame. In character systems of this kind the two successive frames contain the same information, but the interlace system, by presenting the picture twice as often, reduces the visual flicker apparent to the observer. In addition, the retention of the interlace system enables genuine TV pictures to be added to graphics pictures of this kind, provided synchronization is attended to.

Characters are usually generated on a 5 × 7 matrix with one space to the right and three spaces below each character. The example below shows the letter "A":

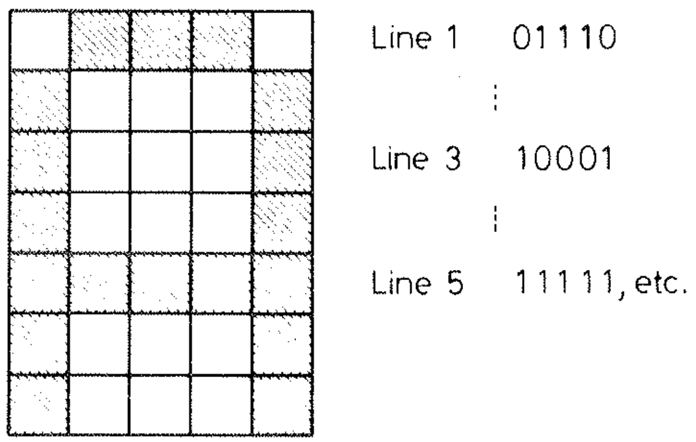

In each interlace frame, the code representing this letter emerges from the memory ten times, once per horizontal scan-line (including three empty lines between the rows), and remains there for one microsecond each time. The line-count, the letter code, and counter of the 6 MHz clock are delivered to a read-only memory, which yields for this combination of input a five-digit output representing the points to be illuminated on the screen. Thus, for example, when drawing the third line of the letter "A" above, the memory would be presented with the code for A (1), the line number (3) and the next six ticks of the 6 MHz clocks will yield the digits 1, 0, 0, 0, 1, 0. (The last zero is for the intercharacter space.)

The actual character generator included in the control system used additional binary digits in the memory to specify the colour channel to which the output is to be diverted for each character. Special characters were also specified giving lower case letters such as a b c, as well as upper case letters A B C, and also some characters were required allowing horizontal and vertical lines and boxes to be drawn. Such "drawing" characters must cross into the margin positions and occupy the entire box, so that continuous lines may be drawn across the screen without gaps.

Appendix II contains the specification which was sent to manufacturers in an effort to obtain controllers of this kind. Twenty of these generators were delivered (Ref. C3), six allowing output in colour, but otherwise identical in all respects. Later chapters describe the use of these generators, showing that they work satisfactorily.

2. DOT GENERATORS AND CURSOR GENERATORS

The ideal graphics generator for generalized graphics in monochrome is a semiconductor memory with 500 rows of 500 to 1,000 binary digits, each addressable individually by a mechanism which allows a one or a zero to be written at will into any randomly addressable position. A separate, fast read-out mechanism, with multiplexing to attain the necessary speed, produces a television picture with a bright-up for each digit in the memory, 1,000 dots per line, 500 visible lines per picture:

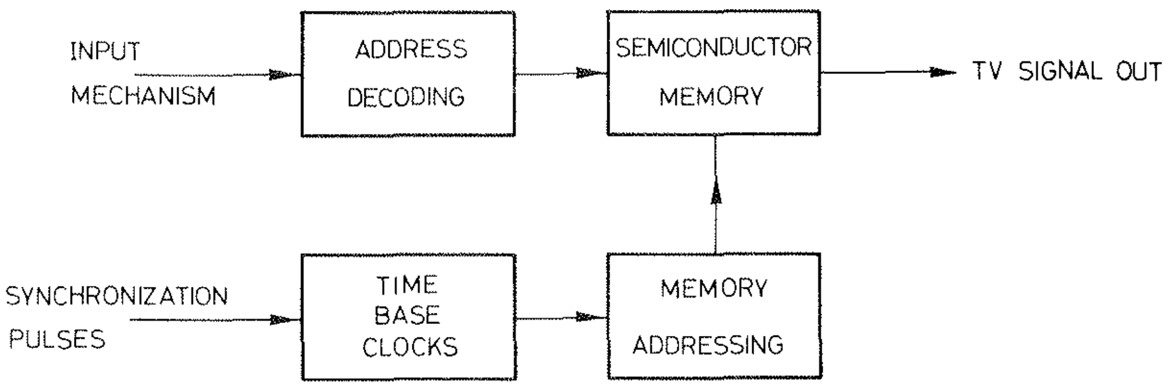

The main defect of such a system is its cost. There are in fact several such arrangements working, and at least one has been manufactured for sale[18-20]. In fact the system referred to has a lower resolution than that specified above (256 × 384) and this lower resolution would not be adequate for fine curves, but its existence shows the technical feasibility of a totally digital solution. The requirement to store 500,000 dots for a single picture, and the fact that more than one such picture may require to be displayed concurrently, made the installation of digital memories of this kind quite impractical.

2.1 The data disk

Another way to store the necessary half-million dots per picture is by means of a magnetic disk. An American company (Ref. C4) makes a complete package for sale, including a driver computer, interfaces, and software.

A number of interesting technical problems have been solved by this company. The speed of acquisition of data from the disk is about 500 kHz, which is a very respectable rate for an input/output device, but nevertheless this is only one tenth of what is required for drawing a TV picture in real time. The extra factor is won by reading from several tracks of the disk in parallel, and copying this information into a register, from which it is shifted out serially, giving the necessary increase in speed. The information must be appropriately distributed on the tracks of the disk, and even such a simple operation as drawing a straight line between arbitrary points can require a rather elaborate data distribution process. It can readily be seen why a computer and special software are required.

In spite of this disadvantage, however, the data disk might have been considered: after all, the software does exist, and the complicated problems of random access to a strictly organized structure have been solved. The thing which makes the data disk inacceptable in our system is the fact that the TV signals from it are necessarily in synchronism with the disk. This conflicts with our requirement to add picture signals, unless all other picture sources are synchronized with the disk. Even then, the use of a second disk is excluded (a disk holds 20 black-and-white, or 6 coloured pictures, or a combination), on the grounds that the two must run at identical speed. Lack of expansibility is a restriction we cannot live with.

Synchronization of a disk is possible; that is, the angular velocity ω can be constrained and the disk can be made to pass a certain point at certain intervals equally spaced with time. Since such a system is stable, it is likely to oscillate, and it is very difficult to prevent systematic excursions in the time derivative $\dot{\omega}$. Such excursions, of course, would cause jitter or at least distortion of the pictures. The unwillingness of the manufacturer to supply circuitry for external synchronization of the data disk confirmed our suspicion that the problem was an inadvisable one to tackle. The data disk was abandoned as being an attractive device but one unsuited to a system of the kind planned.

2.2 Scan-converter tubes

Luckily, a second solution is available. This is the storage or scan-converter tube. The storage tube is a cathode-ray tube of conventional design, except that instead of a phosphor on the front glass, a target electrode is placed in the line of the accelerated beam (shown in Fig. III.1a). The picture is stored as electronic charges by secondary emission on a matrix of silicon dioxide squares deposited by integrated circuit techniques on a silicon substrate (Fig. III.1b). Each square carries a charge. Read-out is performed by changing the target voltage so that secondary emission no longer occurs, and the existing charge-pattern modulates the scanning electron beam by repelling the electrons in it at places having high negative charge. Read-out is slightly destructive (pictures can be made to last as long as 20 minutes), but erasure can be performed by yet a third voltage on the target which corresponds to a fully destructive read-out.

Tubes of this kind are now available from a number of companies in Europe and the United States. Since the target contains squares of silicon dioxide about 10 microns on a

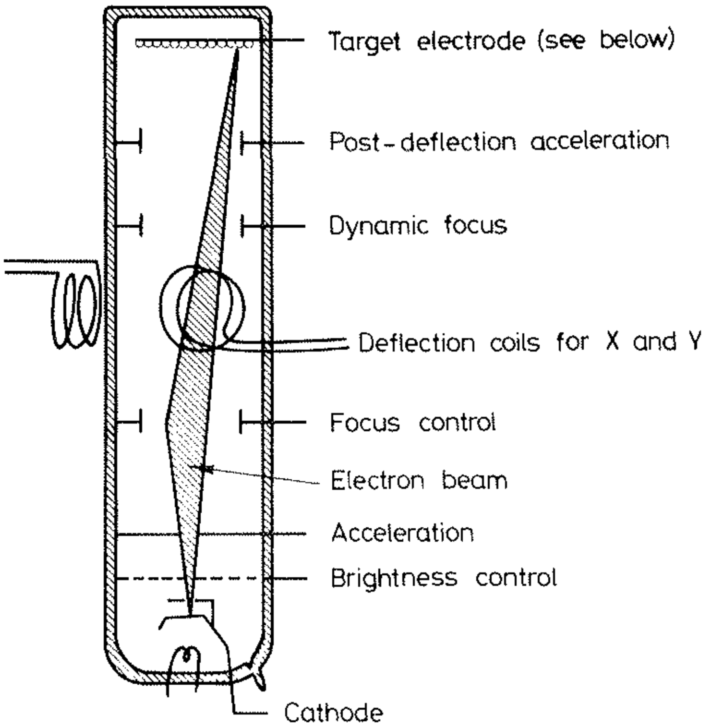

a) Scan converter tube

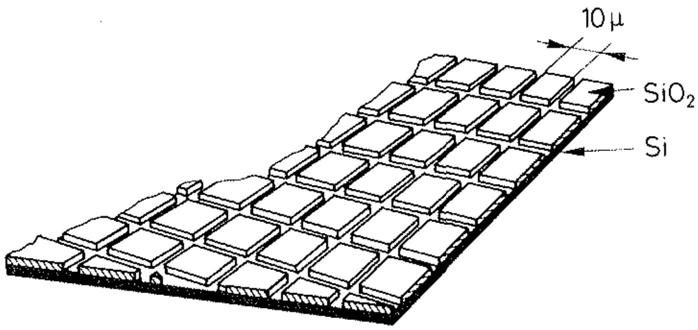

b) Scan converter-tube target

Fig. III.1

side, resolutions are not limited by the granularity of the target, but only by the diameter and shape of the spot. By careful design of the focusing circuitry, resolutions of 500 line-pairs and more are obtainable. Curiously enough no European manufacturer at present markets a complete electronic package for deflecting and modulating the existing tubes. However, two American manufacturers do so, and sell them under the name of "Scan Converter Memory". Samples of both were purchased and tested. Both produced satisfactory television graphics, and the cheaper of the two was chosen for the system (Ref. C5).

2.3 Using the scan-converter

The writing of dots into the scan-converter is done by moving the spot in the X and Y direction with digital-to-analog converters, and producing a bright-up of the beam when the spot arrives in the chosen position. The speed at which dots are written is limited by the settling-time of the digital-to-analog converters for adjacent dots, but for widely separated dots the settling-time of the deflection coils may be a more severe restriction. Half a million dots per second is quite practical, but it is possible to achieve such speeds only if the position of the dots is calculated electronically, as by a vector generator.

2.4 Flyback writing

Since the switch from reading to writing mode requires high voltages to be changed and takes time, it is not possible to view a picture which is being written. That is, the scan-converter cannot produce a TV picture whilst it is being written on. Since an elaborate picture can take an appreciable time to write (say one second) the operator is forced to view a blank screen during this time. This was considered undesirable, and luckily it is avoidable. A television picture is drawn every 20 msec, and the spot takes 19 of these milliseconds to travel on its zigzag path down the tube, and 1 msec to return:

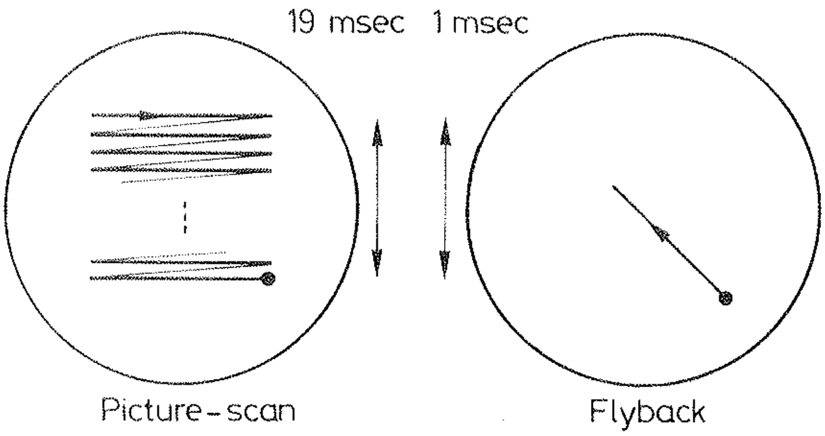

During this millisecond the electron beam is not needed for drawing the picture, and is used instead to write into the target. In practice this millisecond is adequate

- to switch from "read" to "write";
- to write about 300 dots;
- to switch back from "write" to "read";
- to re-position the spot for the reading of a new TV frame.

A vector which crosses the entire picture can be written in three frame times, or 60 msec. Quite elaborate pictures can be drawn in a few seconds, and the operator can see the picture "growing" before his eyes, just as if the writing were being done continuously. It has been found that a drawing time of even several seconds is quite allowable under these circumstances, whereas with a disappearing picture the operator feels lost after as little as half a second in front of a blank screen.

Dot generators with suitable digital-to-analog converters are available commercially, as are vector generators which produce dot coordinates given the X and Y incremental addresses in binary digits. Both are made in CAMAC-compatible form by several manufacturers, and there was no need to build this equipment, though the dot-generator was modified to stop at the end of each flyback period as described in the previous paragraph.

2.5 Interfacing the scan-converter to the computer

The scan-converter memory was interfaced to a CAMAC register, so that mode-switching could be commanded from the computer. The interface provides the necessary commands for "read", "write", "erase", and several other modes, and also allows for manual switching for test purposes. It contains special circuitry for the protection of the rather delicate scan-converter tube from over-current damage by a faulty voltage configuration on its electrodes. Such faulty voltages could result either from an illegal combination of manual commands, or from a voltage transient during switch-on or switch-off of the mains supply.

2.6 Requirement for a cursor

An important requirement in any interactive system is for the operator to be able to point out to the computer some item or position on a screen. In the past, this has often been done with a light-pen. It was generally agreed, right from the start of the project, that while it might be suitable for experiments in computer graphics, a light-pen was unthinkable for the serious control of a large machine. The trailing cable, the exhausting position of the hand, the special constraints on the direction in which it must be pointed at the screen, all conspire to make a light-pen a rather inconvenient instrument. A cursor on the screen itself is more convenient.

The most obvious way to produce a cursor is to draw it on the screen by the same mechanism as the graph or table with which it is to be used.

Unfortunately, having drawn the cursor at some position on the screen, we wish to steer it around by computer or by a manual control. On a screen of the storage type (of which the scan-converter is an example) moving a cursor is effectively impossible. This is because the cursor covers some picture elements, which must be erased in order to re-write it in a new position. The part of the picture previously covered by the cursor must be re-written, and the information for doing this is no longer available.

The only way to produce a cursor for a storage display of this kind is therefore to produce a separate television signal with a picture of the cursor line or dot on it, and add this signal to that carrying the graph or table, before displaying the sum-signal on the screen.

In this way we get the concept of a cursor as a separate picture-generator, producing a separate, synchronized TV picture just like the character generators and scan-converters, and able to be added to any screen which requires this facility by a summing amplifier at the input of the monitor.

The input to the cursor is a pair of computer-interfaced registers, which hold the X and Y coordinates of the point at which the cursor is to be drawn. The X-register is copied every line and the Y-register every frame into counters, which are counted down at line and dot frequencies, respectively. When both counters are zero simultaneously, we have arrived at the indicated part of the picture. The following sketch shows the principle:

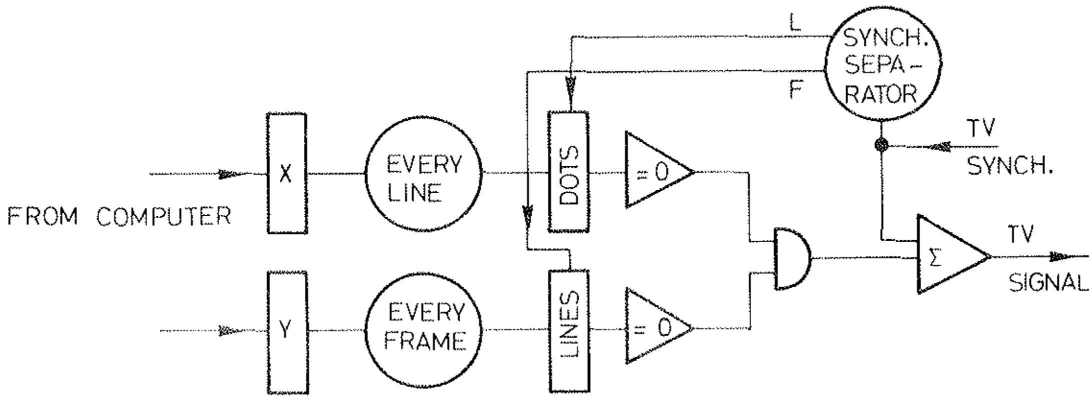

2.7 Modes of the cursor

In practice, the cursor bright-up is more than a single dot. By means of "one-shots" or delay circuits the bright-up can be expanded both in X and Y to give a little square about a microsecond wide and three TV lines high. For convenience, the cursor is also supplied in various other forms, by brighting up for all X-values or for all Y-values, or for both. The following cursor modes can be set:

small square (described above) at (X, Y)
crossed lines at (X, Y)
horizontal line at (Y)
vertical line at (X).

The cursor mode can also be set to make the signal flash on and off about twice a second. This mode is available with all the four forms described above, and is equally annoying to watch in all of them. We would not bother to make it available if we were designing the system again.

3. THE CHOICE OF MONITORS AND SCREENS

Pictures can be presented in a variety of ways, but the cathode-ray tube is the only practical choice. A number of exotic methods were considered originally, but most of them are impracticable in our case.

3.1 The plasma panel

Plasma displays can be purchased, which enable individual pockets of neon gas to be made to glow selectively. A panel 30 cm × 30 cm is obtainable commercially from an American company (Ref. C6). The panel allows the presentation of an array of 512 × 512 perfectly circular dots, and graphs drawn on it have a pleasing orange-red hue and very good sharpness. Unfortunately the plasma panel does not tie in very well with the concept of separating the problem of producing the picture and displaying it. Since the actual picture generation is done on the screen itself, the picture cannot be produced elsewhere except by regenerating it on a plasma panel elsewhere or by closed-circuit TV. The only other way to use a plasma panel is to staticize the TV signal, split it into 512 × 512 dots, and drive the plasma panel with this information. This is practical, but elaborate. It still does not allow the screen to show closed-circuit television pictures from cameras when this is required. The plasma panel is a two-level (light-dark) device.

3.2 Liquid crystal displays

Liquid crystal displays suffer from the same defects as plasma panels, as far as we are concerned. In addition, they are unstable and tend to be temperature sensitive. No liquid crystal display producing a matrix of dots is available commercially. Indeed the only liquid crystal displays of any kind available at the time of writing are numerical indicators, usually of the seven-segment type, or alphanumeric indicators of the starburst type. Both these arrangements will be explained under "large displays" below, but it will suffice at this point to say that they are not suitable for the construction of general-purpose screens.

3.3 Why cathode-ray tubes?

Of those devices currently available, the television monitor is the ideal one for display. This is not to say it will not be replaced by other devices in the next few years, but at the moment there is no other device which will display the necessary 250,000-500,000 dots so cheaply and reliably, and with the additional availability of many grey levels and colours if necessary.

The choice of black-and-white monitors is easy. Monitors of the highest "television studio" quality are more linear than domestic TV sets, and have smaller spot size. They are very expensive and often quite bulky, and since the control room was to contain some 30-40 monitor screens in all, it was decided to make do with monitors of moderate quality in most situations. Appendix III gives the specification of the most commonly used black-and-white monitor screen in our control room.

Colour screens are another matter. The commonest device is the "shadow-mask" tube, in which the screen is covered in a regular pattern of dots of phosphor material glowing in three primary colours, R (red), G (green), B (blue):

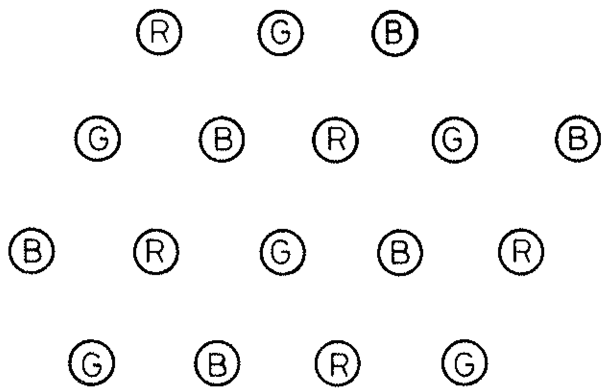

and the tube has three electron guns instead of the usual one. Each gun carries the signal for its own colour, and a mask with a similar pattern of holes is situated just behind the fluorescent screen and so placed that the beam from each gun can under no circumstances reach the phosphor for either of the other two colours.

The resolution of a shadow-mask monitor is limited, like that of all CRT devices, by spot size. However, a further limitation is the fineness of the shadow mask and its associated phosphor-matrix. The most obvious assumption would be that the finer the mesh of dots, the better. In practice this is not so. An improvement in the fineness of the dot structure is only worth having if perfect precision is maintained, both initially and under different temperature conditions. Normal domestic television receivers give about 580 dot-cycles across the screen, and by paying for a high-quality studio tube about 650 dots per line may be obtained.

3.4 The Trinitron

A disadvantage of the shadow mask monitor is the need to adjust the electron beams so that the shadow mask is correctly aligned for "convergence" of each of the three colours and the ability to produce a pure white image at all points of the screen. This process is very complicated and lengthy, because the controls are in the X and Y directions, while the dots of the shadow mask are on lines at 120° to these axes, making the adjustments interdependent. To avoid this, a Japanese company (Ref. C7) has made a tube called the Trinitron, in which the phosphor is in vertical stripes instead of a dot-pattern. The Trinitron is easy to adjust because the correction axes correspond to the deflection axes; it has only one electron gun and is much lighter and cheaper, and was therefore seriously considered for our control room. Unfortunately, being conceived for cheap television sets, the Trinitron tube has only 420 triplets of vertical stripes.

The most extreme criterion of resolution on the face of a cathode-ray tube display is the requirement to represent 64 characters on a line, each character having some 8 dots in its width or 512 dot-representations per line. This is a very difficult criterion to satisfy, as will be seen below.

It should be noted that the resolution of a colour screen can equally be given (and often is) in megahertz. Since the drawing of a single horizontal scan-line takes about 52 μsec (with 12 μsec for the flyback to the beginning of the next line), and since by a sampling theorem of Nyquist the highest frequency representable is half the number of samples per unit time, we have the following bandwidths:

"Trinitron"	4.04 MHz
domestic shadow mask	5.58 MHz
high-quality tube	6.25 MHz
signal requirement	9.84 MHz

Thus it can be seen that even the best monitor tube of this kind is not really perfect, and that nothing less than the best should really be considered. In practice, observers are so impressed by seeing pictures of this kind in colour, instead of the usual black-and-white, that they usually forgive the slightly reduced quality.

3.5 Exotic cathode-ray tubes

Unlimited resolution can be obtained by the use of special tubes. The penetration phosphor tube has red and green phosphor layers separated by a metallic layer a few atoms thick. The energy of the electron beam is used to determine whether the second layer will glow, and therefore modifies the colour. This tube is extremely difficult to use, as the change of energy is accompanied by a change in the electrical characteristics of the tube, requiring special circuitry. Thus in the end, the grain-free structure of the picture is paid for by great complexity. In addition, the requirement for switching high-voltage circuits in order to change colour makes it impractical to perform a colour change in less than 20 μsec. This renders the penetration tube unsuitable for television raster use. Another promising experimental tube exploits the physiological fact that a mixture of a green phosphor with a saturating characteristic and a red phosphor with a quasi-linear characteristic will appear to change in colour from green to red as beam current is increased. [This tube was described to the author, in a private communication, by a worker at the Royal Radar Establishment at Great Malvern in England[21].] This saturation tube is easier to use than the penetration phosphor, and has promise; but it was in no condition of development for inclusion in a working system, and in any case only promises red, green, and a questionable yellow colour.

3.6 Conclusion

Taking all things into account, therefore, the decision was taken to use high-quality studio television monitors, having a 650-line shadow-mask tube, a circuit bandwidth of 10 MHz, and correspondingly strict criteria on spot-size distortions and linearity (Refs. C8, C9).

4. THE PROBLEM OF LARGE WALL DISPLAYS

In a control system of the kind being described, one is never far from a TV monitor, since it is the very essence of the system that all signals are in the standard form and are easily transportable on coaxial cables. However, it is of some advantage to have a large display on the front wall of the control room, which can be seen from all parts of the room and also by visitors looking at the control room from the observation gallery. Attempts have been made to include such displays in our control room.

As for the problem of normal graphics, the production of large pictures can be divided into two classes, one for character information and the other for more general patterns, and as in the other case, characters are easier to produce.

4.1 Possibilities

On a large display, characters are needed which are between 2 cm and 10 cm high. There are various methods of generating such letters:

i) mechanical flap units, as used at airports;

ii) other mechanical devices;

iii) characters made of lamps or light-emitting diodes;

iv) cathode-ray tube devices.

4.1.1 Airport flap units

Mechanical flap units are used at most airports in the world to announce the time of arrival and departure of flights. For each character position to be displayed, a number of flaps, each printed with the top half of a letter or number on the front, and the bottom half of another on the back, are connected radially to the outside of a cylinder:

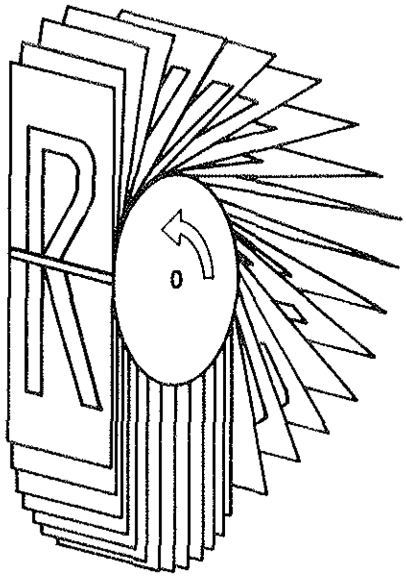

The top of one character is held by a small hook, and when the cylinder is turned one step by a stepping motor, the next character drops into place. A counter keeps pace with the number of pulses given to the stepping motor, and this counter is compared with the desired value of the character, enabling the device to be stopped when the right character is in place. Flap units are a little noisy and unreliable, but they have the advantage over most other devices in that they improve with increase in ambient light, rather than becoming less visible.

4.1.2 Mechanical 5 × 7 dot devices

Other mechanical devices are available in which 35 mechanical relays driving little white flags are arranged in a 5 × 7 matrix which can be driven by suitable electronics to give an alphabet and numbers. Such devices, in addition to being mechanical and therefore inherently unreliable, are very expensive, since 35 separate mechanisms are needed per character. An English company (Ref. C10) has reduced the expense of such a device by representing each relay by a small plastic flap with a piece of iron embedded in it. A line of seven electromagnets (representing a vertical column of a character) travels on a carriage along the line of writing and sets the flaps to a black or white position. The flaps being cheap, and the carriage, which is the expensive element, being shared over a whole line of the display, the advantages of a mechanical device can be obtained at a lower cost.

4.1.3 Electronic 5 × 7 dot devices

Characters can be made of matrices of 35 small lamps. This arrangement has a number of disadvantages: fairly large currents must be switched, making a transistor amplifier obligatory for each dot and therefore raising the cost prohibitively, and also a current is needed to maintain the display. For a large panel this can require several hundreds or even thousands of watts, producing quite a lot of heat and making reliability somewhat harder to maintain. The discrete wiring made necessary by the use of the transistors makes the price of such an arrangement prohibitively high for most purposes; although a large panel of this kind was made for the Olympic Games in Munich in 1972[22], in a situation where money was effectively no object. Letters and drivers of this kind are available in totally integrated form, including light-emitting diodes. Such an arrangement could be considered in spite of the very high price, but the letters are mounted on dual in-line packages of standard type, and the minimum separation between the individual letters in a group or between groups on adjacent packages is too large for successful reading as text. They are useful only to represent rows of individual letters, not words.

Liquid crystal displays are becoming very interesting in some fields[23]. Substances exist which when subjected to electrostatic fields can change from transparent to slightly reflective, or vice versa, or even change colour. Most of these devices are available only in the form of a seven-segment numeric display:

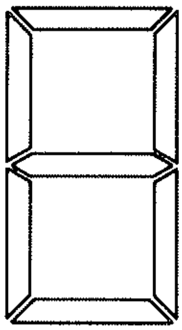

where selected bars change to or from the semi-reflective state, and by suitable illumination can be made bright on a dark background. The seven-segment arrangement can be used only for representing the numbers from 0-9. To obtain the entire alphabet it is necessary to use the so-called "starburst" pattern:

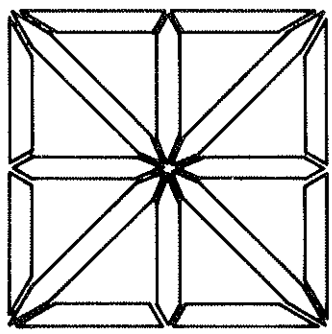

also called a 16-segment display. Very few suppliers produce a starburst liquid crystal display, and the one company which made an integrated circuit decoder-driver which might have made this arrangement feasible took it off the market again because of technical difficulties and a very small demand. The main reason for the small demand is that the starburst alphabet is rather ugly and hard to read

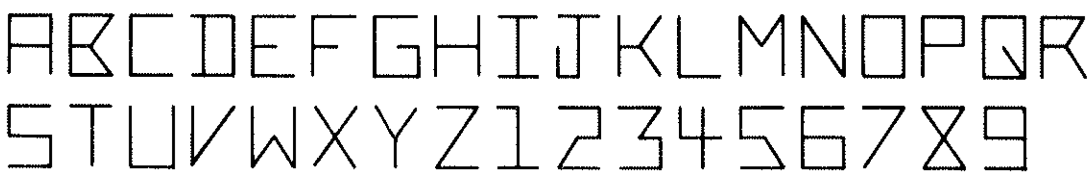

A page of information in starburst patterns is really rather disturbing, and this, combined with the very high cost and tricky technology, caused us to abandon liquid crystal displays after a few essays in the laboratory. This is not to say that when liquid crystal displays with 5 × 7 characters appear on the market in economic form this decision will not be reconsidered.

4.1.4 The multiscreen method

In actual fact the only way to generate large, readable displays cheaply, reliably, and easily is by the use of the cathode-ray tube. The cathode-ray tube owes its success to the fact that the display surface itself is only a powder coating on glass and is therefore cheap, and that the expensive part (the electron gun, the focusing deflection, and modulation arrangement) is shared over all the display elements.

We can make a few large characters on a cathode-ray tube by a simple modification of the character display described in Section 1 of this chapter. By grouping six monitors in a block, two horizontally and three vertically, a useful display about 1 metre square and having some 12 rows of 32 characters can be produced. The characters are about 5 cm high and are made by software. Since the information content of each monitor is low, one character controller can be made to drive all six monitors. The software character driver economizes in character-width by making the width of each character suitable: for example an I is narrow and an M is wide.

Mechanical details are that special monitors were engineered, using cathode-ray tubes of the most rectangular form possible, and having the chassis behind the tube instead of below it as in the conventional monitor:

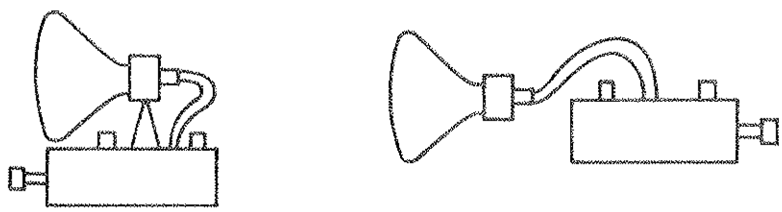

Conventional monitor   Special arrangement

The construction and delivery of a complete block of six monitors of this kind was entrusted to a company which supplied the entire unit ready to use. A sheet of lightly silvered clear plastic covers all six monitors and conceals them. The silvering is about 1 micron thick and is effectively transparent to the letters, which glow through it.

The cost per character of the entire device is about 20% of the cost of an equivalent panel of airport flap units, with the additional advantage of electronic, and therefore reliable, Fast operation. Updating of parts of the display every accelerator cycle (repetition rate ∿ 5 sec) is quite feasible, as is the inclusion of a clock changing every second.

5. PROJECTION OF GRAPHS AND CCTV PICTURES

The multi-CRT panel, while it produces a very nice, large-character display, does not present a general television image from any signal in the system. For this purpose a number of large-scale display devices were considered:

i) conventional projection TV

ii) exotic projection TV devices

iii) hard-copy on transparent material, projected with an overhead projector.

Conventional projection TV works quite well. A special TV tube, about 8 cm in diameter, and having a reinforced cathode and high operating voltage, gives a very bright picture, which is projected onto the wall by a large-aperture mirror. The curvature of the tube-screen and the spherical aberration of the mirror are corrected by a lens between the mirror and the projection screen:

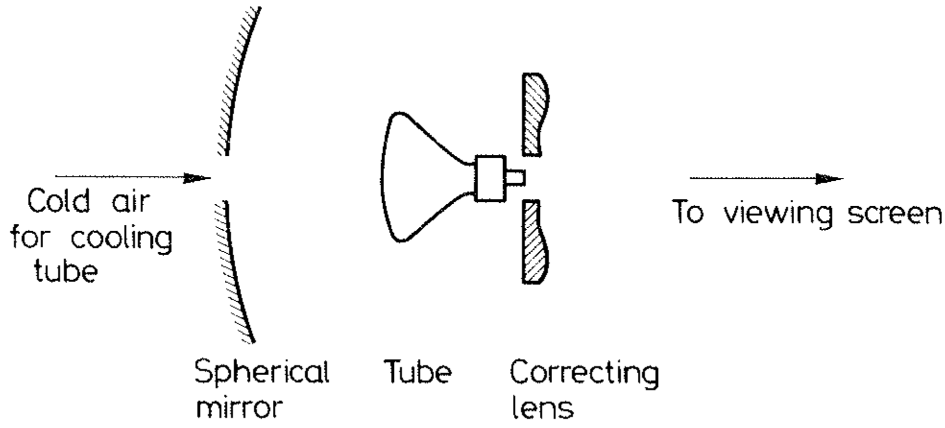

The brightness of such an arrangement is lower than that on the surface of the cathode-ray tube by a factor of 8 $F^2\mu^2$, where F is the aperture and μ the magnification, and in practice this means that projection TV can only be observed in a fairly dark room (20 lux maximum ambient illumination on the screen). Attempts to make the image brighter result in the production of excessive x-radiation from the high voltages employed and the shortening of the tube life below the acceptable 500-hour level.

6. LIGHT-VALVES

Various attempts have been made to obtain brighter pictures by projection. If the cathode-ray beam can be used to switch an external source of light, this latter can be much brighter than the glow of a phosphor. Such a device is the Eidofors, a patented device in which the target inside the CRT is an oil-film. The electric charge built up by the beam distorts the oil-film locally, causing the picture to be built up as a relief-map on the oil. An ingenious optical system using the interference principle is used to form a projection of this relief in dark and light. The Eidofors works, and is available commercially, even in a colour version. The price of a black-and-white version is about ten times that of projection TV, but this is not its greatest drawback. The fact that there is oil in the vacuum chamber of the tube means that the cathode becomes poisoned by its vapour and ceases to emit electrons. The cathode is therefore demountable and the equipment includes a vacuum system, rather like an electron microscope. A mechanic is needed on-site to service the Eidofors every few operational hours, and this renders its use quite impractical in a control room, where 24-hours-a-day service is required.

In an attempt to make light-valves by alternative methods, various ways of attacking the problem are being attempted. Certain crystals turn black on suffering electron bombardment, and this phenomenon, known as "cathodochromism", has been exploited in the so-called "black" cathode-ray tubes. Erasure is obtained by heating the crystal, and the speed of erasure is limited by the thermal strength of the face-plate. The number of write-erase cycles the crystal will accept is limited. Interest in cathodochromic devices was at a peak some years ago, and one or two references are given[24-27]. The author was advised by a research worker at one of the laboratories involved that they spent some years on cathodochromics research and finally abandoned it.

Hard-copy on a transparent (for instance, mylar) sheet can be projected in an overhead projector. We have done this, and it works. It involves a delay and will not be used in the control room.

7. CONCLUSION

On consideration of all the devices available for large-scale display, it was finally decided to order a conventional TV projector. The six-screen device will probably be installed as a supplementary device, to show the salient accelerator parameters continuously. The TV projector will be available on the video-switch, with standby switching to extend the tube life. It will be operated in the back-projection mode.

CHAPTER IV

THE GRAPHICS SYSTEM (SOFTWARE)

In the previous chapters, we have discussed how a modular graphics system, based on picture generators and monitors and held together by a picture switching and summation system, will satisfy the display requirements in the control room. It will only do this, of course, if steered by suitable software, and in this chapter we shall discuss some of the computer programs which have been written to this end.

1. THE USE OF THE INTERPRETER

The constraints put on the design of a graphics programming system included:

- the distributed nature of the picture-generation equipment;
- the existence of a special graphics computer, which was to relieve the rest of the system of this responsibility;
- the existence of the NODAL language, and its multicomputer and real-time abilities.

It was very soon decided that all graphics equipment should be attached to the display computer, and that a considerable simplification would be possible if only this computer contained graphics software. A program to draw (say) a graph, should be resident in the display computer, and the data for drawing that graph can either be in the memory of that computer or can be demanded by it when it needs it:

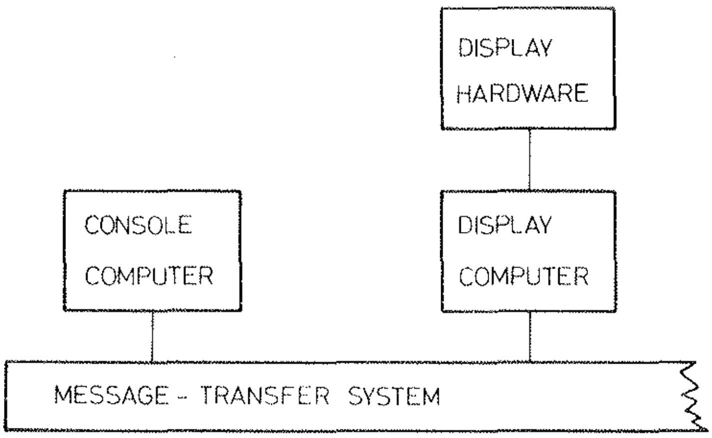

In the normal way, the request for a graph to be drawn will originate at the console. The graph-drawing program will initially be resident in the console computer. It will contain an initial part which amasses information from the real world of the accelerator, performs calculations, and then deposits the coordinates of points on the graph in local storage arrays.

The program which draws the graph is part of the coding in the console computer, but is never obeyed there. When the program is ready to draw a graph, it sends this graph-plotting code, together with the arrays containing the data, to the display computer, where the graph-plotting process is performed. In this way it is not necessary for any computer, except the display computer, to understand display commands, and yet any computer can initiate graph-plotting.

The need to keep the load on the data links connecting the computer low, restricts the frequency with which graphs can be drawn in this way. In general, graphs will be re-drawn at most once every accelerator cycle of 5 sec. The continuous motion of elements on the screen by re-drawing them at 100 msec intervals had already been precluded at the design stage, and the software technique described here does not impose further restrictions.

There is one exception to this: a cursor may be made to move on the screen by means of a track-ball or knob. This cursor is, however, not drawn on the rest of the picture by the normal route, but is generated directly in the console computer and added to the graph at the monitor input. Thus the rapidly repeated conversations involved in re-drawing the cursor repeatedly as the hand moves the ball, involve only a conversation between the console computer and the cursor hardware:

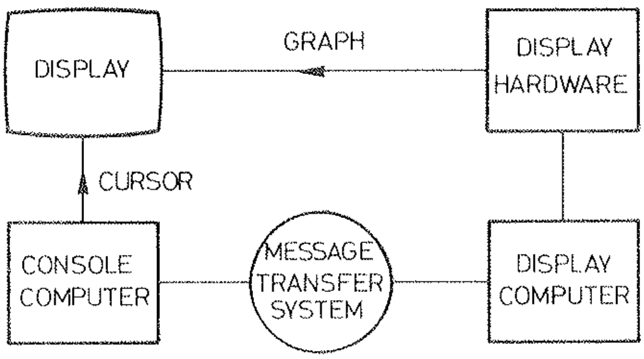

Another quite serious problem which had to be solved at the design stage is that of mapping. Initially, this looks very simple. Graphs can be drawn in a conventional X-Y space, with the axes at the bottom and left, and with some conventional coordinate system, such as 0-1000 in each direction:

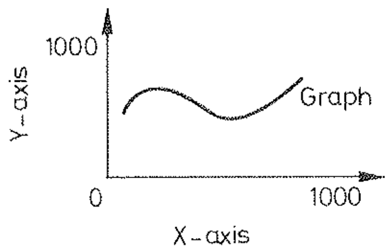

Such a naive approach immediately breaks down if an attempt is made to cover the whole screen, whose width is 4/3 of its height. Circles will look elliptical. Thus either the screen must be 1000 × 1333 or 750 × 1000 counts.

A further complication arises when the character-type of display is considered. The character display consists of 24 rows of 64 characters, and convention decrees that in this case Y be counted downwards:

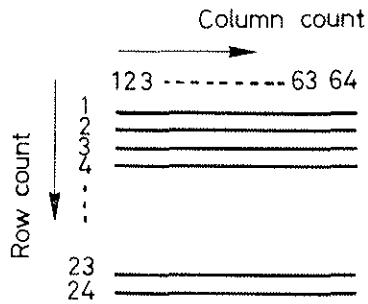

Just to make things even more difficult, it will be noticed that rows and columns are usually counted in a system which starts with 1, while X × Y plots usually refer to origins at zero.

It can be seen that the problem of designing a homogeneous mapping system which applies to character displays, graphics displays, and cursors is formidable, and may involve us in a number of compromises.

The solution which has been adopted is that the ball and cursor and the dot graphics system are all mapped on the system with the origin in the bottom left-hand corner, and that people who wish to combine graphics with character pictures have usually to use a mapping formula. Even this manipulation can be avoided by the use of a special instruction, which says, in effect: "I wish to write characters, starting at the point whose dot-coordinates are X, Y". However, it seems to be unavoidable for a programmer who wants to use both graphics systems simultaneously to understand this problem.

In all other respects, it has proved possible to make the user software easy to use in the most obvious way possible. Figure IV.1 shows the standard layout form used for planning pictures.

2. THE CONCEPT OF A MULTI-LEVEL SOFTWARE SYSTEM FOR GRAPHICS[28,29]

The NODAL interpretive language used for the whole control system can also be used for graphics, provided certain statements are added to enable the graphs to be drawn.

These statements can be of various types, from a simple one enabling a dot to be produced at a point (X, Y), to an elaborate process requiring the collecting of real-time information at fixed time intervals and the production at the end of a fixed period of a histogram marked with axes and descriptive text.

There is always a problem here, in that there are various kinds of users -- the one who wishes to use the graphics systems for elaborate new investigations, and who is therefore willing to invest considerable time and study, and the physicist or engineer who has a problem in a completely different field and is therefore much more concerned with drawing the graph or diagram he needs than investigating the intimate details of computer science. This problem, like that of field mapping mentioned before, can only be solved by a compromise.

The compromise, which corresponds nicely with the mode of solution of the problem, is to build a hierarchy of subroutines, in which each level of the hierarchy uses the routines

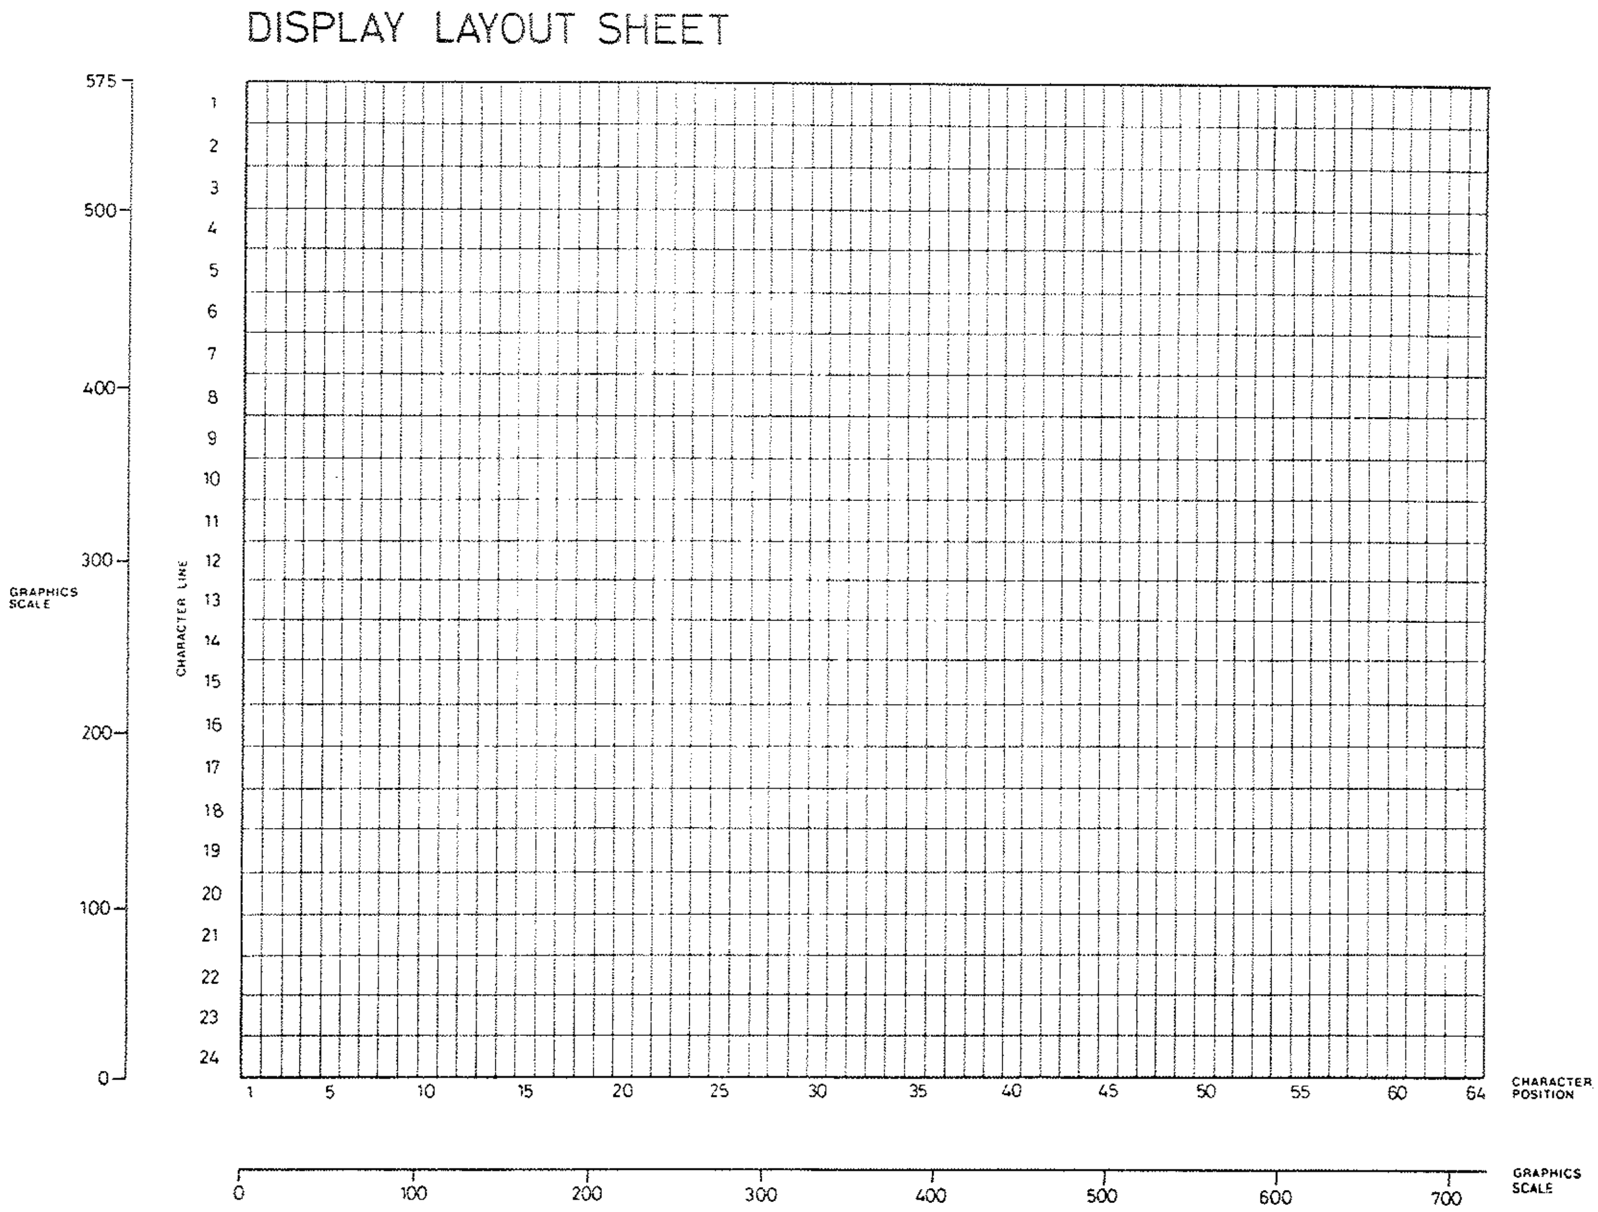

Fig. IV.1 A display layout sheet

in the lower level as bricks. Thus the lowest level routines allow the following operations:

First level

1) Reserve graphics facilities and release them.
2) Connect reserved facilities with screens.
3) Clear the drawing field.
4) Define the part of the screen on which drawing is to take place.
5) Choose a writing position.
6) Draw simple elements, like dots, vectors, circle segments, and text.

Using these simple elements, a patient programmer can construct most graphs. Since the dot memories have their own storage, the elements of a picture can be added one at a time, after an initial clearing of the screen. If necessary the picture can be built up as the data are acquired, points being added when their coordinates become available.

Most programmers do not want to go to this kind of trouble. They would like to write a program to acquire their data, store it into an array, and then call a plotting routine. Such a general-purpose plotting routine must be given only the minimal facts to be useful, and with this small number of inputs it becomes inflexible. Thus some higher-level bricks have been designed which

Second level

1) accept the ranges of the variables and perform scaling;
2) plot a histogram (vertical bars);
3) plot a dot histogram (this is an unjoined graph of Y for equal increments in X);
4) plot a graph for equal increments in X and Y;
5) plot a scatter diagram from X, Y pairs;
6) plot a joined graph from X, Y pairs;
7) draw axes, complete with tick marks and annotations.

Users of the accelerator who are doing experiments in machine physics will often work at this level. Half-a-dozen statements of the kind described above suffice to draw quite elaborate structures on the screen, and changes are easy to make.

The next level of subroutine for graphics tends to concern interaction between a graph and an operator: a number is culled from the screen by pointing at it with a cursor, for instance, or control is given to a knob of a variable chosen in this way. Bricks belonging in this category are, for instance, as follows:

Third level

1) Use the track ball and cursor to choose an area of the screen to be enlarged, and then re-draw that part of the picture (this is the so-called "zoom" facility).

2) Give knob control of a variable whose name is on the screen, and report back the real-time value of that variable.

3) Rotate or otherwise modify picture elements on demand.

4) Plot a repetitive graph in real-time, with its history in colour. Thus the current graph might always be drawn in green, the previous graph appears in red, the previous one to that in blue, and all older graphs are lost.

Structures like these can be constructed by hand from second-level subroutines, and indeed the third-level routines available are constructed in this way, but the amount of thought required is quite large, and once the job is done, it is an advantage to have it on file. It should be noted, however, that a user for whom the definition of a third-order subroutine is slightly wrong has the alternatives of building something for himself out of second-order subroutines, or he can modify a suitable existing routine. This latter procedure is usually easy, as all routines above the first order are written in the higher-level (interpretive) language.

3. HIGH-LEVEL ROUTINES FOR GRAPHICS

The fourth and last level of graphics software is that which is actually used for the process of control. Such programs are those for

1) emittance plots,

2) beam position diagrams,

3) phase-space diagrams,

4) vacuum status histograms,

5) synoptic diagrams for services,

6) time-plots of variables,

and so on. Such programs are used for one purpose only -- that for which they are designed. They will never be called as subroutines by people other than the author, and they may therefore disobey all the rules of good planning and good programming. Indeed good design may make fault-finding and maintenance easier, but the main criterion for such programs is that they satisfy a physical purpose efficiently. The approach we have been using for the production of such fourth-order programs is that specialists in the various fields are presented with well-debugged routines of the first three orders and given advice on how to use them to solve their problem, be it in accelerator physics, vacuum technology, power engineering, or any other speciality. In this way the techniques of programming can be separated from those of accelerator physics, just as in earlier chapters we saw that the computer interface enables human interfacing to be separated from accelerator control.

CHAPTER V

INTERACTIVE DEVICES FOR THE OPERATOR

In the previous chapters we have discussed methods of presenting visual information, mainly in the form of tables of figures, but also as histograms, graphs, diagrams, and TV pictures from cameras.

In this chapter we shall discuss the methods by which the operator closes the loop, that is, the devices which the operator uses to inform the system of his decision. Up to now, this would have been an easier task. Control devices were mostly one-way arrangements, the knob and the push-button being typical. Some of the devices described here break new (or fairly new) ground by being two-way devices which both feed the computer and are fed by it.

Most of the earlier applications of computers fought shy of computer control of the machine elements. Observation was permitted, because it was felt that computer errors entering into an observation process could only spoil a display, and not the behaviour of the equipment. This attitude was justified in part by the fact that the reliability of a computer complex was usually lower than that of the equipment it was replacing.

With modern technology the computer may well be more reliable than any other method of information collection, if indeed any other method exists. The multiplexing of a large number of signals and their transmission and eventual decoding and display, if done by special-purpose electronic hardware, may well involve more equipment than the computer which would do the same job. Furthermore the computers, being off-the-shelf devices, may be easier to maintain and repair, and thus lead to a higher load-factor than is obtainable from a special-purpose system.

It is, however, not the reliability question which has forced computers on the accelerator control community. The fact is that computers have become necessary in view of their ability to collect, organize, compute, and display; and that once they found their way into the control system, it was only a matter of time before they would be used to control as well as display. Thus in the SPS control system, while there are less control signals than monitoring signals, yet control through the computer system is now a normal part of the scheme.

Some control is done directly from the computer. In this type of control, called "closed-loop control", the computer collects data from the equipment, and uses them as a basis for some corrective operation. Such closed-loop sequences are used in the accelerator, for example to correct the waveform of the main power supply from magnetic field information, in order to keep the magnetic field close to the necessary theoretical curve. For speed, and in order to keep the load on the data links low, such loops are usually confined to one computer and its associated equipment; but on an experimental basis, or temporarily, the console computer and the entire system may be involved. Closed-loop control is only possible if a practical algorithm exists and can be computed in an available time and amount of computer memory. For cases where this is impossible, operator or "open-loop" control must be resorted to. The expression "open-loop" is really a misnomer, since the loop is actually closed by the eyes, mind, and hand of the operator.

1. TOUCH-BUTTONS

1.1 Conventional buttons

For simple selection between a number of different choices, buttons may be used. Conventional control rooms use large numbers of such buttons. They usually drive relay-latches and thence change the state of the equipment, but some of the buttons are two-state devices. In the case of computer-controlled equipment the difference disappears. All buttons cause a machine interrupt whose source the computer can identify, and the action to be taken is a separate, computer-driven process.

The labelling of conventional buttons presents a problem. The purpose of the button is usually written on or above the button itself, and this serves to identify its function. In the case of a two-state button, the surface of the button is often illuminated, or a lamp adjacent to the button is used for this effect.

In all these arrangements, there is one button per function to be controlled. This seemingly obvious arrangement has two big disadvantages: first of all, the number of buttons for an elaborate structure such as an accelerator can reach well into the thousands; and secondly, many buttons may have to be installed for expansion or for rare use and then clutter up the panels without performing any useful function for most of their lives.

1.2 Touch-buttons

To overcome these disadvantages, buttons have been devised with programmable legends. Such buttons effectively disappear when not in use, and the computer program can create the buttons it needs for the particular operation currently being performed.

A considerable advantage of the programmable button concept is that the arrangement has the aspect of a teaching machine. The user chooses from a number of possibilities, which in our case is 16, and having made this choice he sees the button he touched disappear, as well as those for the 15 choices he did not make. The computer now makes available to the user up to 16 new choices, which are all relevant to his current situation, and so on. At no point in the process is the user confronted with irrelevant buttons. In practice there is rarely any need for all 16 buttons to be active at once, and so the user can glance over the legends rather than having to read them all.

The way in which the buttons have been implemented, at CERN and elsewhere, often leads people to think of an X-Y digitizing device. The similarity is misleading. Touch-buttons are a quantum or digital, not an analog device, and the family resemblance is to a button, not to a joystick.

1.3 Physical principles for buttons

There seem to be two distinct principles on which buttons with programmable legends can be implemented:

1) A physical button is supplied with legend facilities.

2) The buttons are drawn on a monitor screen, and some method is used for detecting the presence and whereabouts of the button the operator wishes to identify.

Method 1 has been implemented. The Goddard Space Centre at Houston, Texas, USA, uses commercially available push-buttons with 16 miniature projectors in each. A small piece of microfilm has 16 button-legends on it, and switching on 1 of 16 miniature lamps makes the appropriate legend appear on the face of the button:

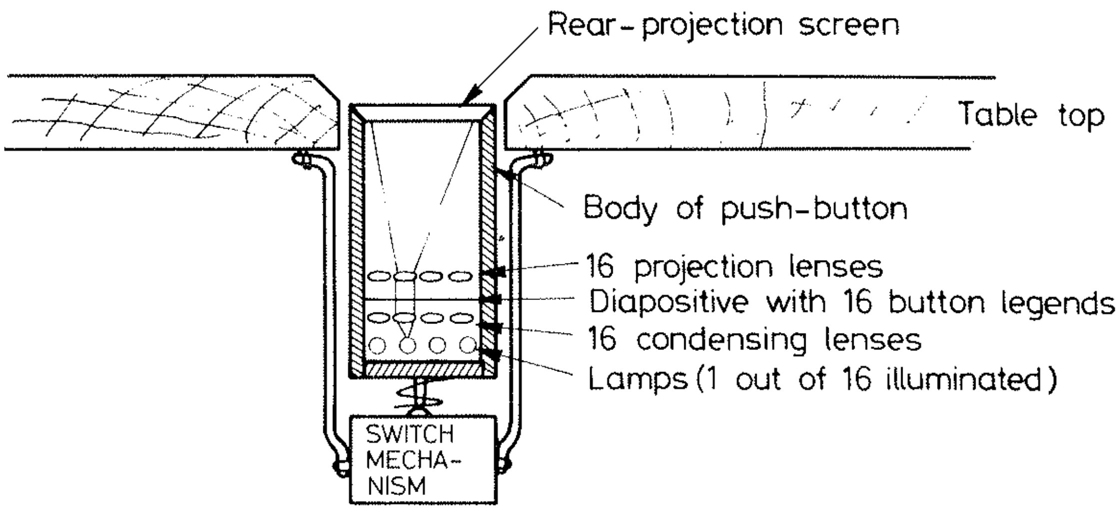

As could be expected, this kind of button is fantastically expensive. Only the space-program could contemplate many panels full of them. However, the interest lies in the fact that this was probably the earliest solution of the problem.

All other buttons are of type 2, in that the legends are conveniently written on the face of a cathode-ray tube, and used as buttons in some way or other. The earliest method of choosing such a legend was by a light-pen. This was used at the Los Alamos Meson Physics Facility in this context and then abandoned. It is often used for non-control applications, as a tool in computer graphics processes. For completeness the light-pen will be described briefly below.

The tip of the pen contains a small photocell, which is connected to the electronics through an amplifier and a cable. A small button on the pen serves to connect it temporarily. The pen is pointed at the CRT phosphor and the button is pressed. A pulse is sent to the computer when the spot on the screen next finds itself under the light-pen. Since the instantaneous position of the spot is either available in X and Y registers (in a random display), or can be kept in a pair of X and Y counters (in a television raster display), this position can be read into the computer and used:

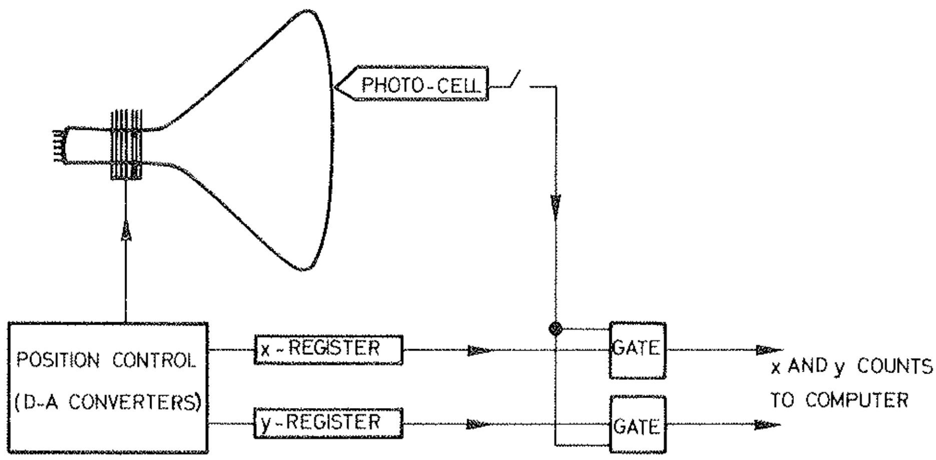

Light-pens work. The main reason for not using them is that the cable gets in the way, pointing them accurately is an acquired art (they must be held perpendicularly with respect to the screen), and that they usually cover the very feature it is wanted to indicate at the very moment the operator wants to see it. Pushing the button is an additional operation while the pointing is being done correctly. A light-pen might be useful in some contexts, though it is usually better replaced by a ball and cursor. It is certainly not a satisfactory substitute for a button.

The best light-pen is certainly a finger, which the operator always has by him, and several arrangements based on finger-detection have been tried, as follows:

1) The interruption of light-beams with the finger. The light-beams originate from lamps and are detected with phototransistors. This device has been used for data collection on a beam line at the Rutherford Laboratory in England. It is not cheap to do properly, and it suffers from the disadvantage of being a little bulky physically[30].

2) Acoustic waves are made to travel across a panel of glass[31]. The wave is stopped at the point where a finger touches the glass, and this is used to identify the location of the fingertip on the panel. This method also is fairly elaborate and bulky. It has also been abandoned by its originators in favour of another device (No. 3 below).

3) A mesh of wires is suspended above a mutually perpendicular set. The pressure of a finger brings one wire from each set into physical contact. The wires and supports can be made so thin and delicate that they are inconspicuous. The whole is protected by a plastic sheet. This method, although presumably having the unreliability of a home-made switch, has been used in air traffic control and at the Stanford Linear Accelerator. The author was told by the Stanford team in a private communication that wires do break in use, but not often enough to render the device unusable.

4) An arrangement can be made which works by picking up the general 50 Hz electrical hum in the room from the tip of the finger. Using a field-effect transistor per button this method can be made to work. However, there may be problems in constructing a transparent panel with sensitive points, sensitive in spite of wear and dirt.

5) A fifth method is available which had not been described before. The author suggested it in a paper on the touch-buttons: mechanical buttons, made of transparent material, are placed over the screen of the tube; each is connected to a switch mechanism beside the button, so that the writing on the face of the tube can be seen through the transparent material:

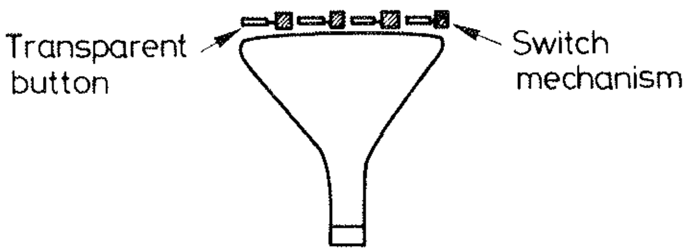

The only drawback from which this arrangement suffers is that the mechanism of the buttons obscures parts of the screen, which can therefore not be used for any other purpose. For this reason mechanical buttons were not used in this project. However, this idea was held in reserve and would have been used if the method described next had given trouble.

6) A set of capacitors is made on the surface of a sheet of glass by printed-circuit techniques. Each capacitor looks in principle like this:

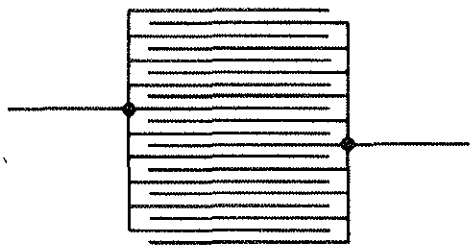

and the structure, being some 50% transparent, is barely visible. Printed lines, 80 μm thick and apart, are used. The capacitors are protected from actual contact with the finger by a thin film of varnish. The proximity of a finger increases the capacity sufficiently for detection. Since the capacitor has considerably more capacity than the leads to the circuit, touching the leads can be made to increase the capacity by an insignificant amount. In practice, by intelligent design of the printed-circuit layout, touching the buttons can give a 10% increase in capacity, while touching the leads or an adjacent button gives less than 1%. The difference is not difficult to discriminate. The actual detection is done by means of a phase-locked loop circuit per button, with a master oscillator for all the buttons:

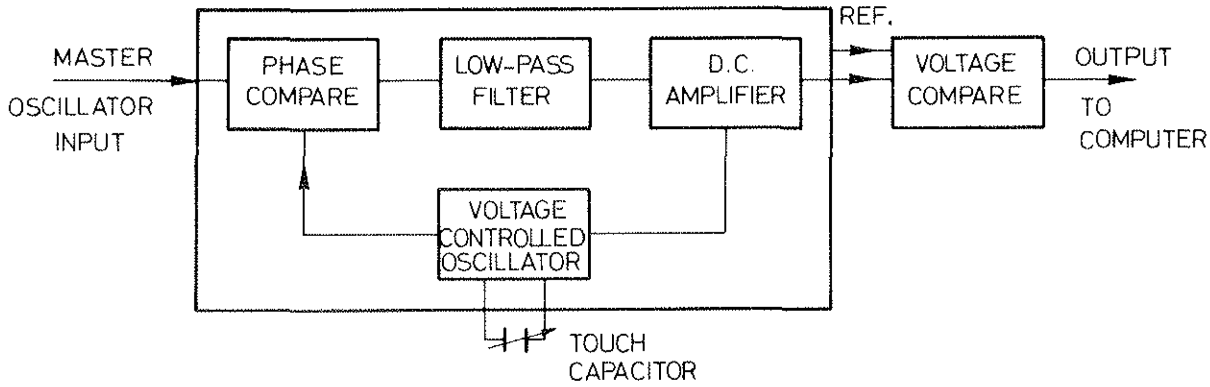

The oscillator, running at 120 kHz, locks in frequency to the master oscillator, and a change in capacity causes the phase, but not the frequency, to alter. A change in the error-correcting voltage is detected by discrimination against a reference voltage. This method, working as it does at constant frequency and amplitude at the sensitive point, is completely immune to all kinds of interference effects.

1.4 Conclusions

Method 6, the capacitive method described above, has been successfully installed. Simpler methods of detection for the change in capacity have subsequently been tried both at CERN and elsewhere, but no published record of reliable success yet exists. The screens are now being marketed in Denmark by one company (Ref. C3), and a complete copy of the device is being prepared for marketing by another (Ref. C11). See also Refs. 32 and 33.

2. THE KNOB AND THE BALL

Operator input of quantities to the computer system and thence to the accelerator can be by a number of methods:

i) by typing from a keyboard,

ii) by setting values on thumbwheels,

iii) by turning a knob,

iv) by steering a ball or joystick.

The first two methods have limited utility. When the value to be set is precisely known, there is usually some reason for this: either the value results from a computation or it is obtained from a book or table. In the first case, the computer program which computes the value may also perform the work of setting it in the hardware via the computer system, rendering manual input unnecessary. In the second case, either typing or the setting of a thumbwheel will do, but typing is often quicker and easier. The electronic thumbwheel, which does carrying of the tens automatically both on incrementation and decrementation, is often quite convenient; but for continuous control of a variable, where in any case some output is being watched on a screen or other indicator, the knob or track-ball is the device of choice.

2.1 Encoders

To read into the computer the degree of rotation of a spindle being turned by the hand, an encoder is needed. Two basic methods exist: a variable resistor followed by a digital-to-analog converter, or an optical encoder. The resistor is simple and fairly cheap, but results obligatorily in an absolute encoder. This can be difficult for the software to use. As will be seen below, the optical type of encoder can also be made absolute, but the relative kind is always more convenient except when the encoder is being used for precision measurement, in other applications than this one.

An optical encoder consists essentially of a light-source and a light-detector, separated by a transparent disk on which patterns are printed in black. As the disk rotates, the patterns interrupt the light arriving at the detector, and a series of pulses is emitted which informs the computer through suitable circuitry, of the direction and amount of rotation.

Absolute detectors can be made by having light-sources and detectors at various radial distances on the disk, and disposing binary patterns around it:

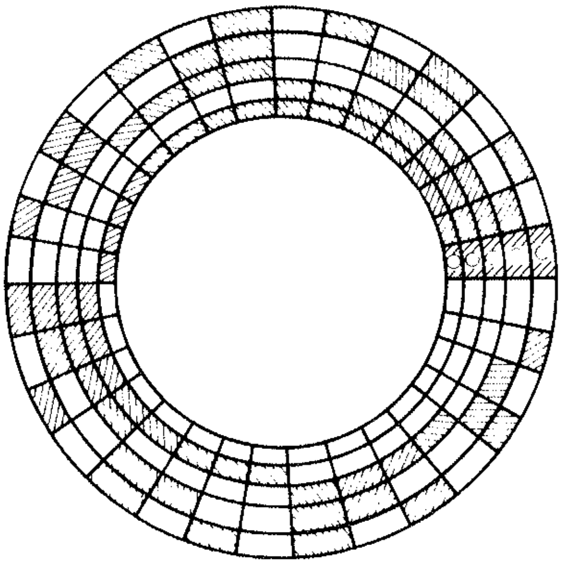

The disadvantage of a normal binary code in a case like this is that if the detectors have any non-radial alignment, the transition from one value to another, for instance from 31 to 0 in the case above, can be indeterminate. To counter this, the order of the binary numbers is scrambled in a systematic way, such that any two adjacent numbers differ only by the addition or subtraction of a single binary digit. The following pattern, which is encoded in a so-called "Gray code", satisfies the requirement, and is usually used:

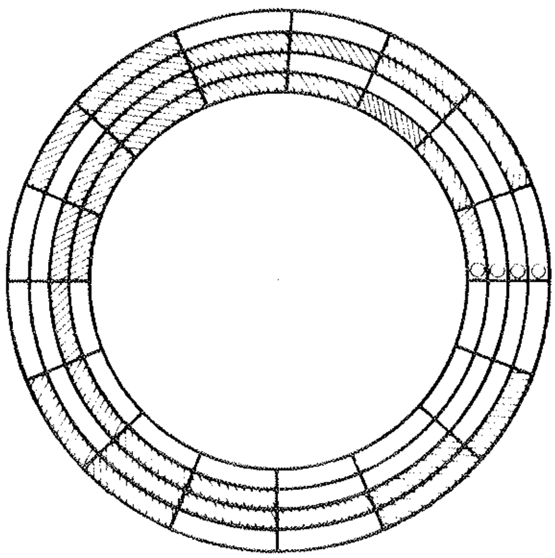

This kind of encoder is often used for the precise measurement of angles in a single revolution. As mentioned before, it is absolute, and therefore the measurement of the angle between two positions requires a subtraction, which may cause difficulties if a complete turn has been exceeded.

A relative encoder shifts this problem to a counter, or even to the computer. A large-capacity register can be supplied, which can accommodate the counts from many turns of the spindle. The encoder disk is printed with equally spaced radial lines, and two pairs of lights and detectors are used, so that each may emit a square wave with a fixed number of pulses per revolution. Relative encoders can also use mechanical, magnetic, or capacitive methods to emit the necessary pulses. The two detectors are so arranged that they emit pulse trains separated in phase by 90°.

Figure V.1 shows how these pulse trains are detected and counted. Effectively, one wave is delayed by 90° and compared with the other. The one which gives coincidence is that which was originally 90° ahead in phase, and the pulses from this one are counted. In practice, the 90° phase shift is obtained by multiplying the frequency of each square wave by 4 and giving a digital delay of one pulse. Thus the counter is incremented or decremented by four times the pulse rate given by the printing on the wheel.

2.2 Computer-controlled knob

An encoder as described above can be mounted on a panel, connected to the computer through an interface, and used as a control device. Encoders of this kind are available commercially, as are up-down circuits with sign detection.

The knob which is described here, however, does more than this: it can be commanded by the computer to behave in a number of different ways so as to simulate various control devices usually rotated by hand. The three modes of operation in the original model were the following:

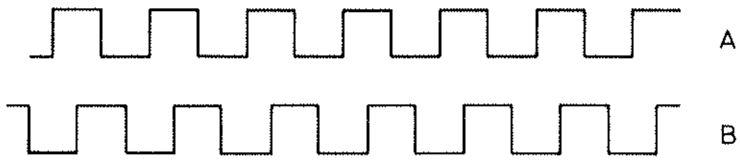

a) Pulses from encoder when the spindle is turned in the positive direction (A leads B by 90°)

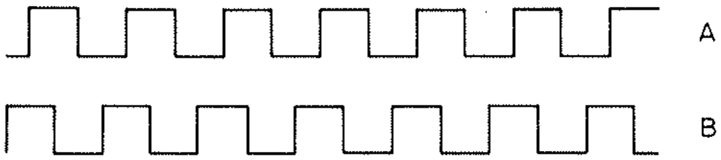

b) Pulses from encoder when the spindle is turned in the negative direction (B leads A by 90°)

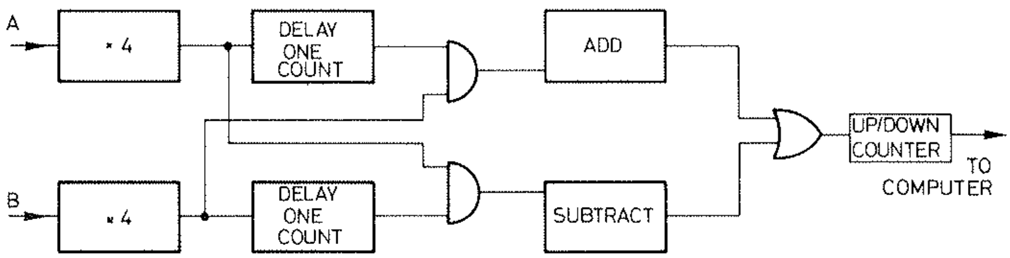

c) Method of counting direction and magnitude of rotation

Fig. V.1

1) A knob with programmable resistance to turning, all the way from free movement to complete locking. This is useful both for giving the operator "feel" of the operations he is controlling, and also for limiting his actions, when necessary.

2) A knob with spring-return to zero, with a small triangular pointer, apparently attached to the knob. This can be used for controlling the velocity of a variable, rather than its position.

3) A multi-position switch, with detents and visual indication of the currently selected position. The number of positions per revolution has been set at 16 so that each index position of the switch shall cause the counter to increment by an exactly integral number (512/16 = 32), making the switch suitable for use in manually controlled selection processes.

2.2.1 How the knob works

Figure V.2a shows how the knob is arranged to be both an input and output device for the computer. Figure V.2b shows the mechanical arrangement. The encoder is permanently attached to the operator knob through a gear train. Other devices are attached through gear trains containing electromagnetic clutches. The energizing of a clutch causes a particular device to be attached to the otherwise freely moving knob, thus restricting its movement in a suitable way. Figures V.3a, b, c, and d show these phases diagrammatically.

The resistance to movement of the knob may be controlled by energizing an electromechanical brake with a current supplied from a digital-to-analog converter whose digital input is supplied from the computer. A digital value from 0 to 255 gives a proportional current, and this range of currents allows the knob to run freely at zero current and to be locked completely for maximum current. The brake used appears to introduce smooth proportional resistance at intermediate currents.

The knob may be turned into a switch by energizing a clutch and thus attaching an indexer. The indexer consists of a spring-loaded ball running on a wheel with 16 detents.

The spring-return-to-zero motion is obtained by energizing yet another clutch, and attaching a device using a cam and two springs, to obtain the necessary progressive return force as the knob is moved away from the home position.

Visual indication of the behaviour of the knob is obtained by lamps. One of these is in the centre of the knob, and shows when controlled "feel" is available. The others shine through the translucent front panel and are only visible when energized. They are on arms mounted concentrically with the knob and are driven by gear trains through the appropriate clutches so that each lamp only moves and shines when it is needed. The gearing is such that the indicator appears to be rigidly attached to the knob.

The entire device is encased in a cylindrical body, 287 mm long and 146 mm in diameter. Three knobs can just fit side by side in a 19-inch rack, though in actual fact one knob per console is thought to be sufficient.

2.2.2 How the knob is used

An interpretive command MKNOB (a,b) sets the mode of the knob.

a = 1, b = 0 gives an indexed switch

a = 2, b = 0 gives a spring-return knob

a = 3, b = N gives a knob with resistance N (0 N 255).

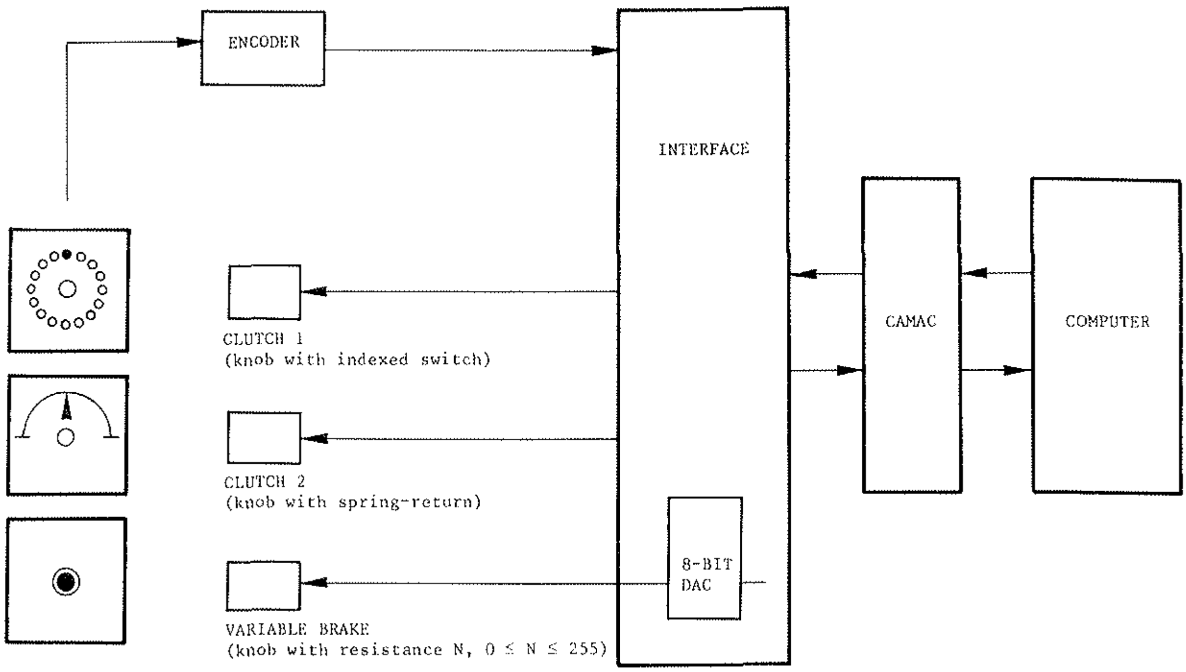

a) Over-all connections of the computer-controlled knob

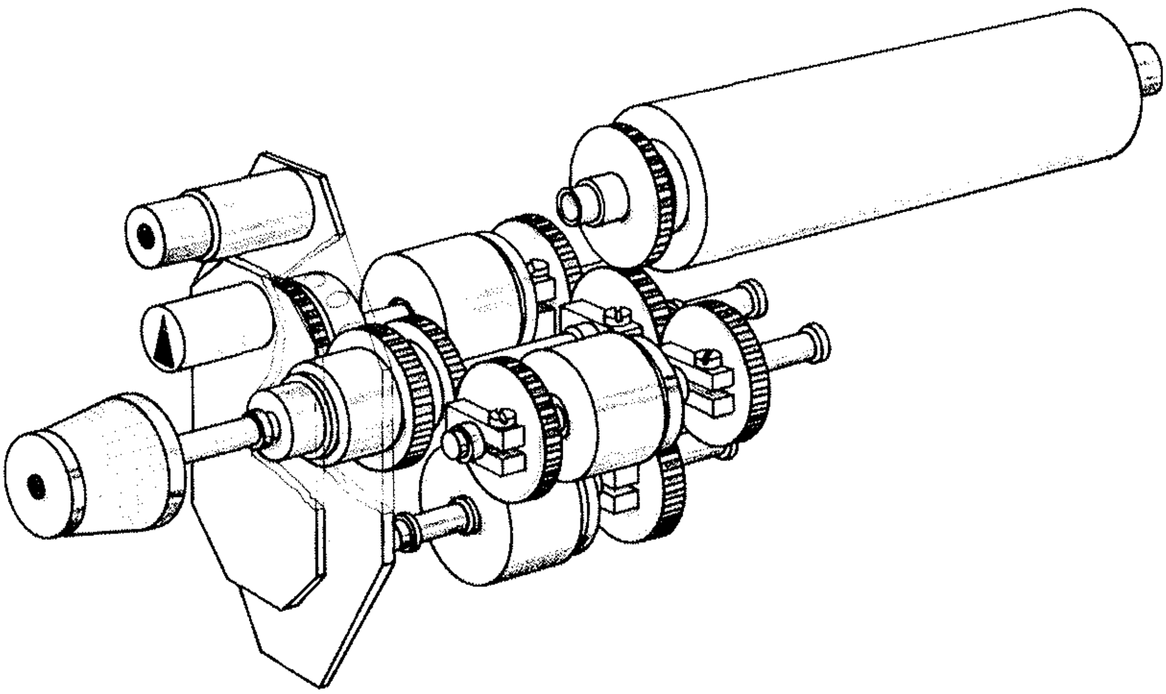

b) Computer-controlled knob (housing and translucent front panel not shown)

Fig. V.2

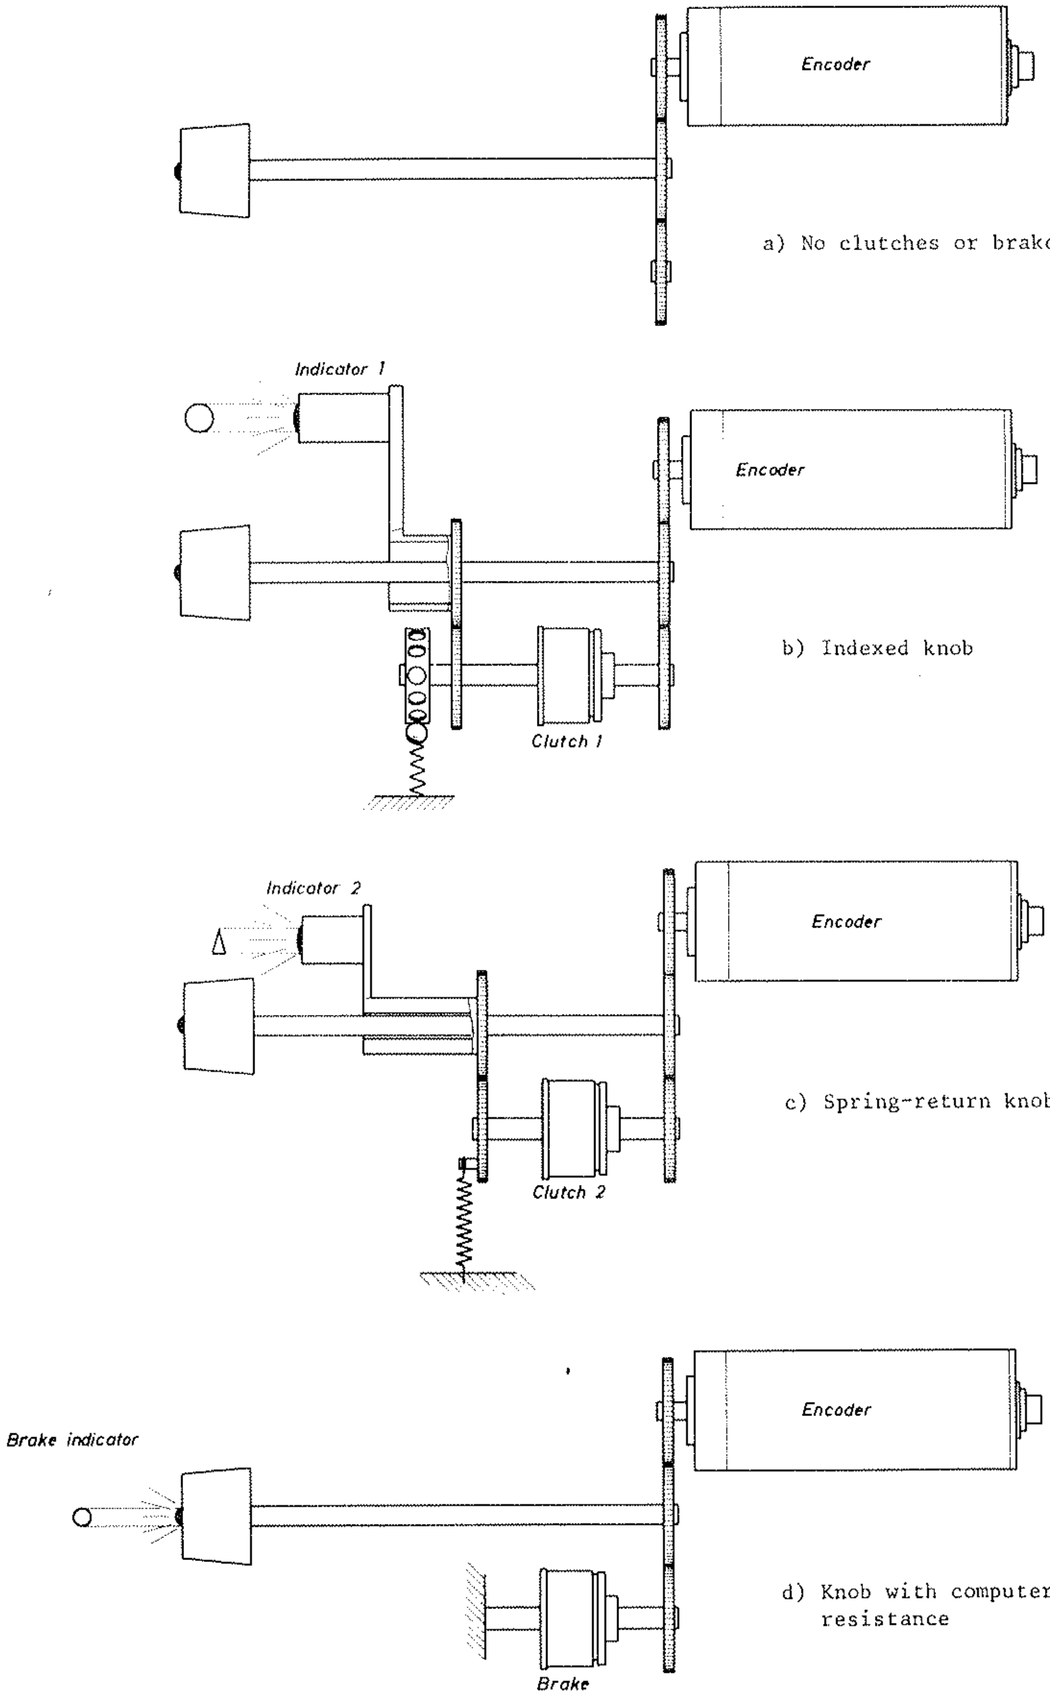

a) No clutches or brakes energized

b) Indexed knob

c) Spring-return knob

d) Knob with computer-controlled resistance

Fig. V.3

In the third mode, the brake can be used to give the impression of end-stops. The command SKNOB (a,b) allows the knob to be rotated only between the limits a and b. Attempts to exceed them cause the knob to lock, but a turn back into the permitted range (allowed because there is a small degree of backlash in the gears) releases the knob. In practice the knob feels as though it is moving between hard metal stops.

Read-out from the knob is done by the interpretive command:

SET X = KNOB

The knob counter-register can be preset before the operator turns it by the command:

SET KNOB = Y

This capability gives the relative facility mentioned at the beginning of the discussion on decoders.

2.2.3 Conclusion

The knob works and is beginning to be considered by applications programmers as a useful day-to-day tool for input of information to the computer, and thence to the control system.

2.3 The track-ball

To steer the cursor around the screen a manual steering device is needed. If two-dimensional control is needed, we can use two knobs, a joystick, or a track-ball. The track-ball was chosen arbitrarily as being convenient and pleasant to use. It is a continuous incremental device for the control of position (unlike the joystick, which is usually used for the control of velocity).

The ball is embedded in the table top. Two encoders are made to roll on its underside. Each encoder is interfaced to the computer separately, and in the same way as those for the knob. Thus relative measurements of the X and Y components of the operator's hand-guiding motion are available. While the ball may be used as an input device for variables, it is most useful as a steering device for the cursor, in which the relative components of cursor motion can be directly read into the computer.

The figure below shows the ball arrangement:

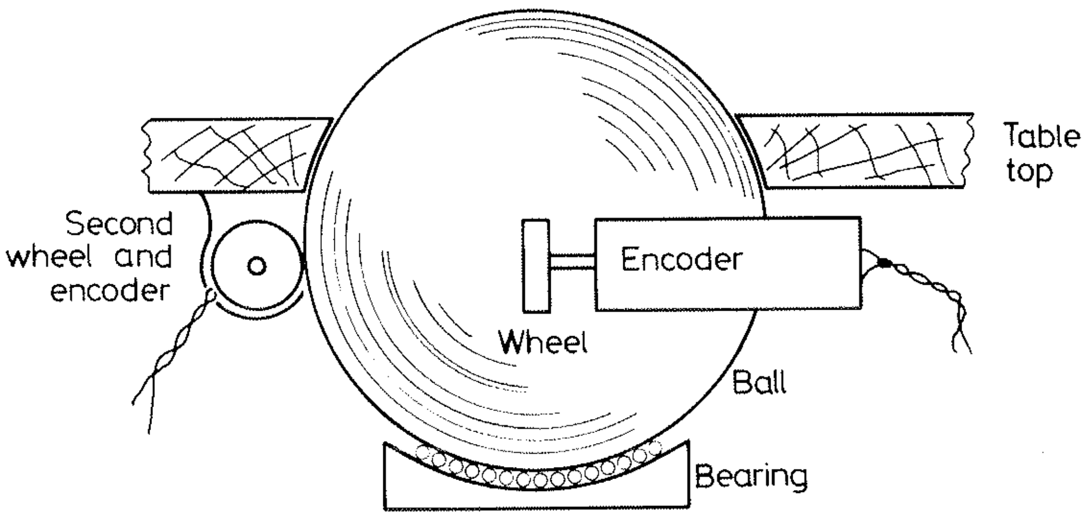

A rather pleasant and efficient ball of this kind, about 10 cm in diameter, was made at CERN and is installed in the consoles. No great superiority is claimed over commercial models, but the ability to use the standard encoders and interfaces makes the ball effectively appear to the computer as two knobs, and this is a great convenience for software. The ball is bigger and more pleasant to use than many others. An interrupt button is supplied beside it and is useful in many applications where the operator wishes to indicate the position of a point on a screen to the computer.

3. THE WORD-STORE

3.1 Introduction

Touch-buttons are a comprehensive and useful control device. The ability to launch an operation or make a choice by simply pushing a button covers all control requirements except those involving control with visual feedback, which in general are best performed with knobs and balls.

The touch-button scheme has one basic drawback in a situation which is constantly changing. It is not trivial, even in an interactive interpretive program, to add new elements to the touch-button tree. The logical structure of the tree is fixed from day to day, and the variety of control available is drawn from the actual operation of the buttons.

It was suggested in the early days of the project that if any flexibility was required beyond that available with touch-buttons, it could be obtained by directly typing in statements in NODAL. This has indeed proved to be the case. A useful program written and tested on-line can be stored in a file, and the statement, e.g.

RUN MYFILE

causes this program to be run every time it is needed. The word-store described here reduces even this simple operation to the push of a button, as well as facilitating various operations occurring in program development.

3.2 Description of the word-store

The hardware of the word-store consists of a plasma display panel allowing the representation of 8 rows of 32 characters. A row of eight buttons identifies the individual rows of characters. Two additional buttons complete the hardware, which is interfaced with the computer in the usual way through CAMAC. Thus any row of characters can be written from the computer and any one of the 10 buttons can cause an identifiable interrupt:

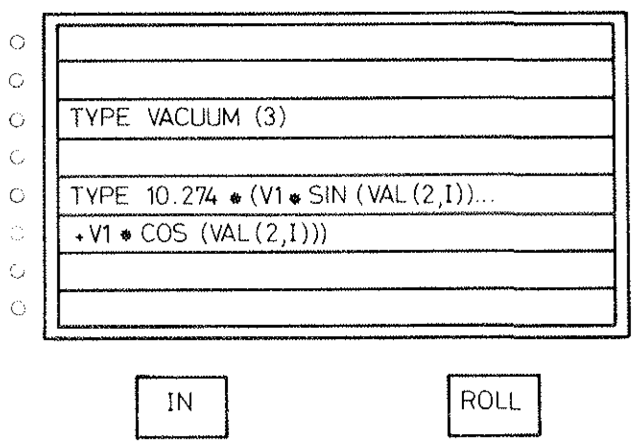

In use, messages displayed on the word-store can be emitted by pushing the appropriate identifying button, provided the word-store is in its standard condition. Thus if, as in the diagram above, the third row of the word-store has the legend "TYPE VACUUM (3)" the depression of the button to the left of this legend causes the NODAL statement TYPE VACUUM (3), followed by a new-line character, to be input to the interpreter just as though this string of characters had been typed. This has the effect of replacing the typing of a string by the pushing of a button. Since in practice the typing of the string may require some thought in its preparation, the ability to do it once only and to repeat it at will without further thought can be useful. Clearly up to eight statements can be emitted at the push of a button at any time.

3.3 How to put information into the word-store

The IN button is depressed. This button lights up, showing that the word-store is in a non-standard condition. Whilst the IN button is illuminated, words displayed on the word-store cannot be sent to the computer.

The operator now decides which row of the word-store is to be replaced from the typewriter keyboard. He pushes the appropriate identifying button and the old message disappears from view. (This old message is not permanently lost, and can subsequently be retrieved by the use of the ROLL button, which will be described below.)

Simultaneously with the disappearance of a row of text, a flashing cursor (part of the standard plasma display hardware) appears in the first character position of the row of blanks produced. Depression of subsequent buttons causes other rows to disappear, the cursor always moving to the beginning of the row most recently selected.

At this point, anything typed on the keyboard will be sent into the empty row with the identifying cursor, and will appear visibly on the plasma panel. The IN light goes out, showing that the word-store is now in standard condition. The transfer is not performed until the receipt of the "new-line" character, and the message to be put into the word-store can therefore be edited and refined on the on-line screen by use of the normal editing facilities available in NODAL.

The "new line" character is transferred to the word-store with the rest of the line, and though it is invisible in the display, it will be emitted with the rest of the line whenever the identifying button is pushed. The storage of the "new line" character can be avoided if necessary (for instance when the text represents only the first half of a program line), by terminating the type-in with three consecutive dots prior to the new-line symbol. The word-store program recognizes three consecutive dots as a valid terminator, and stores them visibly at the end of a message, but omits to send them to the computer when the identifying button is pushed. Rows 5 and 6 in the diagram above show how, by pushing two buttons, a long message may be simulated, with a single "new line" symbol at the end of the entire message.

The light under the IN button is extinguished whenever typing is complete. The word-store is then in its standard condition and ready to emit its messages.

Messages which are caused to disappear when a row is selected are not lost. The previous eight rows rejected in this way can be retrieved successively by the use of the ROLL button. Each depression of the ROLL button causes a pointer to move cyclically round the stack into which discarded messages have been pushed. The message opposite the pointer is displayed. Since the stack is only eight words deep (an arbitrary restriction imposed by storage limitations), the retrievability of discarded messages is limited. The ninth-oldest message is lost. Nevertheless, the availability of a stack enables the word-store to hold 16 messages, 8 of which are visible and instantly available.

When the IN button is unlit, messages can be issued interchangeably by the keyboard and the word-store, which have exactly equivalent status:

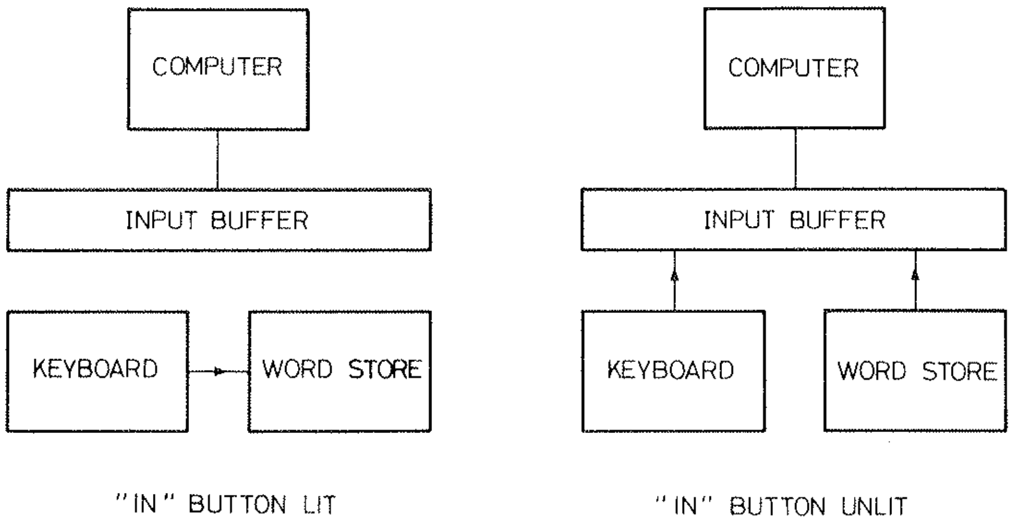

Since the messages in both cases go through the input buffer of the software teletype handler, mixtures of messages can be issued. Thus data can be added from the keyboard while standard parts of the message are issued from the word-store. Messages can be taken out of the word-store, edited from the keyboard, and then replaced in the word-store. To facilitate this operation, a % sign can be typed on the keyboard, which in NODAL convention causes all material following it to be treated as commentary and ignored. In this way material can be moved from the word-store into the teletype buffer for editing without being executed.

4. THE ELECTRONIC THUMBWHEEL

4.1 Introduction

Thumbwheels are used for setting up parameters in a number of situations in conventional control rooms. Such a thumbwheel normally has a series of indicators, each showing a digit and each having a serrated wheel for setting. In addition, a zero-set button is also usually supplied:

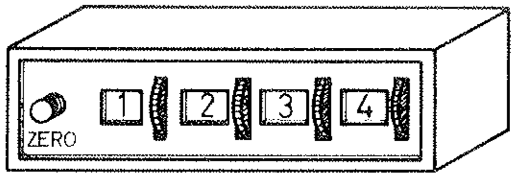

Thumbwheels of this kind are nowadays supplied with integrated circuitry allowing direct read-out in BCD code, and therefore suitable for direct computer input. When a delay or other parameter is to be set to a value which is known numerically, this kind of device is ideal. Setting up is simple, and a visual read-out of the value currently set is constantly at hand.

In many situations, however, the conventional thumbwheel has a number of drawbacks. It is not quite so easy to increment the set value by a given number of counts, as there may be a carry to perform manually. The addition or subtraction of a specific percentage requires a calculation followed by a re-set. Finally, in situations where incrementation is to be continuous while a display is watched -- quickly while the output is very wrong and slowly when the desired output is approached -- the traditional thumbwheel is totally unsuitable.

4.2 The electronic thumbwheel

The electronic thumbwheel looks similar to a conventional mechanical one, but is in fact quite different and has none of the disadvantages of the latter. The number-wheel is replaced by a light-emitting diode display, and the actual thumbwheels by rocker switches. The zero-set is omitted:

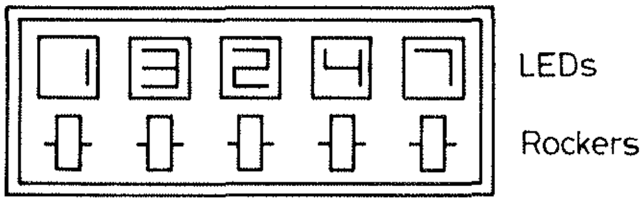

4.3 Operation

In operation, each rocker-switch has the ability to increment or decrement its own decade.

An instantaneous push on the top of any rocker-switch causes the value of the digit above it to be incremented by 1. Holding the top of the switch down causes continuous incrementation at a speed of about five counts a second. The bottom of each rocker-switch causes decrementation in the same way. All incrementation and decrementation is with automatic carry. Thus the number shown on the indicators can easily be set to any arbitrary value, and can also be incremented or decremented by units, tens, hundreds, etc.

The display reads out continuously from a BCD counter, which is incremented or decremented by an oscillator controlled by the switches. This counter is interfaced in both directions to the computer, so that it can be both read out of and written into. The counter goes from zero to 16383, and negative values are forbidden. Any attempt to set values below zero or above 16383 results in a zeroing of the entire register, which explains why no special zero button is necessary. A moment's push on the most significant rocker-switch sets the indicated number to zero.

Not only can the BCD register be incremented by unity at a rate of five times a second, but it can be incremented with a range of higher frequencies by a frequency-controlled oscillator controlled by a word used for this purpose. This incrementation register can be loaded from the computer, and the number in the main display register is growing or shrinking continuously all the time this register has a non-zero content. Thus control of the register contents can be obtained by a knob used as a velocity control, without introducing excessive computer traffic.

Finally, the read-out of the register is done through a BCD-binary encoder, so that no computer conversion of the reading is necessary.

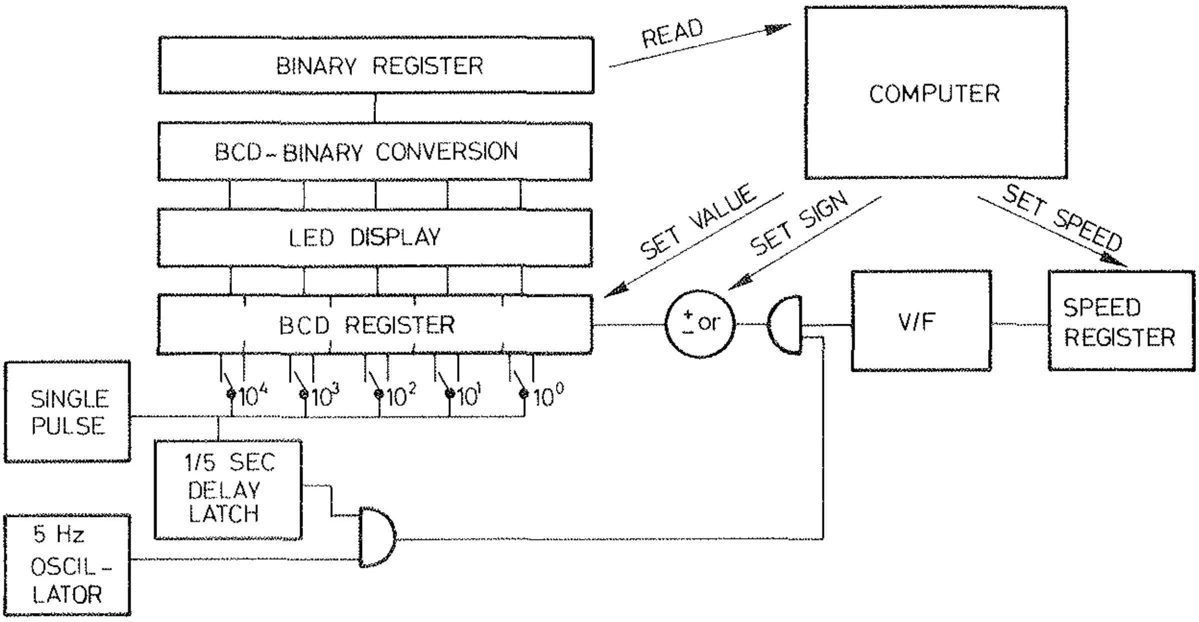

4.4 Conclusion

The electronic thumbwheel has been built and used. It is currently being employed on the console for the dedicated purpose of setting triggering delays for oscilloscope time-bases. Many of the advanced features, such as the auto-incrementation described above, are not yet in use.

5. THE ALLOCATION SYSTEM FOR OPERATOR SCREENS, AND THE MEMORIES ASSOCIATED WITH THEM

5.1 Screens

Throughout the entire design process of the central control system, the principle was adopted that the operator should not be burdened by having to understand the precise structure of the graphics system. Thus rather than numbering the TV screens, it is preferable to equip each one with an identifying button. The computer can then ask the question "Which screen?" and the operator's answer is to push the identifying button, answering, in effect, "That one". Even in the case where the screen in question is out of reach (as in the case of the wall display, or a screen on another console), it is still possible to use a dedicated touch-button marked "wall display" or push a touch-button marked "console 2", followed by a push on the equivalent screen on one's own console.

Thus in practice, it is quite possible to make a program to clear, move, or copy any picture within the system, simply by asking, by touch-button selection, for the appropriate operation, followed by parameters given from the identifying buttons on screens.

The process for moving a picture from one screen to another, for example, involves calling a touch-button page called "video-switching" which looks like this:

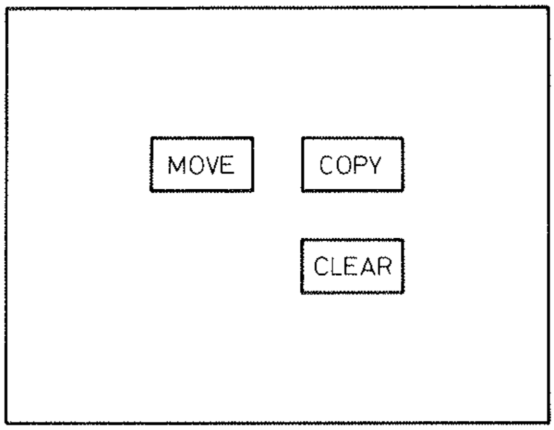

On pushing the button marked "move", a message appears which says "Which screen is source?", and on identifying the screen carrying the picture to be moved the operator is asked "What destination?". When the destination is identified by pushing its button, the video-switch program automatically connects to the destination screen all signals which originally went to the source screen. Since this is a "move" not a "copy", this operation is followed by a disconnection of the signals originally connected to the source screen. In the case of a "copy", the disconnect phase is omitted, while of course in the case of a "clear" only the disconnect phase is performed.

To perform these operations in every case, it is necessary for the program to have at its disposal the complete state of all the current connections in the video-switch. This is in fact available, as the software controlling each physical switch carries the updated state of that switch at all times.

5.2 Memories

The memories, as was mentioned in an earlier chapter, are a shared resource. In this way they can be compared to tracks on a computer disk, which are all alike, and are allocated by a program requiring them and returned to a pool when they are released or when the reserving program is aborted. The user program simply reserves a dot memory, or character memory, or a cursor (or all three), and writes into them with normal graphics statements. These memories are subsequently released by the picture-clear mechanism described above.

5.3 Modifications to original scheme

It was originally intended to include the colour memories in the switching scheme, but this was partially abandoned. In the most recent version of the system, a colour memory is permanently attached to each colour screen, saving considerable complications for the commonest applications. It is still possible to add dot-pictures to the individual colour channels, and the picture can be exported to black-and-white screens by the method described above.

Another place where the original scheme has had to be modified for simplicity of programming is the allocation of memories during programming. The original idea was for the program to ask "Where would you like the picture to appear?", and for the operator to reply by pointing to a screen. This was abandoned in favour of an arrangement where the screens are numbered (from 1 to 11 on each console), and a screen-assign statement in the program follows the statement(s) reserving the picture memory. In this way a program can run and produce a picture on a preselected screen without operator intervention. The picture can always be moved subsequently by the "move" or "copy" procedure described above, and the software is clever enough to disconnect all screens when it gets an instruction to "clear" a picture, even when what is identified is not the original or parent picture.

5.4 Conclusion

This picture-allocation arrangement is not ideal, but it gives access to all the equipment at all the desks. The hard-copy process will in due course be added to the facilities in the video-switching program. This is the obvious place for it to be introduced.

Some operations are not possible; for instance, moving the colour picture, which is not feasible since one cannot disconnect the hard-wired lines to its colour memory (copying of the picture is possible, because it does not involve breaking a connection). When an impossible operation is demanded, the program has been arranged to ignore the instruction. Knowledgeable operators soon realize their mistake, and the ones who do not understand the limitations of the system simply give up. This arrangement seems to be quite satisfactory.

CHAPTER VI

SOME COMPLETE INTERACTIVE SYSTEMS

Whilst it is not intended at this point to start a complete discussion on how an accelerator will be controlled with the equipment described in the previous chapters, the intention here is to give brief outlines of some interactive processes which are typical of those that will control and monitor the various subsystems making up the accelerator. These processes are normally programmed, in the NODAL interpretive language, by the experts in charge of the various subsystems, but it was necessary to provide some specimen programs to show the capabilities of an interactive system of this kind. The first example was made by the author, the second with a collaborator, the third with the author's advice and help, and the last two entirely by independent workers who had mastered the system. It is thought that there will be dozens or hundreds of such programs in due course, but it is difficult to conceive of principles not included in these few first ones.

6.1 Control and monitoring of variables

The touch-button tree is used to choose a page of variables for control. The names of these variables are displayed in the left-hand column, and each variable is followed by information about it, in columns representing "nominal value", actual "read-back value" obtained from reading the accelerator hardware, and a "setting value".

A cursor, in the form of a horizontal line, is under the control of the track-ball. By rolling the cursor up and down, a variable can be chosen, and pushing the button next to the track-ball selects this variable. The program responds by changing the colour of this row of writing to show it has been selected.

At this point the knob becomes available. The light on it comes on to show it is alive (and also to show its mode) and turning it causes the "setting value" to change. The setting can now be controlled to have any desired value, and the operator then pushes the button next to the knob.

The "setting value" now changes colour, to show that it is being sent to the hardware, and that a read-back from the equipment is being done. When enough time has been allowed for the setting operation to be performed in a remote computer, a read-back is requested. If the read-back is sufficiently close to the desired value, the colours in the row go back to normal and the ball-cursor operation can be repeated to select a new variable for control.

If for some reason the setting value cannot be achieved by the hardware, the read-back figure will change to a distinctive colour. The same occurs if the setting was never even attempted, having been observed by the remote computer to be out of the permitted range.

This rather simple program could be the basis of most control and interaction, but more specialized systems are being introduced by some users, in which the information, instead of being in rows, is distributed over a mimic diagram. A two-dimensional cursor (crossed lines) is used, but the principle is otherwise the same.

6.2 The function editor

First version

The voltage waveforms of an accelerator cycle, which are used for powering correcting elements around the ring and for other purposes, are generated by electronic hardware from a series of time vectors, each vector giving a time period in milliseconds and a voltage slope in volts per millisecond. A series of 64 of these number-pairs defines the voltage at every moment of the accelerator cycle. The 64 straight lines can be made to approximate to any chosen continuous curve. The hardware is limited as to the time duration of any particular vector or the maximum slope attainable, and the software computing the vectors is expected not to produce vectors which are impossible to realize.

The function editor is a program which allows an operator to construct an arbitrary function of this kind, and subsequently to edit details of it. It is also to perform checks and refuse to send the edited data to the hardware if they fail.

Two initial design decisions were made for this program: one was to compensate for the limited quality of a scan-converter picture by allowing the operator to "zoom" any chosen part of the picture to arbitrary size, and the other was to avoid the choice between a large number of mathematical functions from buttons by allowing such functions to be typed in from a keyboard in interpretive language.

The function is constrained to be defined by up to 64 pairs of numbers in two arrays called X(I) and Y(I), where I has values from 1 to N ($N \leq 64$).

The function is displayed as a graph on a black-and-white screen, and a cross-shaped cursor under the control of a track-ball is superimposed.

A screenful of touch-buttons with legends like this is available to the operator:

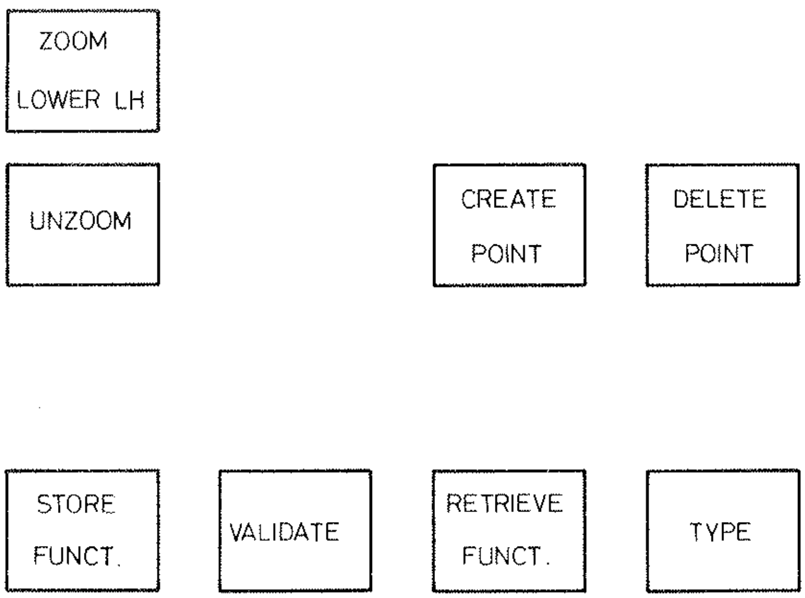

The function of "zooming" on part of the picture is performed by pointing at the bottom left-hand corner of the area to be zoomed and touching the "zoom" button, when the legends on the screen disappear and a single button marked:

ZOOM UPPER RH

appears. At the same time a pair of lines is drawn where the cursor was, and the cursor is free to be moved by a ball to show the upper right-hand corner of the area to be enlarged. When the right point has been chosen, the single touch-button is chosen, and the touch-screen returns to the original condition. The picture is now shown in the "zoomed" condition, and a cursor on it allows further zooming to be done again if desired. The axes are re-drawn and labelled with the correct coordinates for the enlarged portion of the picture. The following three drawings show the sequence for a typical zooming operation (the cursor is drawn in lightly to distinguish it from the curve).

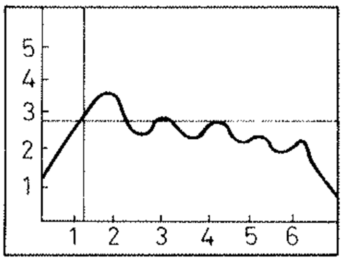

Choose lower LH corner

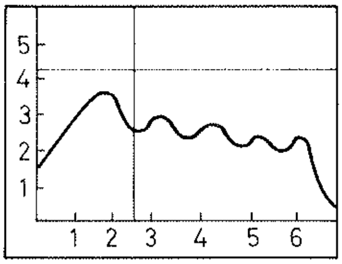

Choose upper RH corner

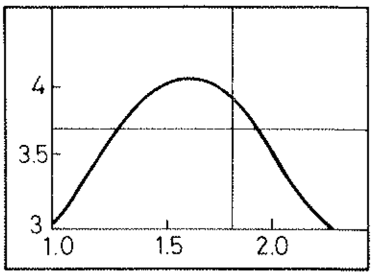

Zoomed picture
(Can be zoomed again)

The "unzoom" button restores the original picture, however many zooming operations have been nested. (It would be very attractive to restore the previous picture each time the unzoom button was pressed, and this would indeed be easy to implement, but the 128 words of computer storage needed at each stage seemed a trifle extravagant for this feature.)

For the "create point" and "delete point" features, only the vertical line of the cursor is used. The pushing of the appropriate button causes the nearest point to the current horizontal cursor position to have a new point interpolated linearly to its right or an existing point removed. The curve is immediately re-drawn after this operation.

All the time a curve is shown, the coordinates of the point on the curve which is on the vertical cursor are shown on a plasma panel, as is its ordinal number: for instance

14 2.746 3.124

means that the cursor is currently pointing at point number 14, whose coordinates are X = 2.746, Y = 3.124. This allows the curve to be investigated with a far greater precision than inspection of the graph would allow.

All changes to the curve save for the insertion and deletion of points are implemented by means of the "type" button. On touching this, control is passed to the keyboard, and a statement may be typed in in the interpretive language. For instance,

FOR I = 10, 20: SET Y(I) = Y(I)*1.01

will raise the part of the graph between the 10th and 20th points by 1%, while

SET X(17) = X(17) + 0.03

will move the 17th point slightly to the right. In this way arbitrary control can be obtained over any point or points of the curve, including any arbitrary algebraic operation. These expressions may also include the variable N, being the number of point-pairs in the curve, and the index NX, which is the ordinal number of the point nearest the cursor position when the "type" button was touched.

Thus, for example, the

SET Y(NX) = Y(NX)*1.01

will increase the point under the cursor by 1% every time it is performed, and this operation can be repeated at the touch of a button by putting it (or any other statement like it) into the word-store.

The "validate" button gives control to a function which tests the whole function for realizability, and the "store" and "retrieve" buttons lead to touch-button trees which move the completed data to and from the individual computers which are to use them for generating the necessary waveforms.

Second version

An improvement to this program for a particular purpose was devised by S. Van der Meer, who needed a fixed number of X-values and did not require complicated functional relationships. He allows both the X and Y cursors to control the "create" function, and produces an arbitrary point at any fixed X-value by simply pointing at it. The original curve is drawn in green, and the changed part of the curve is immediately overprinted in red (two scan-converters are used). Thus two curves are seen at any moment. The original in green, and the modified one in red. At the touch of a "set" button the new curve changes to green and the red curve disappears.

The requirement for a fixed number of points makes the "delete" operation unnecessary in this case, and the precision is attained simply by the ability to zoom arbitrarily.

This version of the program is far less powerful but is adequate for its purpose (editing the starting functions for the main magnet power supplies, which are in any case subsequently re-calculated at each cycle by closed-loop control in a local computer). It has the advantage over the original, more flexible editor of being simpler and easier to use, and this description of it was included to show how a user could modify the ideas to suit a specialized purpose.

6.3 The alarm system

The alarm system uses a computer with a mass memory (drum store) to absorb coded alarm messages and present them in plain text on an operator screen. This problem is complicated by the fact that at certain times the number of alarm messages may be very high.

Attempts will be made in due course to reduce the need to display a very large number of alarm messages, but for some time into the foreseeable future the list of alarm messages is expected at any time to be longer than can be presented on a single display screen.

The alarm system therefore consists not only of a colour screen, but also of an ancillary black-and-white screen and a set of touch-buttons for manipulating the large amount of information which is hidden on the drum-store and making chosen parts of it appear on the screens for examination.

Since the alarm system must be always available, none of this hardware is shared with any part of the graphics system. Both screens and the touch-button device are dedicated peripherals of the Alarm computer, and indeed the touch-buttons have so few different legends that the display under them has been dispensed with. The 16 legends are permanently affixed, and a series of 16 lights under them is substituted for the cathode-ray tube, each light being switched on by computer control when the appropriate button is in use by the program.

In use, the Alarm computer shows a list of alarm messages on the large colour screen. The newest message is at the bottom of the screen, and the messages roll up the screen as newer ones arrive.

At the touch of an "examine" button, this "new alarm" display moves to the small black-and-white screen, and examination of those alarms which have rolled off the top of the screen can proceed on the large colour screen. This arrangement ensures that, even when hidden parts of the alarm list are being examined by the operator, new alarms are still visible as they arrive at the bottom of the black-and-white screen.

Fresh alarm messages are presented in black writing on a red background, and are therefore very visible. The operator can acknowledge any message, to show that he has noticed it, and such an acknowledged message changes to red writing on a dark background, which reduces its visibility and makes the unacknowledged messages show up brightly by comparison. Alarms can also be cancelled, and both operations are performed by specialized touch-buttons, which are alive during the "examine" phase of the program.

Since the alarm system, for simplicity and cheapness, has no cursor and ball, the message to be operated upon by these "cancel" and "acknowledge" instructions, as well as by "roll up" and "roll down" instructions used for examining hidden parts of the list, is always the line of writing in the middle of the screen. A permanent cursor (in the form of a star of distinctive colour) is drawn to mark this line, and to cancel any particular line of writing the "roll up" or "roll down" button is touched repeatedly until that line is in the middle of the screen. Touching the "cancel" or "acknowledge" button at this point performs the appropriate operation on the chosen message. Extra buttons marked "move to top" and "move to bottom" are supplied to roll the information by half a screen for convenience, and these, together with an "automatic operation" button, changing back to the mode in which the latest message is on the colour screen, form the full complement of alarm controls.

The rest of the alarm system, which changes the colours of displayed messages automatically when new alarms come in showing that previously defective conditions have been repaired, and which cancels messages which have been seen by the operator and are now obsolete, is very elaborate, especially in view of the fact that the information is drum-based and only the displayed information is in the core of the computer at any time. However, the simple display scheme described is all the operator sees of the alarm system.

6.4 The observation of analog waveforms

Analog signals for the main control room arrive in a distribution rack in the central computer room. These signals come, in the main, along the 100 kHz twisted pairs in special cables (as used in cable television and in the data links for the message-handling system). The signals may originate from the six auxiliary buildings or from the injection and extraction systems and their respective beam lines.

Considerable multiplexing may already have been done on these signals before they reach the control centre, but in the control room there remains the problem of how to present these signals on oscilloscope screen for observation and comparison.

6.4.1 *The instruments*

Each console has two oscilloscopes, which are absolutely standard laboratory instruments of the modular plug-in type. Each oscilloscope has three plug-ins; a time-base and two chopping amplifiers. Thus each oscilloscope could run from a different time-base and display up to four waveforms with the full bandwidth of the cables bringing the information. The oscilloscopes have the ability to store waveforms as pictures of variable persistence though the storage feature can be switched off.

In practice, only three of the four possible waveforms are shown on each screen, to avoid presenting too confused a picture.

6.4.2 *Information routing*

All analog signals which arrive in the control centre are available on a "cross-bar" switch driven by the display computer through its CAMAC and multiplex interfaces. The components of this switch are identical with those used for the TV picture-switch mentioned earlier. Thus by issuing suitable computer instructions, any signal can be connected to any trace on any oscilloscope. To economize on switching equipment, a number of restrictions are imposed, such as some signals being available on only two desks, and no signal being available more than once on the same display.

6.4.3 *Signal labelling*

A plasma panel dedicated to analog signals is supplied on each console, and this is mounted below the oscilloscopes. The panel can display up to eight lines of alphanumeric information, and the top six of these are used to show the operator the name of the signal currently connected to a particular oscilloscope trace, as well as calibration factors and units for this signal. Each line of text has an associated "identify" button next to it. The identify button is used to specify its signal as the subject of a connect, calibrate or disconnect operation.

6.4.4 Signal selection

Signal selection is performed by a dedicated touch-screen. Control is given to this touch-screen from the main operator position, and the console is then temporarily available for signal selection. The analog switching touch-screen is programmed as a hierarchical "tree" in the usual way, meaning that choices are offered successively of subsystem, location, and signal, with a separate hard-wired backspace button so that the operator can flick through a series of related signals without following the entire hierarchy each time. The appropriate assign button is then used to connect the chosen signal to an oscilloscope trace and write its name and parameters on the appropriate line of the plasma panel. For economy, certain groups of signals available from auxiliary buildings share cables. A demand for a signal, when another is using the cable needed for conveying it, cannot be met successfully. The operator must either release the line by disconnecting the offending signal or abandon his attempt, and both possibilities must be allowed for.

The method which has been adopted is to supply a "release" button, whose depression, followed by the touch of an "identify" button, releases a line and its associated signal and text. This release procedure is performed by an interrupt program of fairly high priority, so that the releasing of a line can be done without disturbing other work.

Any signal on the oscilloscope screen can be temporarily replaced by a calibration signal (square wave, 1 V peak to peak) from the same geographical source. This operation is performed by the program when a "calibrate" button is pushed, followed by the appropriate "identify". The name of the signal is temporarily replaced by the word CALIBRATE. The *status quo* is restored by pushing the "identify" button once more, when the signal and its name reappear.

The calibrate and release buttons are the "identify" buttons of the bottom two rows of the plasma panel, there being eight rows available, only six of which are needed for signals.

6.4.5 Synchronization

The triggering line of each oscilloscope is hard-wired to a programmed timing unit which gives an accelerator datum time delayed by a programmable delay. The datum line is the injection time of the accelerator beam, and the delay is set by an electronic thumbwheel (described in Chapter V, Section 4). A real-time program receives an interrupt every time the thumbwheel moves and updates the delay in the timing unit. There are two electronic thumbwheels on each desk: one for the triggering of each oscilloscope. The thumbwheel gives the number of milliseconds between injection and the triggering of its oscilloscope.

6.4.6 Other adjustments

The controls on the plug-ins of the oscilloscope are all available to the operator. He can adjust any of them, but must leave (or replace) the Y-amplifier gain to a standard position if the calibration factor, written on the plasma panel next to the signal name, is to have any meaning.

6.4.7 Conclusion

The switching system described gives direct viewing of most analog signals in the control room without a proliferation of patch-boards and oscilloscopes for dedicated purposes. The system described does not present signals of high frequency, and for this purpose specially connected high-speed instruments are still necessary. However the number of such signals is very small (mainly for a wide-band pickup in the beam and for the RF system). Other high-speed signals are converted to TV signals locally and then piped in via the video-switch, again avoiding the need for special display equipment in the control room.

Figure VI.1 shows the panel layout of the analog observation system.

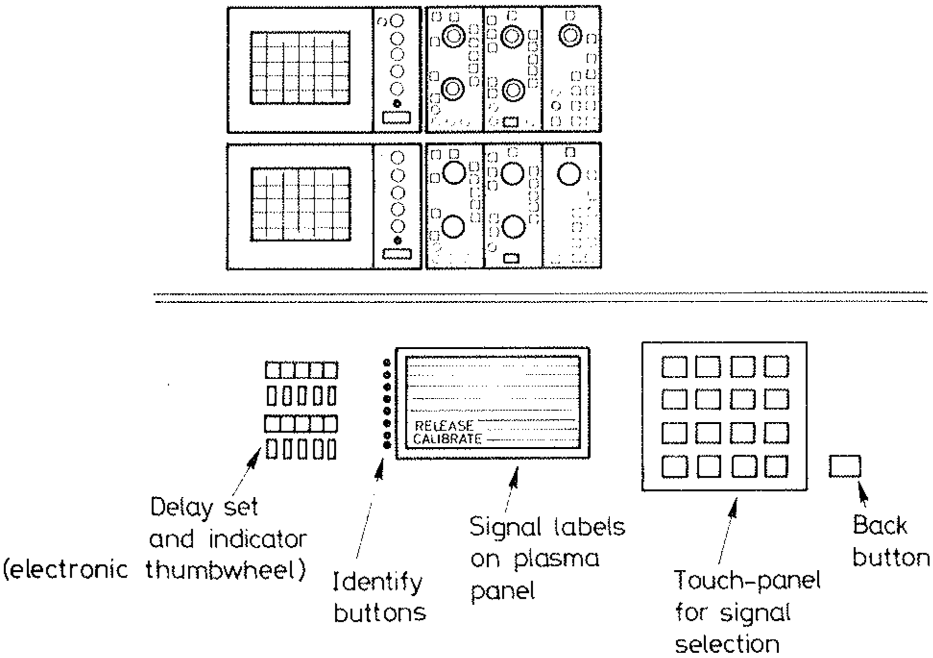

Fig. VI.1 Analog observation -- panel hardware

CHAPTER VII

SOME SUGGESTIONS FOR FUTURE DEVELOPMENTS

1. GRAPHICS

The speed of drawing on the scan-converters is limited to the speed of computation of the individual points of a graph by the vector generator. Very much faster vectors can be drawn by directly applying analog voltages to the input of the scan converter. This is done by generating Lissajous figures with suitable sine waves, and modulating the brightness inversely from the traversal speed of the beam so that all vectors will have the same brightness, independent of their length. Thus the relations:

$$X = 1 \cos \theta \cos (\omega t)$$

$$Y = 1 \sin \theta \cos (\omega t)$$

$$I = 1 \sin (\omega t)$$

give the coordinates and brightness of a vector from the origin at an angle θ at any time, and suitable d.c. shifts of the voltage allow such a vector to be drawn anywhere on the screen. Similarly, the relations

$$X = R \cos \omega t$$

$$Y = R \sin \omega t$$

$$I = R$$

allow a circle to be drawn in the same time as a vector.

Since drawing is only to be done during the TV flyback time, a test has to be done to ensure that the entire vector or circle can be drawn before picture-scanning recommences for the next frame. In practice this means that all drawing must cease during the last 5% of the 1 msec flyback time.

It appears from experiments already performed that very perfect circles and vectors can be produced at 20 kHz, and this represents a gain of about 100 on the process of producing vectors by digital means. In addition, a simple clocking arrangement enables the circle to be preferentially blanked, so that the generator can be made to produce not only circles but also individual octants, which are useful for drawing characters. For example, the number 6 can be drawn with a complete circle, a straight line, and two additional quadrants:

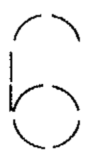

A character generator of this kind, written in software, is already used for the scan-converter memory, but is very slow. The proposed method would make it much more useful.

Even better quality graphics could be produced on a TV raster by a completely digital method. Semiconductor memories big enough to hold an entire dot-picture are rather expensive, but prices are dropping to the point where such systems are already being sold for reduced definitions. A complicating factor is that the drawing of vectors in random-access memory requires in general the re-calculation of each point in real-time, and this certainly requires a dedicated microcomputer. The speed of the analog method described above could never be approached, but the very high quality and stability of a digital picture might make the slowness and expense palatable. Certainly digital systems are available commercially, but none of them (except the data disk) are of a high enough definition to draw smooth curves with no apparent steps.

The character controller built for this project has proved to have several drawbacks: the lower-case alphabet, whose realization caused many design constraints, turns out to be difficult to program and not very much used. On the other hand, the four colours supplied are minimal, and an eight-colour system would be preferable. The graphics symbols supplied with the character generator are rather limited, and techniques are now available for replacing the programmable read-only memory which was used for these symbols by a read-mostly memory. With suitable software of considerable complexity such a system might come nearer to a useful composite graphics arrangement for graphics and characters. The size of the software might well be unacceptable and a microcomputer per memory might be needed here also.

2. CONTROL DEVICES

While the touch-buttons work very well, the mechanical buttons discussed in Chapter V.1 should probably be developed for use in other than central locations, as they use a simpler interface.

A new computer-controlled knob is currently being built, which instead of having spring and detent mechanisms can simulate these and many other operator feedback motions.

The new knob contains a silent, high-torque motor, and three clutches, which are driven by digital-to-analog converters and drive amplifiers. By energizing one or other of these clutches the knob may be driven or pushed in either direction, or connected to a brake. By suitable programming, any relation between force and angle, or even between force and rate of change of angle, can be obtained. The prototype for this device has provision for the possible insertion of a microcomputer to control such functions, should computer I/O loading be a problem.

In addition, the knob has two concentric circles in the front, each having 64 light-emitting diodes which are individually programmable. These can be illuminated selectively to give patterns like this:

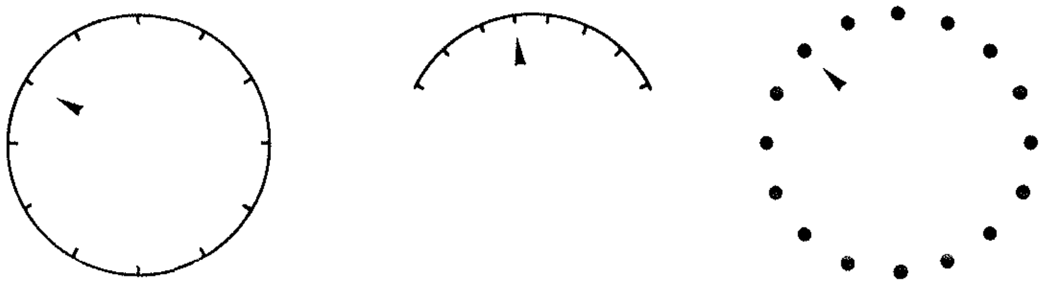

and thus show the currently available operating mode of the knob.

3. SYSTEMS

A new piece of hardware is available on the market (Ref. C12), which fits exactly into our system, thanks to its being based on TV norms. This is a microfiche reader, with TV read-out and automatic seeking of any page of information by computer control.

The microfiche reader contains some 750 microfiches (10 × 15 cm) each capable of holding 96 photographic copies of A4 pages of information. A TV image of any half-page is available within at most 5 seconds of giving its address, and this TV image can of course be seen at any console.

This device was, originally, conceived as an information retrieval system, but in the SPS control room it could be much more than this. Not only could it be linked to the alarm system, to give an appropriate page of information at any desk on request, describing faulty equipment in detail when its defectiveness is reported, but it could be used to give further information on subroutines which the user wishes to incorporate into on-line interactively written programs, such as functions and data-modules.

The most exciting application of a programmable microfiche reader of this kind, however, is in combination with the touch-buttons. By allowing the operator a number of choices after reading each displayed page, a logical tree of selections on the display screen may be combined with a logical tree of choices on the touch-button screen. In this way a true teaching program may be written for training accelerator operators. This is a purely organizational and software project whose potentialities in terms of the application of new ideas are very great.

Finally, there is what might be termed the ultimate integrated control device. This arrangement was not made for the SPS accelerator since the latter is a slow-cycling device with a period of 5-10 seconds and with no major use for a fast "hand and eye" feedback device. The integrated control device works as follows.

A gear-wheel made of clear plastic is held above and parallel to the face of a display tube, and supported at its periphery in such a way that it can rotate. Such rotation can be effected by means of a knob glued to the centre of its upper surface. The gear-wheel is coupled to another having the usual encoder on its spindle:

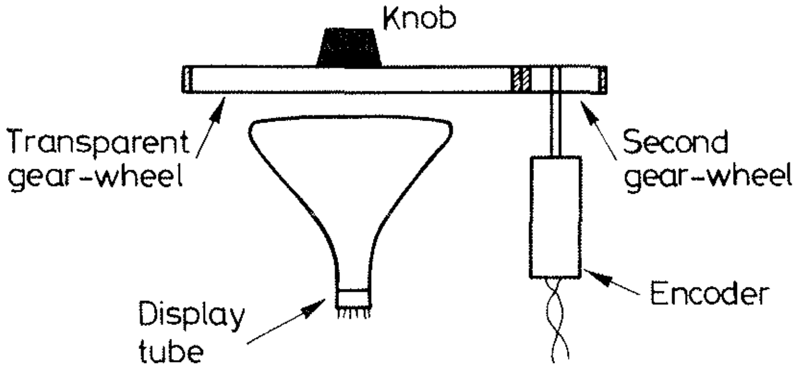

If the display is driven with suitable graphics hardware (including a radial cursor) a correctly calibrated scale can be drawn on it by the computer, and the various parameters of desired variable, set variable, actual reading, and permitted limits can all be displayed on this scale.

The replacement of the simple encoder shown by the computer-controlled knob described previously enables the device to be used also for the selection of variables and programs. This integrated control device (together with a single button) enables the whole gamut of other control devices, including touch-buttons, to be dispensed with. A simpler machine than a large accelerator could be controlled entirely with the integrated control device, by using it successively to choose programs and input parameters to them.

REFERENCES

1) The 300 GeV programme, CERN Report 1050, January 1972.

2) P. Adams, K9 computer software system, Report RHEL/R240 (1971).

3) B.L. Hartway, J. Bergstein and C.M. Plopper, The central control room man-machine interface at the Clinton P. Anderson Meson Physics Facility (LAMPF), Proc. Particle Accelerator Conf., San Francisco, 1973 (IEEE, New York, 1973), p. 633.

4) C.S. Taylor, Linac operation and controls, CERN MPS Linac Note 71-14 (1971).

5) L.A. Klaisner, P. Mellick, K. Rich, S.R. Smith and M. Storm, Central control system of the NAL accelerator, Proc. Particle Accelerator Conf., Chicago, 1971 (IEEE, New York, 1971), p. 424.

6) K.A. Fox, M.P. Pasturel and P.S. Showman, A human interface for automatic measurement systems, Hewlett-Packard Journal, Vol. 23, No. 8, April 1972.

7) M.C. Crowley-Milling, A digital control system designed for operator convenience, Report DNPL TM/80 (1971).

8) E.A. Gurd, E.A. Drawneck and B.S. Walker, A multiple console interactive graphics system, Proc. Internat. Conf. on Computer Systems and Technology, London, 1974 (Instn. Elect. Engrs. Conf. Publ. No. 121, London, 1974), p. 173.

9) J. Standeven, S. Mihailovic and D.B.G. Edwards, Multiconsole-computer display system utilizing TV techniques, Proc. Instn. Elect. Engrs. 115, 1375 (1968).

10) R.J. Scott-Kerr, Computer-aided control at Pembroke power station, Electronics and Power, March 1972, p. 87.

11) J.T. Hyman, The computer control system for the SPS, Proc. Internat. Conf. on High-Energy accelerators, Stanford, 1974 (USAEC, Washington, 1974), p. 503.

12) M.C. Crowley-Milling, J.T. Hyman and G.C. Shering, The NODAL system for the SPS, Report CERN Lab II-CO/74-2 (1974).

13) Handbook of human engineering data (US Naval Training Device Center, Port Washington, NY, 2nd ed. (rev.) 1952, reprinted 1960).

14) Human engineering guide to equipment design (McGraw-Hill, Inc., New York, 1963).

15) E. Edwards, Ergonomics in control, Proc. Instn. Elect. Engrs. 120, 1181 (1973).

16) J. Croney, Anthropometrics for designers (B.T. Batsford, Ltd., London, 1971).

17) Der "vermessene" Mensch, Anthropometrie in Kunst und Wissenschaft (Heinz Moos Verlag, Munich, 1973).

18) The Nordcom system, Norsk Data Elektronikk A.S. (see Ref. C1).

19) G. Menzies, H.A. O'Donnel, A.J. Perry and C.J. Robinson, Interactive video-graphic display system for use with cathode-ray tube scanning devices, J. Phys. E 8, 676 (1975).

20) A. Muser and J.R. Lutz, Visualisation polychrome par mosaïque, Nuclear Instrum. Methods 113, 131 (1973).

21) Private communication from Dr. C. Clutterbuck, Royal Radar Establishment, St. Andrews' Road, Great Malvern, Worcs., WR14 3PS, United Kingdom; and
Private communication from Dr. C.D. Reid, Atomic Weapons Research Establishment, Aldermaston, Berks., R67 4PR, United Kingdom.

22) C. Dean, Computer-controlled scoreboards at the Olympic Games, Electronics and Power, Aug. 1972, p. 322.

23) Anon., Liquid crystal materials and device developments, Electronic Engng., Aug. 1974, p. 30.

24) I. Gorog, Cathodochromic image displays, Appl. Optics 9, 2243 (1970).

25) P.T. Bolwijn and R.A. van Doorn, On the selective erasure of cathodochromic image displays, J. Phys. D 5, 896 (1972).

26) R.C. Duncan, B.W. Faughan and W. Phillips, Photochromics and cathodochromics, Appl. Optics 9, 2236 (1970).

27) Z.J. Kiss and W. Phillips, Cathodochromism in photochromic materials, Phys. Rev. 180, 924 (1969).

28) C. Christensen and E.N. Pinsen, Multifunction graphics for a large computer system, Proc. Fall Joint Computer Conf., Anaheim, Calif., 1967 (AFIPS, New York, 1967), p. 697.

29) C. Brewer and C.C.M. Parish, An experimental data display autocode (EDDA), Central Electricity Generating Board Report RD/L/R/1617 (1969).

30) E.M. Mott, A transparent touch-screen device for interactive graphics displays, Report RHEL/R248 (1972).

31) B.L. Hartway, A compact programmable control panel, Proc. Particle Accelerator Conf., San Francisco, 1973 (IEEE, New York, 1973), p. 531.

32) Proc. Internat. Conf. on Displays, Loughborough, 1971 (Instn. Elect. Engrs. Conf. Publ. No. 80, London, 1971).

33) F. Beck and B. Stumpe, Two devices for the operator interface of the new CERN accelerator, Report CERN 73-6 (1973).

COMMERCIAL REFERENCES

C 1 A/S Norsk Data Elektronikk,
Økernveien 145, Oslo 5, Norway.

C 2 Traitement Informatique Techniques Nouvelles S.A.,
1, rue Gustave-Eiffel, 91420 Morangis, France.

C 3 Christian Røvsing A/S,
46b Marielundvej, Herlev, Denmark.

C 4 Data Disk Inc.,
686 W. Mande Ave., Sunnyvale, Calif. 94086, USA.

C 5 Princeton Electronic Products, Inc.,
PO Box 101, New Brunswick, NJ 089902, USA.

C 6 Owens-Illinois, Inc.,
PO Box 1035, Toledo, Ohio 43666, USA.

C 7 Sony Corporation,
7-35 Kitashinagawa 6-Chome, Shinagawa-Ku, Tokyo 141, Japan.

C 8 Conrac Corp.,
600 N. Rimsdale Avenue, Covina, Calif. 91722, USA.

C 9 NV Cobar Barco Electronic,
Sevenslaan 8500, Kortrijk, Belgium.

C10 General Electric Co., Ltd.,
1 Stanhope Gate, London W1A 1EH, UK.

C11 Ferroperm Aktieseskab,
Stubbeled 7, Trørød 2950 Vedbaek, Denmark.

C12 Stabletron, Ltd.,
266 Hunts Pond, Fareham, Hampshire P014 4PH, UK.

APPENDIX I

TECHNICAL SPECIFICATION FOR THE CHARACTER GRAPHICS CONTROLLERS USED IN THE CONTROL CENTRE

I N D E X

1. INTRODUCTION

The control room of the 300 GeV proton-synchrotron being built on the Swiss-French border by CERN will have a unified display system based on 625-line raster pictures.

This system will use conventional RGB monitors which are outside the scope of this description. Displays will be generated by various methods, one of which is the subject of this invitation to tender. The units required are stand-alone, rack-mounted generators for three-colour graphics signals in TV raster form. The environment of the unit is shown in Figure 1.

2. SCOPE

This specification covers the requirements for design and manufacture of a colour character subsystem. This subsystem will accept individual 10-bit characters for display and control on a parallel interface. It will store characters in its own internal memory, and from this memory it will continuously regenerate three, 625-line, 50-frame-per-second signals, which will constitute the output of the unit.

The system will comprise

(a) Memory with refresh circuitry

(b) Clocks, timing signals and control logic

(c) Input and output interfaces

(d) Any necessary power supplies

Two types of controller are specified. Colour, and black/white. The two types are to be completely compatible, the black/white controller being a simplified version of the colour one.

Four colour controllers and eight black-and-white controllers are required. Additional units of either kind may be required subsequently.

3. SPECIFICATION OF THE COLOUR CONTROLLERS

The purpose of the unit is to represent the 800 characters currently in its memory continuously on a 625-line, 50-frame-per-second television raster, using a character generator in the form of a suitable read-only memory. Seven of the ten bits in each character will be used to specify the character itself, by means of a slightly modified version of the ASCII representation. Two additional bits are used to specify the colour (by diverting the output on to one or more selected output channels) and the final bit is used to reverse the role of background and character, i.e. to give black characters on a background of the chosen colour.

The unit allows chosen items in the memory to be replaced from an input line.

3.1 Memory

The unit supplied should have a memory capable of holding 800 characters of ten binary digits each. Read-write circuitry should be supplied, capable of reading out the entire contents of the memory every 20 milliseconds, and of replacing randomly-chosen individual characters at a rate of at least 1000 per second.

3.2 Character Generator

The character generator should handle 96 different characters. Fifty-six of these should comprise all the characters in the ASCII set, except for eight of the teletype control characters. Forty additional characters will be defined by CERN, which undertakes to supply the contractor with a precise list of these within one month of award of the contract.

The character generator will produce characters in a matrix with 10 dots on each of 12 consecutive TV scan lines. Each character from the ASC II set is to use a 7 x 8 box in the corner of this matrix, but the special characters may fill the matrix.

3.3 Control Circuitry

Sufficient control circuitry must be supplied to produce a raster signal from the information in the memory, testing the character generator for translation. The picture generated will display the 800 characters in the memory in 20 rows of 40 characters each.

Write circuitry must be supplied, to enable any character in the memory to be replaced, while read-out is taking place, with only temprorary loss of the output signal. The position in memory where writing is to take place is marked by a "cursor", which is a signal not stored in the memory, but added in a suitable form to the video output. The current time-position of the cursor is to be controllable both by relative and absolute commands in the character stream input to the device. Relative commands move the cursor one position in any of four directions. The absolute command allows the cursor to be re-positioned to an arbitrary location.

The clock driving the device must be capable of being synchronised from an external source, or alternatively, it should be possible to replace the clock function from such input. Synchronisation pulses for both frame and line must be generated for addition to the output.

3.4 Inputs

Input to the device is to be on a 10-bit parallel interface with a separate "busy" line. TTL levels are preferred. Input data will consist either of control signals for positioning the cursor, or of characters to be written into memory.

Additional inputs are a 75-ohm line for the external synchronisation mentioned in 3.3 above, and an input socket for the attachment of a local keyboard for test input.

3.5 Outputs

Five outputs are required from the device. Each should appear on a 75Ω BNC coaxial connector. Signals should be 1 volt, positive-going and synch, where included, should follow CCIR conventions.

3.5.1 "R" output containing red signal + cursor signal

3.5.2 "G" output containing green signal + cursor signal + composite synch.

3.5.3 "B" output containing blue signal + red signal

3.5.4 "W" output containing red, blue and green signals, as well as cursor and composite synch.

3.5.5 "S" output containing composite synch only.

3.6 Power Supplies

The controller must supply all its own power requirements from a nominal 220V 50Hz single phase supply and the chassis and working components must be isolated from this supply by a transformer. Due to other equipment on the same supply, the voltage could vary in the following manner :-

Total variations (> 5 seconds) ± 10%

cyclic variations (order of 5 seconds) ± 5%

fast transients (milliseconds) ± 10%

The frequency may vary by ± 1Hz and harmonics introduced by rectifier equipment on the same supply may have a RMS amplitude of up to 10% of the supply voltage.

The equipment must satisfy the safety requirements of the IEC recommendations.

3.7 Ambient Characteristics

The equipment must operate satisfactorily with an air temperature between +10° C and +35° C and a relative humidity up to 85%. The equipment must not be damaged by storage at a temperature of between -15° C and +50° C.

3.8 Mechanical Construction

The controller should be in a form suitable for mounting in a 19-inch rack, and should have an integral power-supply unit and any necessary cooling fans. The unit should have an on-off switch and indicator and fuse, all the input and output connections mentioned in the specifications above, and enough test-points to enable the correct functioning of the unit to be confirmed without dismantling it.

CERN prefers the use of its mechanical standards, and a choice of plugs and sockets from a standard range. Details of these standards are given in Appendix 1.

3.9 Documentation

Accurate and sufficient documentation must be supplied with the controllers. Two complete sets of documentation are required for each type of controller, and this must include circuit diagrams and anotated mechanical layouts of the circuit cards. The documentation must contain sufficient information for the maintenance and modification of equipment.

4. SPECIFICATION OF THE BLACK-AND-WHITE CONTROLLER

The black-and-white controller is a simplified version of the colour controller described above. It differs from it only in having less memory, and in producing only one output stream, which allows the omission of some switching circuitry. It should accept the same signals on the 10-bit input interface, and with identical effect, except that the two colour bits will be ignored. Of the five outputs mentioned in 3.5 above, the black-and-white controller will produce only the "W" output and the "S" output.

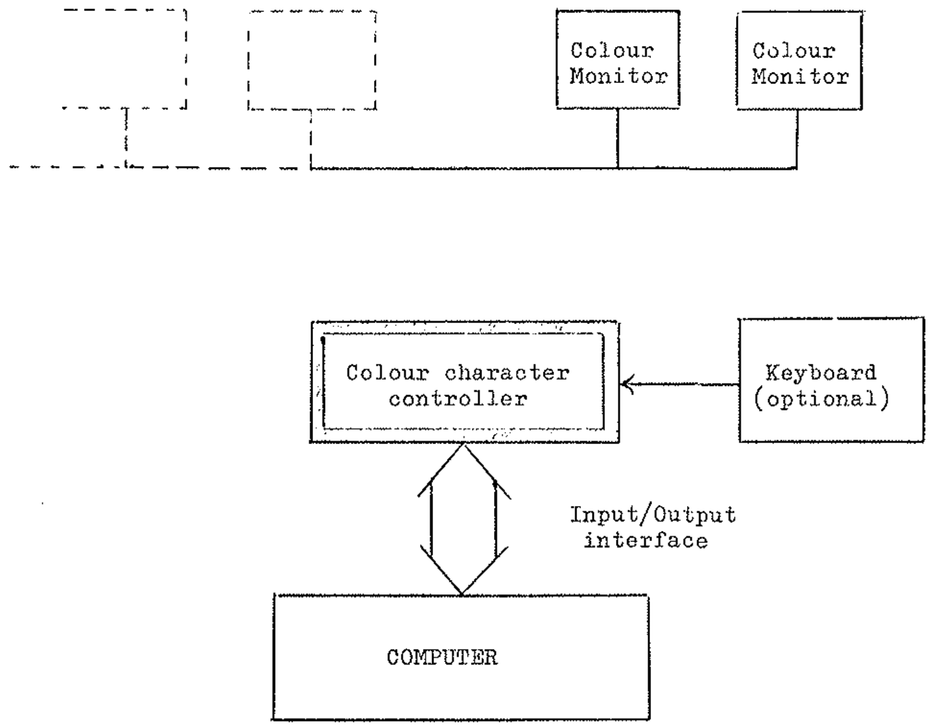

Fig. 1. Environment of colour character unit.

APPENDIX II

TECHNICAL SPECIFICATION FOR THE CONSOLES USED IN THE CONTROL CENTRE

1. Introduction

Four consoles are required for the main control room of the proton synchrotron to be built by CERN on the Swiss-French border near Prevessin (AIN). These consoles will be used to hold that electronic equipment which is directly concerned with operator interaction. Three of the consoles are used by operators for control and monitoring, and are identical. The fourth is used for the control of personnel access and differs in some respects. Additional electronic equipment will be installed in a computer room nearby, in racks which are not the subject of this specification.

2. Shape and Dimensions

The over-all shape of the consoles is shown in Figs. 1 and 2. All electronic elements are mounted in standard 19 inch racks, with the possibility of extra support, such as glides, for units which are deep and heavy. These racks and supports are part of the consoles.

Each unit of each console has three surfaces in which 19 inch electronic chassis can be mounted: a viewing surface, 10 electronic units high*, a controlling or touch surface, 5 electronic units high and, in the case of the three operator consoles, an electronics area, accessible from the back of the console, which is about 22 units high. The table-top is used for the mounting of a small number of control devices which must be kept horizontal. It should be strong enough to be sat on without suffering damage yet not more than about 3 cm thick. The table-top should be precisely 710 mm from the floor, and 350 mm deep. The 5-unit touch-surface should be at 20° to the horizontal and the 10-unit viewing

* an electronic unit is $1\frac{3}{4}$ inches or 44.4 mm. See standards 6-C-1, D1 attached

surface at 10° to the vertical. The wall under the table is constrained to be perpendicular to the touch-surface, and the top of the console is perpendicular to the viewing surface. If these dimensions are observed, male operators in the normal range of size will be able to sit with their knees under the table-top and look horizontally over the top of the console.

The back of the console will probably be vertical, though a slight inward-downward slope for aesthetic purposes is acceptable. In any case the vertical rails for electronics will be mounted at least 12 cm inside the back covers, to allow clearance for plugs and cables.

The individual 19-inch sections of the three operator consoles are 15° apart in plan view, and as close together as possible where the touch-surfaces meet the back of the table-top. The operator consoles each have five, 19-inch sections. The personnel access console has four, 19-inch sections and they are in a straight line. Figs. 3 and 4 give the major dimensions for the two types of console.

All panels covering the console, other than 19-inch ones carrying equipment, should be instantly removable, yet sufficiently well attached not to come off by accident. Magnetic fastenings for example, might be suitable. The means for opening should be invisible or at least inconspicuous.

2.1 VENTILATION

The console also act as a duct for cold air escaping to the under-floor space from the room, which is under slight pressure. The wedge-shaped volumes between the 19-inch sections of each console act as ducts with triangular cross-section. They extend from floor level (holes are cut in the floor-tiles for this purpose) to within about 30 cm of the inside top of the console, and their faces are removable just like the outside panels, so that maintenance personnel can crawl into the wedge-space and have access to the back of the electronic equipment. (see Fig. 5).

The air enters the consoles by grilles cut into the bottom of the five back panels covering the electronic equipment. These grilles should allow air to enter through at least 50% of the bottom 15 cm of each of these panels. Similar grilles should be provided in the back panels of the personnel access console.

In the case of the personnel access console, which has no wedge-shaped ventilation ducts, louvres should be provided near the front of the

.../...

top panels for the flow of air.

3. MATERIAL AND CONSTRUCTION

The consoles may be either of steel or aluminium alloy. The supporting frame must be sufficiently rigid for the 19-inch mountings to carry heavy, deep electronic chassis without deformation, although slides will be mounted for extremely heavy and deep chassis, and back-supports for such slides should be provided. (See drawings).

The frame must be so designed that it will not get in the way of electronic units inside, nor impede personnel crawling inside to service equipment in the console. The weight should be well-spread, as the consoles will be placed on a normal suspended computer floor.

The panels must be thick enough to be rigid, and should be covered in a durable finish whose colour can be chosen by CERN. Those parts of the frame which are visible must also have a suitable finish. The table-top should preferably be of wood covered with Formica, though other finishes would be considered.

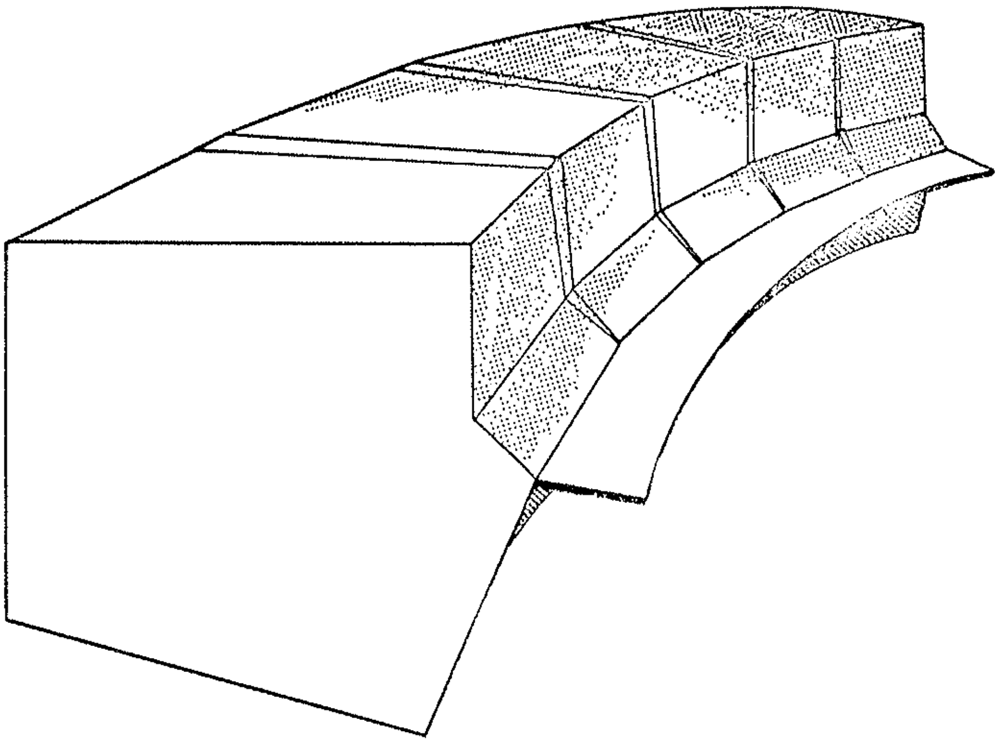

Fig. 1: Over-all view of operator console

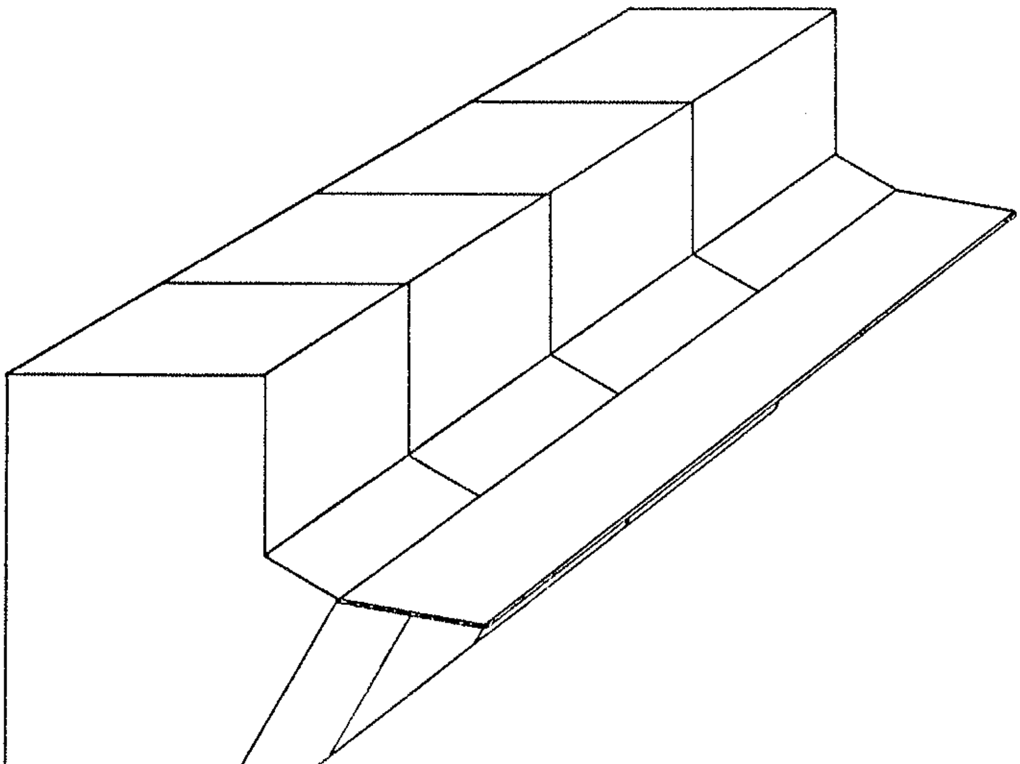

Fig. 2: Over-all view of personnel access console

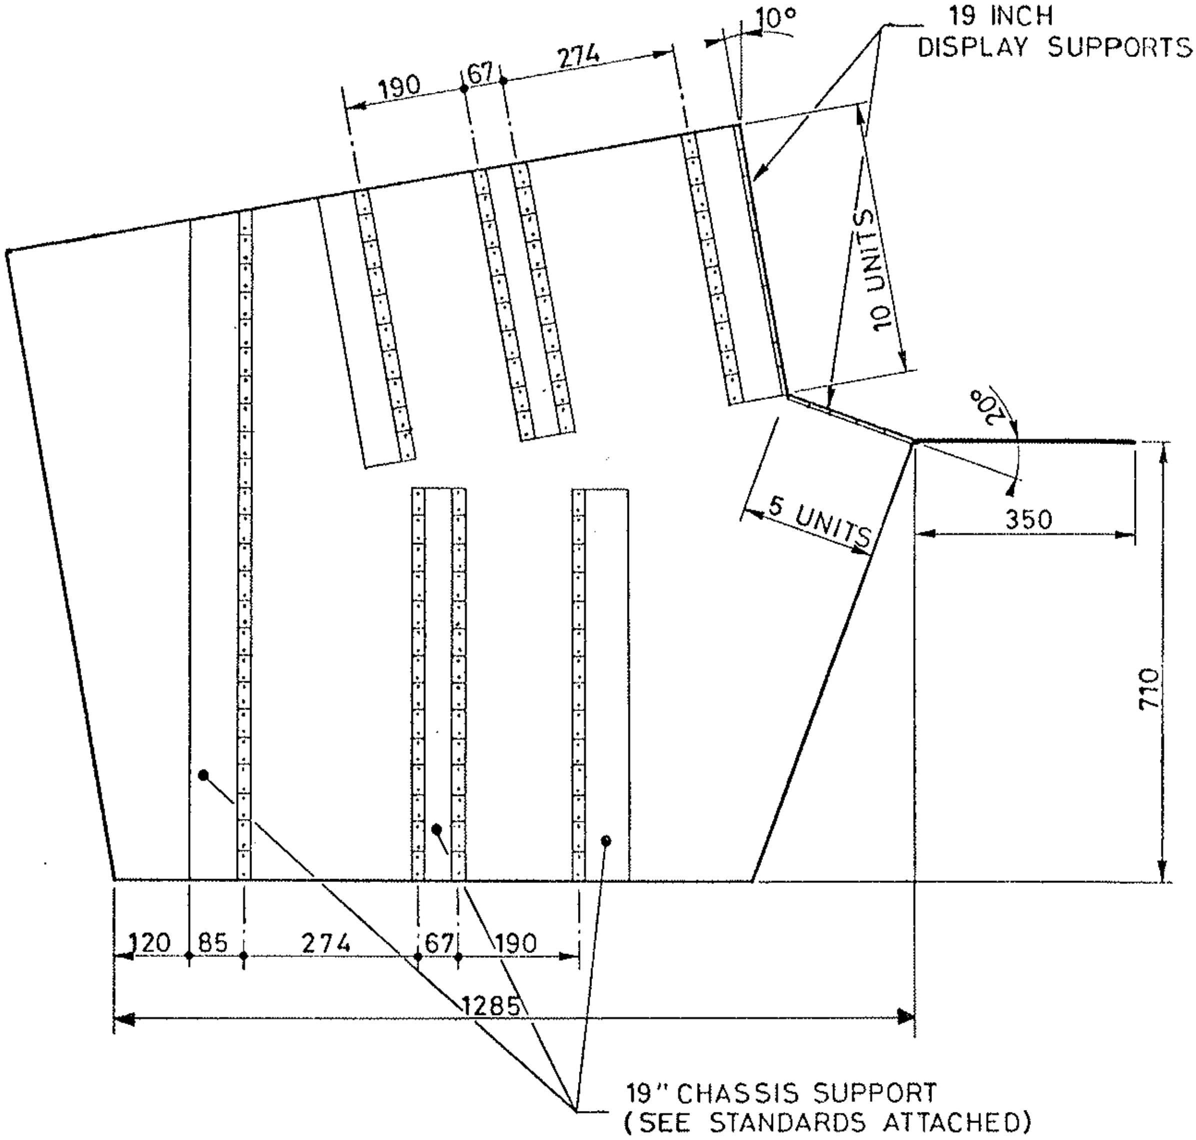

Fig. 3: Major dimensions of operator console

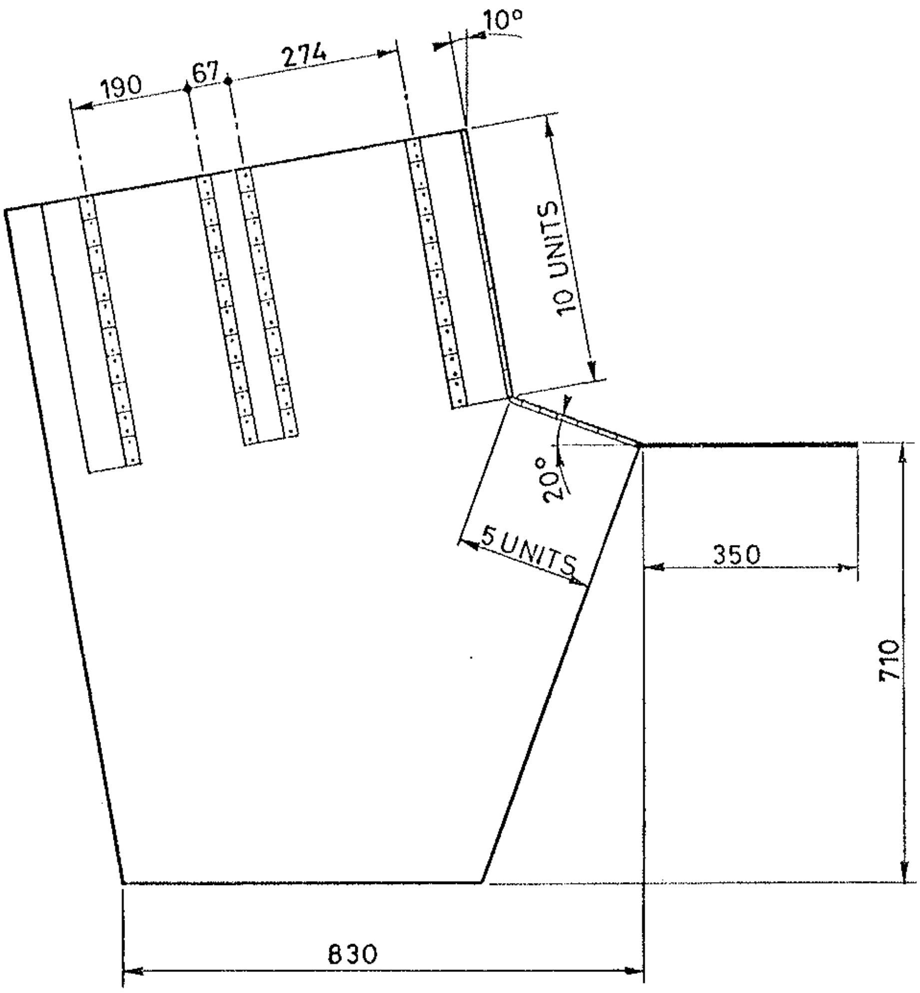

Fig. 4: Major dimensions of personnel access console

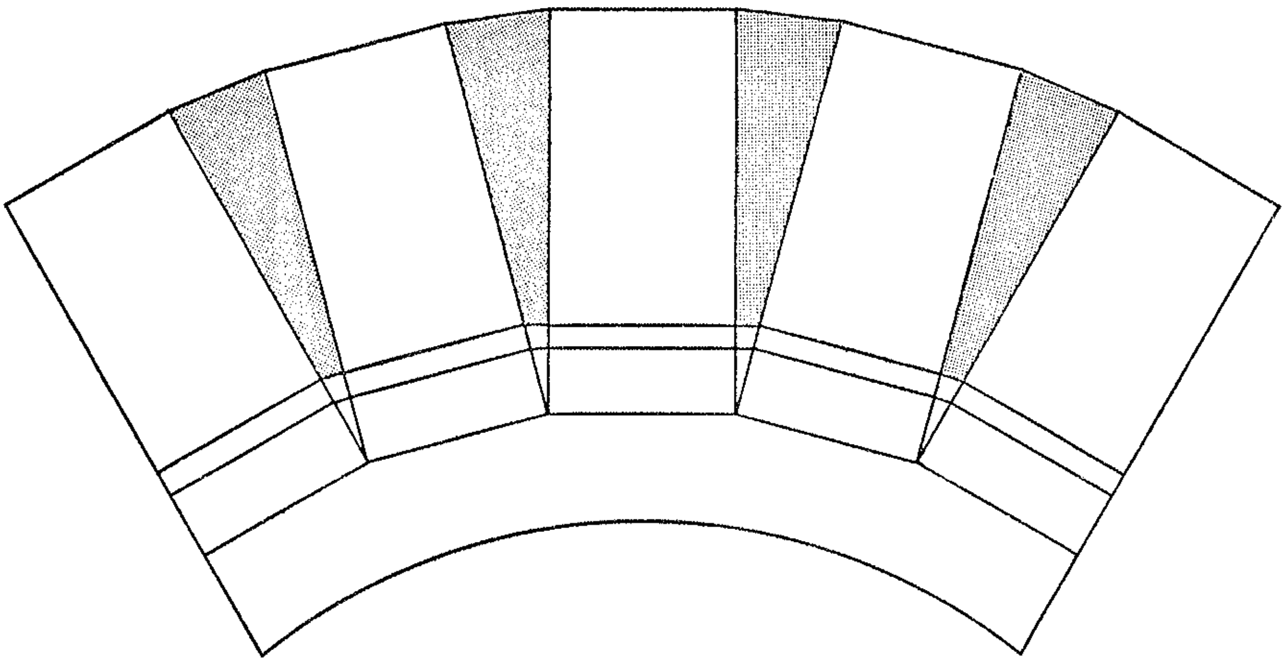

Fig. 5: Plan view of operator console showing wedge-shaped ventilation ducts

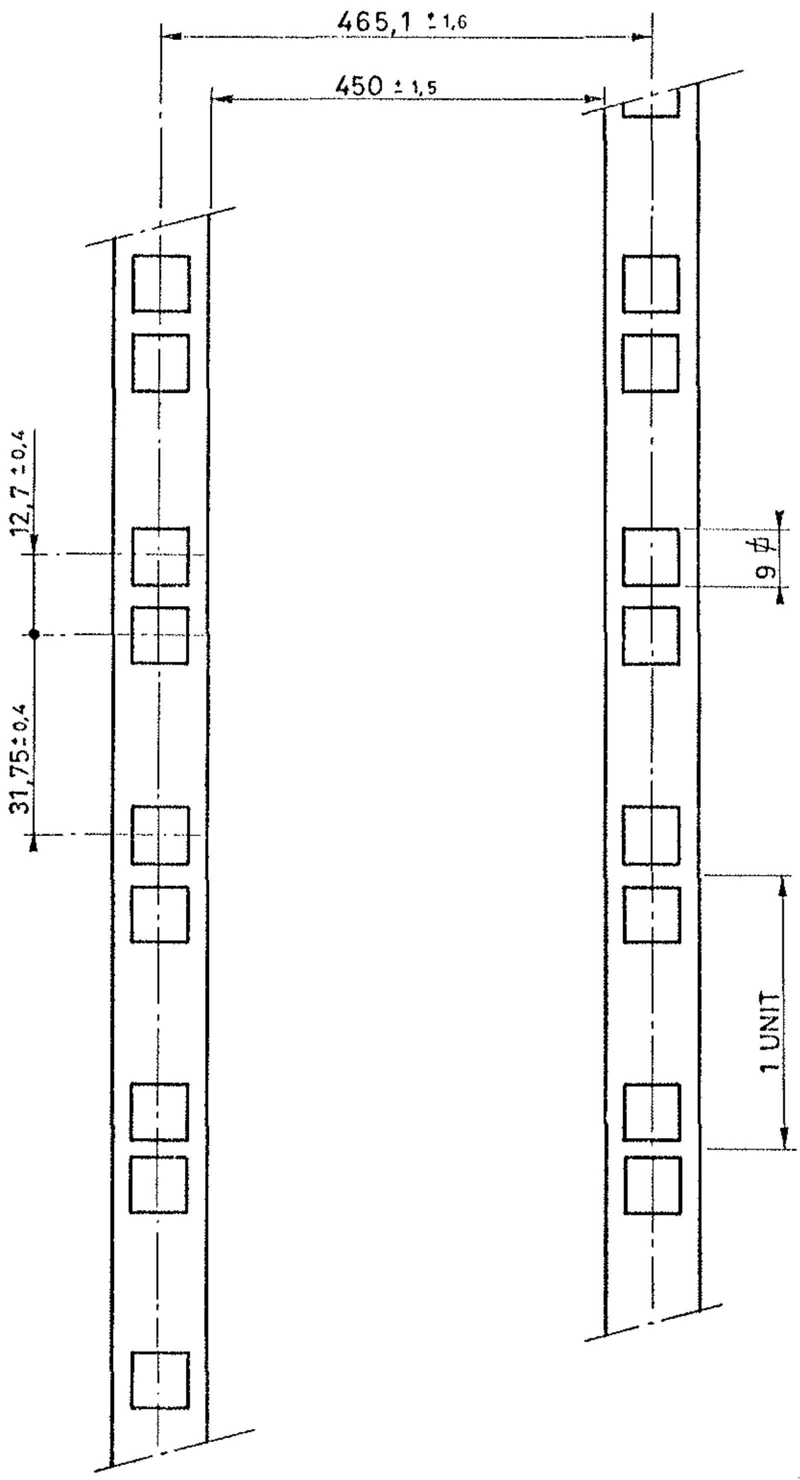

Fig. 6: 19-inch rack dimensions

APPENDIX III

CONRAC II, MODEL ENA MONITOR

The ENA monitor is a transistorized, compact, economy video monitor, designed for continuous operation in a wide variety of closed-circuit and video display applications. It is a complete, self-contained unit which is operated from composite video.

TECHNICAL SUMMARY

ELECTRICAL SPECIFICATIONS

Input Power:	45 watts under normal viewing conditions at 120/240 volts, 60 Hz (525/60 U.S.) or 50 Hz (625/50 CCIR). All performance specifications are met while the line voltage varies from 105 to 130 or 210 to 260 volts A.C. A three wire line cord 6 feet long is provided. Primary Fuse: 3/4 Ampere, Slo-Blo (120-volt operation) 3/8 Ampere, Slo-Blo (240-volt operation) Power Connector: Tower #H-1069-1G
Video Signal:	0.35 volt peak-to-peak minimum for 50 volts at kinescope, sync negative at input.
Video Input Range:	0.35 volt peak-to-peak, 2 volt peak-to-peak composite.
Video Input Impedance:	High impedance bridging which can be terminated by an internal 75 ohm ±5% resistor through a switch located on rear apron. Return loss is greater than 35 dB.
Video Response:	Flat within 3 dB from 100 Hz to 10 MHz.
Video Input Connectors:	Type BNC
Differential Gain:	<5% for 40 volts of kinescope drive.
Pulse Rise Time:	30 volt rise in less than 50 nanoseconds.
DC Restoration:	The black level is constant within 7% from 10% to 90% APL.
Resolution:	650 T.V. lines at CRT center, 500 T.V. lines at CRT corners.
Linearity:	Within 2% of picture height when measured on the vertical axis and within 2% of picture height of active scan area on the horizontal axis using the EIA Ball Chart.
Ultor Voltage:	12KV nominal.
External Sync:	External sync operation is an option.
Kinescope:	ENA9 - 230MB4, ENA12 - ST4500A. All tubes are supplied with a T-band for implosion protection.
Environment:	50°C maximum.

MECHANICAL SPECIFICATIONS (Dimensions do not include connectors and shafts.)

Model	Width	Height	Length	Net Wt.	Shipping Wt.
ENA9/C (Cabinet)	9-7/16"	9-1/16"	9-1/4"	12.7 Lbs.	15.7 Lbs.
ENA9/N (Chassis)	8-11/16"	8-3/16"	8-3/8"	9.5 Lbs.	12.5 Lbs.
ENA9/2R (Dual Rack)	19"	8-11/16"	*8-15/16"	29.0 Lbs.	32.4 Lbs.

*Measured from Mounting Surface (Add 5/16" for Over-All Length)

KINESCOPE: 230MB4 (9") or ST4500 (12")

CHASSIS FINISH: Gold, Irridite Aluminum