

KRIEG IM AETHER

Vorlesungen an der Eidgenössischen Technischen Hochschule in Zürich
im Wintersemester 1978/1979

Leitung: Abteilung für Übermittlungstruppen, Divisionär A. Guisolan

The world's largest accelerator and how it is controlled

Referent: Michael Crowley - Milling

4-1

THE WORLD'S LARGEST ACCELERATOR AND HOW IT IS CONTROLLED

Michael Crowley - Milling

I would like to talk first of all about why we need high-energy accelerators. It may seem very odd that one has to have one of the largest bits of technology in Europe to study the smallest things in the universe, but this all comes from a relationship between particles and waves which was first defined by de Broglie in which the apparent wavelength of a particle λ is given by the relationship $\lambda = hc/E$, where h is Planck's constant, c is the velocity of light and E is the energy of the particle.

If we want to look at something small we look at it with a microscope, but when we get down to something that is comparable with the wavelength of light, the microscope is of no use to us. We then use the electron microscope, in which one uses the energy of the high-energy electrons to give the equivalent of a very short wavelength, shorter than the wavelength of light. Now, when we are talking about atomic particles, the proton has a radius of the order of 10^{-15} m and if we want a wavelength equivalent to that length, then E comes out to be equivalent to approximately 1 GeV or 10^9 eV. So obviously, if we are going to look at the structure inside a proton, we want energies of many GeVs. How do we get these enormous energies? Well, in the early work, where one wanted one or two million volts, one could actually generate that voltage by various means, either cascade generators or electrostatic generators, and just apply the voltage between two electrodes, so that a particle injected at one electrode would be accelerated between them, to pass through a hole in the other electrode so that it could be used for experiments. But once you get above a few million volts you cannot do that. The apparatus becomes extremely large even for 10 million volts, which is about the highest that is commonly used in direct accelerators. Therefore we have to go to indirect means. If we can make a particle go round a circle, which can be done with a charged particle by using a ring of magnets with a magnetic field in the vertical direction which bends the particles round, and then arrange to have one or more regions where you have effectively a voltage gradient, each time the particle comes round it can gain an additional bit of energy. It is rather like trying to throw a stone with a sling. You whirl it around your head and each time it goes round you move your hand a bit to give it a bit more energy until it is going fast enough and then you let go.

So this is the sort of thing one has to do to get these high energies. Now the limit on what one can do is set by the magnetic field one can have in these magnets. Normally, leaving out developments such as superconducting magnets, you cannot use a magnetic field of much greater than two Tesla and, since the radius of curvature for a given field depends upon the energy, this means that if you want to get a higher energy from a given machine, you have to make it a larger radius and put more magnets in.

In the early fifties, as the physicists needed higher and higher energies, CERN was formed because the apparatus was getting so expensive that the individual countries did not want to allocate the greater part of their budget for science to building their own large accelerators. So they banded together to form CERN, where first of all an accelerator to give 25 GeV was built, starting in 1954, ending in 1959, and from then on experiments were carried out at CERN. About 1963 it was decided that still

Adresse des Autors:

Michael C. Crowley-Milling
CERN European Organisation for Nuclear Research
1211 Geneva, Switzerland

4-2

higher energy machines would be needed, and between 1963 and 1971 discussions were going on as to what it should be, where it should be, and so on, and ended up in 1971 by the member states agreeing to build what was then a 300 GeV machine at CERN. The 25 GeV machine had a radius of 100 m and for this 300 GeV machine one has to go to a radius of 1.1 km. Now in fact we have improved on this because of advances in the magnets and various other ancilliary things and now it is a 400 GeV machine, but with the same radius.



Fig. 1 View of the CERN site and the town of Geneva

4-3

Fig. 1 shows the site, the town and lake of Geneva, with the Mont Blanc in the background. The position of the 400 GeV machine is shown by the white line. When you bear in mind that, in order to get to 400 GeV, the protons have to go round this circle about 100,000 times and still stay within a narrow vacuum tube, you can understand that the stability of the position of the magnets is extremely important. In order to get this stability, we had to build the accelerator in a tunnel underground where one could build it in good molasse rock. This meant building it at an average of 40 m beneath the surface, because the top 20 m or so consists of unstable moraine.



Fig. 2 Tunnelling machine breaks through to complete the ring

In order to build this tunnel, we used a tunnelling machine which had a head 4.5 m in diameter. This head rotated slowly, forcing cutters into the molasse. Fig. 2 shows the break-through after going round the 7 km circumference circle. It was a triumph of the geometer's art because it came through within 15 mm of where it was supposed to be after 7 km, and it did not deviate anywhere along its path more than the 7 cm which was the tolerance.

4-4

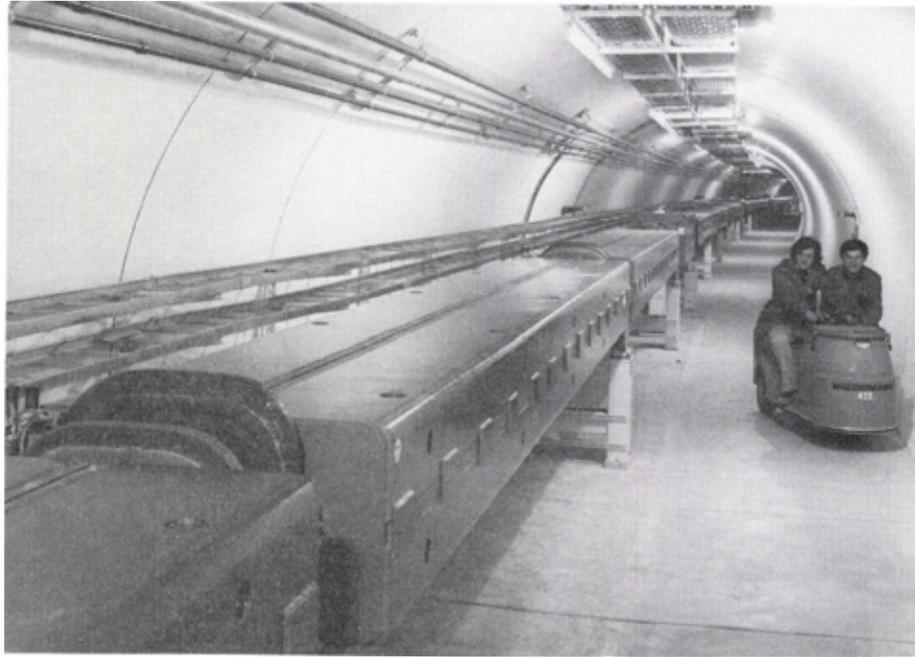


Fig. 3 The completed tunnel with magnets in place

Fig. 3 shows the finished tunnel and the two types of magnets. The long objects in the foreground are the bending magnets to bend the protons round, and the shorter ones are the focusing magnets to keep them focused. In between them, periodically, there are short sectors where there are correcting magnets and beam position monitors, etc.

4-5

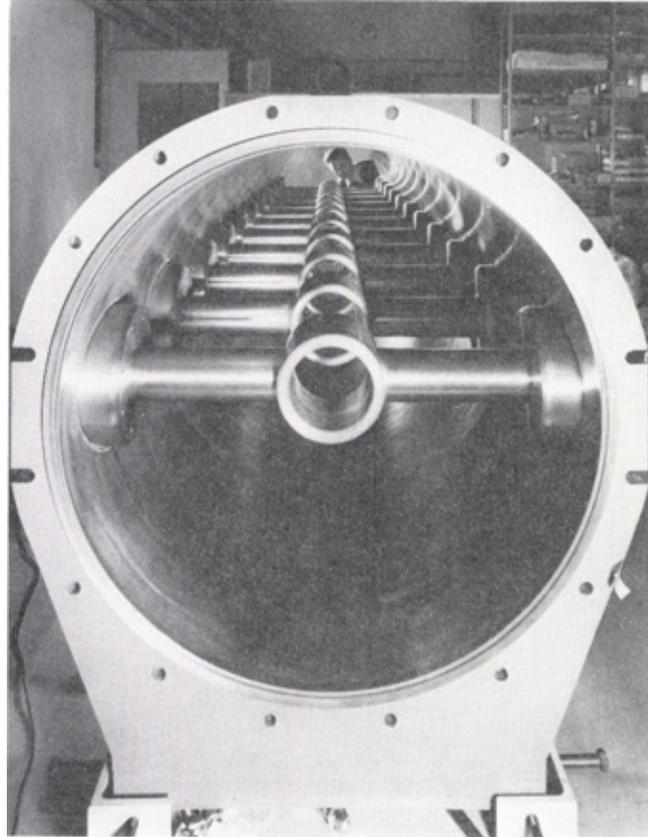


Fig. 4 Internal view of cavity section

The two types of magnets do the job of taking the beam of protons round the tunnel. Then we have to accelerate it, and for that purpose we use radiofrequency power because one has to have a situation where the protons, every time they come round, see a field accelerating them and not a field decelerating them. This is done by having a cavity system of the type shown in Fig. 4, where there are so-called "drift tubes" and the geometry of this cavity is such that when it is excited at 200 MHz, the maximum field is between these drift tubes. This field reverses every half cycle of the 200 MHz, and the protons are formed into bunches, so phased, that the particles are going across a gap when the field is in the right direction to accelerate them, and are then passing inside the drift tube, where they are largely shielded from the field, when the field is in the decelerating direction. So by this trickery we persuade the protons that they are in an accelerating field the whole way. Each cavity is 20 m long, and there are at present three of these cavities, which give a total of about five million volts to the protons each time they come round. One of the cavities is shown in Fig. 5. The power is fed in at one end and the remaining power is fed out into a water load at the other end.

4-6

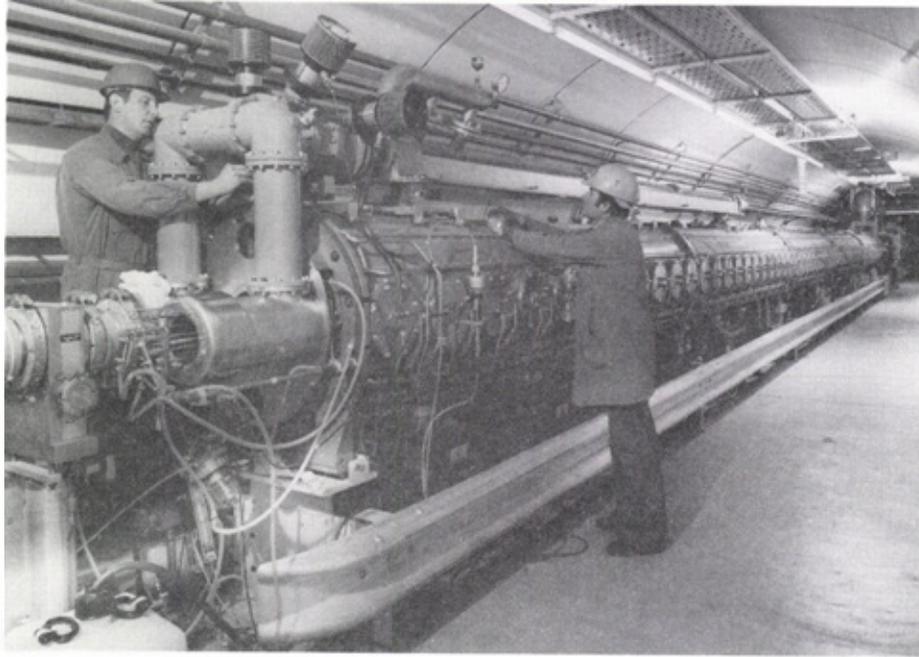


Fig. 5 Radiofrequency cavity in the tunnel

Having got the particles into the accelerator and having accelerated them up to 400 GeV, we then need to take them out in order to be able to use them for experiments. As I said before, the amount you can bend the protons depends upon the magnetic field and their energy, and by the time they have got up to 400 GeV they are very very stiff and it takes a lot to deflect them. So the whole of the equipment shown in Fig. 6 is needed to deflect the protons out of the accelerator towards an experimental area. First of all there is about 12 m of electrostatic deflection, which causes the particles to be extracted to move a few millimetres from their normal orbit. Then there are some thin septum magnets, in which there is a current-carrying strip about 4 mm wide with about 20,000 amps in it, and the magnetic field from this bends them a bit further, followed by more vacuum tanks containing thicker septa with correspondingly more current, which separate the beams just sufficiently for them to occupy separate vacuum pipes as can be seen at the right hand side of the photograph.

4-7

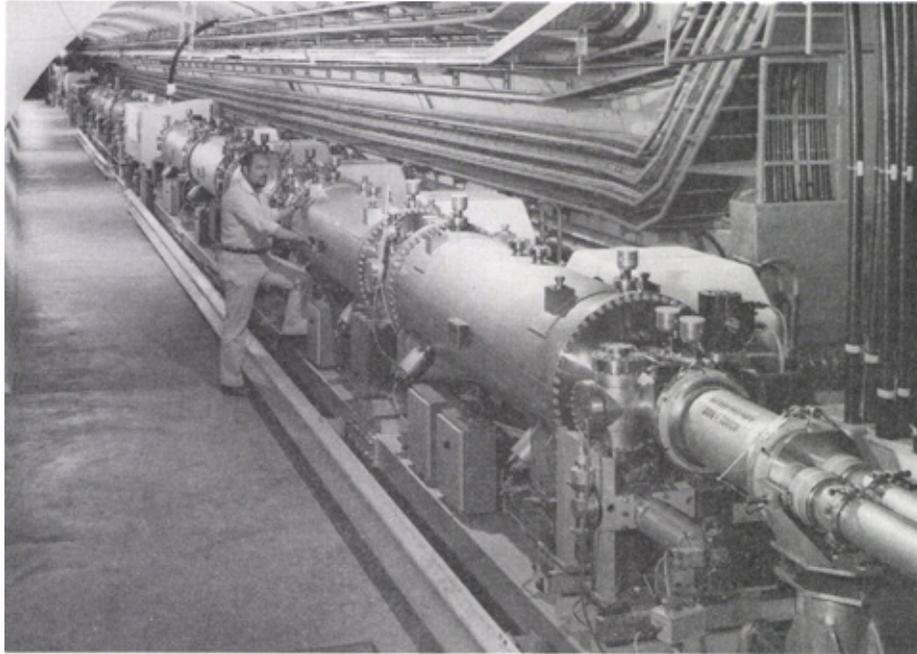


Fig. 6 Extraction equipment in long straight section

Now I do not want to deal very much with experiments here, but just to show that these themselves are fairly large engineering objects, I will give two illustrations. Fig. 7 shows one of the experiments, which consists of large numbers of iron discs, toroidally magnetized, with scintillation counters which record the position the particles pass through. This is actually an experiment for looking into the interaction of neutrinos with iron, neutrinos being rather peculiar particles that have no mass and no charge and only exist by virtue of their energy of motion.

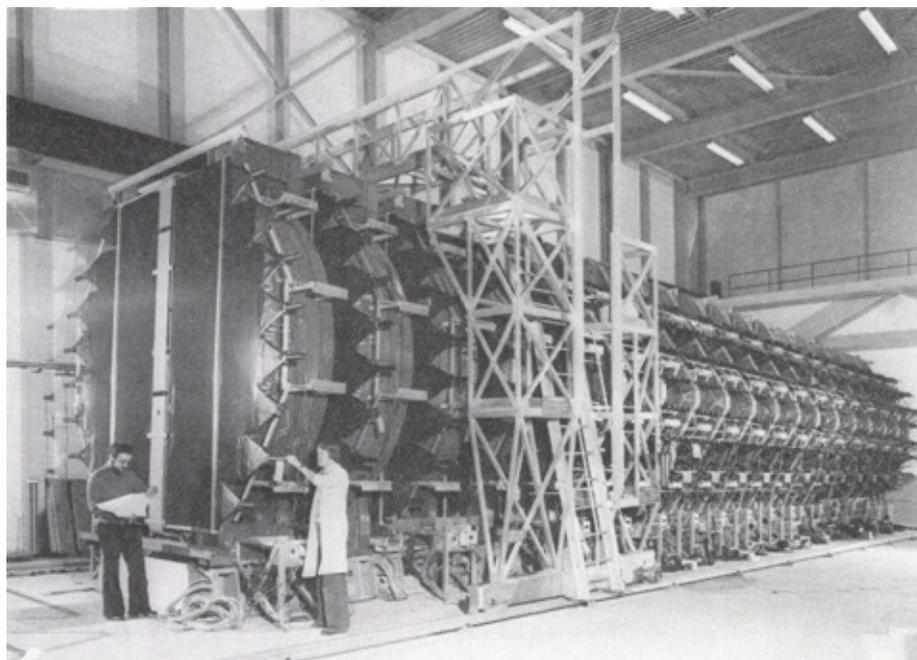


Fig. 7 Experimental apparatus to detect the interaction of neutrinos in iron

4-8

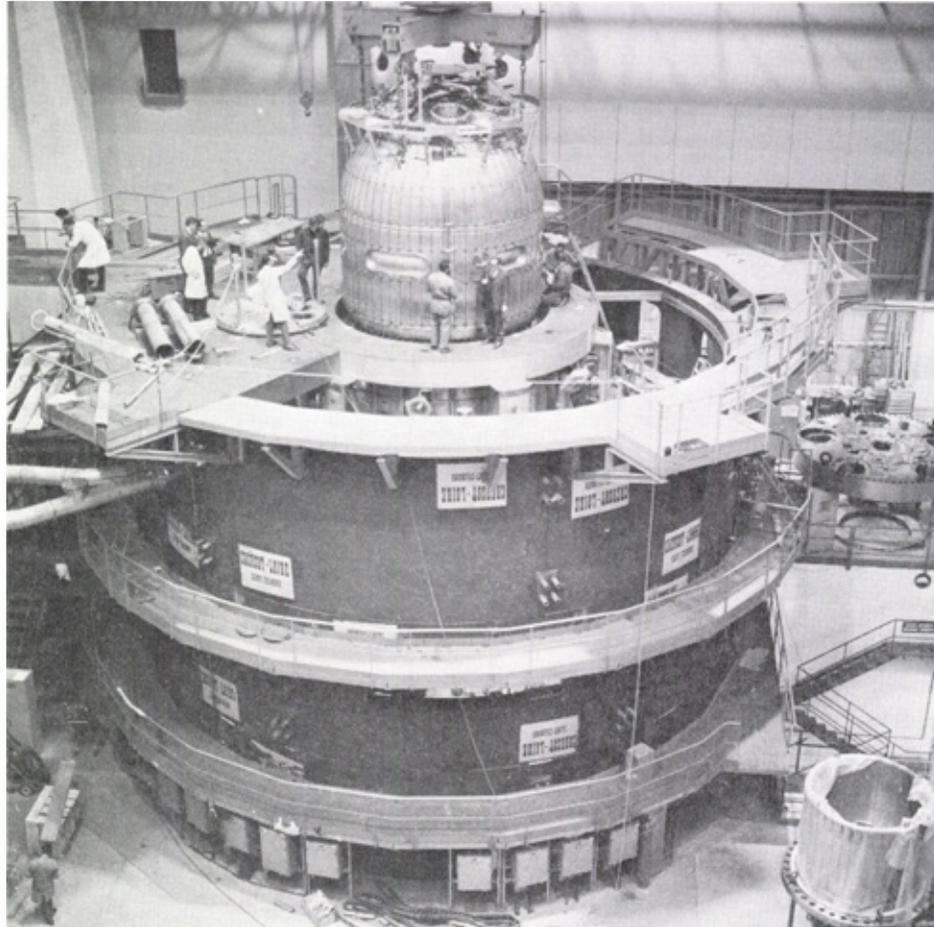


Fig. 8 The liquid-hydrogen container being lowered into the magnet structure for the large bubble chamber detector

Fig. 8 shows another large piece of experimental apparatus which is a hydrogen bubble chamber. One way of detecting the passage of a charged particle is by passing it through a superheated fluid. For simplicity of the analysis, one wants to use the simplest of all target materials, hydrogen. So one normally uses liquid hydrogen. There is about 30 cubic metres of it and it is held at about 5 atm pressure just below its boiling point, about -270°C . Then, when the particles go through, a plunger in the chamber is suddenly pulled down to reduce the pressure in the hydrogen, which immediately starts to boil. But it first boils along any nucleating centres which are the ionization tracks of the particles. So, in the first few milliseconds, a row of tiny bubbles is formed along the path of any charged particle and then you flash lights and use multiple cameras to photograph these tracks and then immediately compress the hydrogen again so that the boiling stops. This is quite a tricky business with 30 cubic metres of liquid hydrogen; it is quite a large device and quite important and difficult from the safety factor point of view.

4-10

the largest one from the core point of view and which is used for solving large mathematical problems that other ones can send to it. Now when we set out to do this, which was in 1971/72, we sent an enquiry round to about 80 manufacturers of computers in Europe and some outside Europe asking for their suggestions about supplying a network like this, particularly what one would do to transfer messages between the various computers, because in such a network this was the heart of the system. The response to this was very discouraging. The computer manufacturers wanted just to supply computers and they had not got round to thinking of networks by then and one or two specialized firms who were interested in network work had not the ability to supply suitable computers. So we split the job and on one hand we sent out specifications for an ideal minicomputer to do this job and on the other hand did a very detailed design of a packet-switching message-transfer system which we then sent out for bids for implementing it. The interesting thing about the minicomputer tender was that the lowest bid that most nearly came to our specification was from a Norwegian firm, Norsk Data and this was a computer that they were just starting to produce. They had only a prototype at the time, but they had an earlier machine which had an extremely good software system and the new machine was going to be using exactly the same software. My experience is that a firm is able to keep target dates on making hardware but not necessarily on software and so we decided to take the slight risk of this computer not quite being in production. As it happens, we were justified in taking that risk, because they managed to supply them pretty well on time.

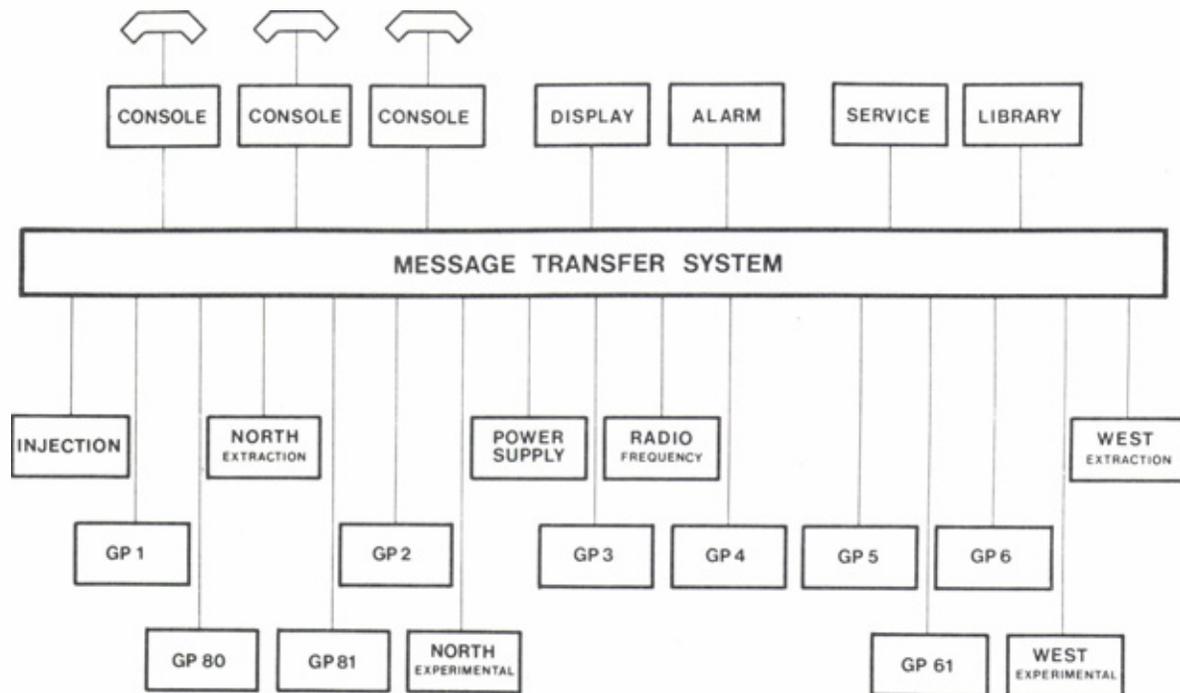


Fig. 10 Control computer system for the SPS

The computers alone are not enough to make a control system; one has to join the computers to the equipment that has to be controlled. Since there was some delay in selecting the computer we were going to use, we wanted to use a computer-independent interface system. Now the only one available at the time, and to some extent the only one available now, that is truly computer-independent is the CAMAC standard. This was developed for instrumentation in high-energy physics. It is not ideal by any means for control -- it is rather fragile and it is rather expensive -- so we decided to use the CAMAC as the primary interface but to do most of the connection to the equipment with a secondary interface multiplex system which we developed ourselves. The control unit for this multiplex system plugs into a CAMAC crate and has a bus system into which up to 32 stations can be put, each station having eleven slots for interface modules. Instead of the normal method of interfacing, which is favoured by the computer manufacturers where they prefer to supply so many digital inputs and so many digital outputs, so many analogue inputs and so on, leaving the customer to make the junction boxes to sort out which goes where, the new modules were tailored according to the apparatus they were

4-11

to control. For example, one module might have six status acquisition, three digital outputs, one digital input and one analogue input, which could be the full requirements for a power supply or a vacuum pump controller. This, of course, was possible in our case because we had so many units of a given type. We had, for example, a thousand vacuum pumps, so it was obvious economical to design a special unit for that rather than use part of the facilities of three general purpose units, which one would have to do if one had stuck to CAMAC.

So far I have only mentioned the hardware, but that is really the simplest part of designing a computer control system and so I would like to go on to the particular troubles one has with a control system for a large accelerator. It is not like designing a process control system for, say, a steel works where you have laid down well beforehand what is to be controlled; so many rolls, so many reheating furnaces and various things like that, for which the programs, once they are written and working properly, are probably not going to be changed much for many years. First of all, a new accelerator like this, on quite a different scale from the earlier accelerators, uses many new parts, many new ideas and the people developing the new equipment often do not know how they will want to control it until they have actually built most of it and tried part of it out, so that one starts with a disadvantage that the information required for the control system will be late. Secondly, an accelerator is not a fixed object. Roughly a tenth of an accelerator's operating time is taken up with experiments on the accelerator itself to try and make it go to higher energy or give more intense beams, or even accelerate different particles. Therefore the control system has to be extremely flexible to allow one to change things rather easily, and this is difficult with a normal type of control system which is programmed in the normal way. So we had to think of some way of getting round these two problems, and this was done by providing a very simple means for the engineers and physicists designing the equipment to write their own programs to control it, rather than the conventional way where the designer of the equipment makes out a specification for a program to a professional programmer. Often, after the program has been written, the engineer finds that it does not do exactly what he wanted, or his apparatus does not behave exactly as he thought, and the program has to be modified. This usually results in the software being late, and needing far more people than was originally expected. We had to find some means of overcoming this problem and the method we chose was to use an interpreter for a high-level control language. One of the main advantages is that the source program is kept and not converted in compiled code that is unintelligible to the user, so that interactive debugging and editing is made very much easier. There was no language available that was suitable. BASIC is normally the language that is interpreted, but it is rather inadequate in its facilities. The language called FOCAL, which is a DEC proprietary language, is rather better in this respect, but of course we could not use FOCAL with Nord computers, so we developed our own which we called NODAL, which is very much more powerful than any of the interpretable languages, and caters for multiprocessing programming.

For a control system, it is not enough to have a good language and a simple means of writing programs and debugging, but you also have got to get into contact with the apparatus to be controlled and this involves some sort of data tables or data base. Another decision we made was to use a fully distributed data base and to distribute it not only geographically but also by function. Whereas it was very difficult to get the hardware designers to define the way their complicated apparatus should be controlled, it was relatively easy to get agreement early on how the individual basic pieces of equipment should be controlled. For example, in controlling a power supply it is necessary to define the steps that have to be made to check the interlocks and switch on, run it up to the required value and various precautions and checks that have to be made: such things do not change whatever the power supply is being used for. Similarly, for things like vacuum pumps, water pumps, etc., it is necessary to define how to carry out the pressure measurements and flow measurements, valves opening and shutting, and so on. These sort of things can be settled early in the program. So we wrote a number of little programs which we called data modules and each of these data modules is essentially a handler for a particular type of equipment with the appropriate data table. The data table for a power supply, for example, keeps the addresses of the hardware, it keeps the demanded value, the maximum allowable values, conversion factors, calibration factors of the measuring devices, and so on. A diagram of the system is shown in Fig. 11. The programs for the interpreter are fed into text buffers and the interpreter, which is a re-entrant program, deals with these in order of priority. When the interpreter finds the name of a system variable in the program, it looks up the NODAL name table. This allows one to give the same type of equipment different names for different purposes. For example, you may have 200 auxiliary power supplies that are being used for many different things, but for the maintenance engineer of the power supplies they can just be called Auxiliary Power Supplies, number 1 to 200. For the man who is setting up, say, an injection beam line and is using one of those power supplies for supplying a magnet that is doing the vertical bend on the beam, he can call that power supply "Injection Vertical Bend". When the interpreter comes across such a name, it looks up in the NODAL name table which data module to call and pass on the information. The data modules use a common driver to operate through the CAMAC and multiplex systems. With this arrangement, the basic system and the data modules were written by about four professional programmers, but the whole mass of applications programs for running the accelerator was written by about 60 of the professional and technical staff, as a part time activity, a lot of whom had never written a program before, or if they had done so, it was just a batch FORTRAN program. This certainly did what one journal described as "breaking the software barrier".

4-12

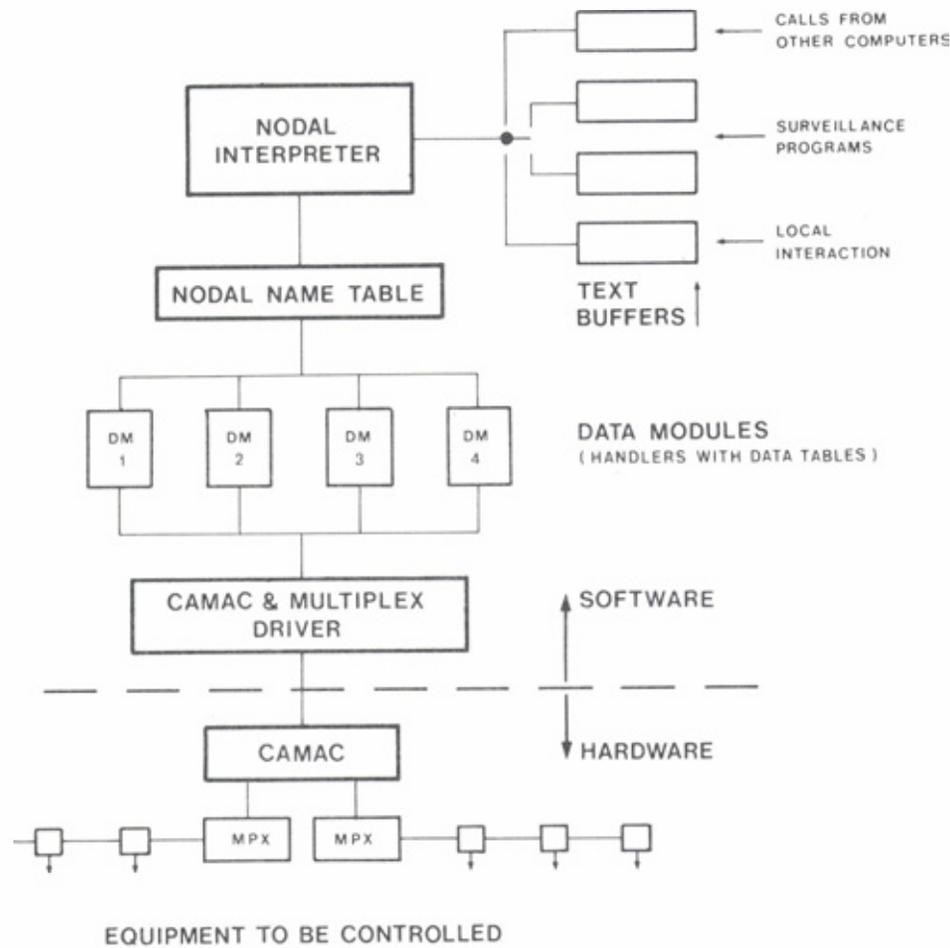


Fig. 11 Simplified schematic of software system

Of course another most important thing is how you actually carry out the control. The early accelerators had specialized controls where everything was brought back by cable to the control room, which then involved quite short distances, and one had specialised racks with separate controls for RF, injection, extraction, and so on, and the operator went physically round to set and look at various things. Here we have gone to completely the opposite extreme. We have consoles that are completely general purpose. They are peripherals of the console computer and we have four of these in practice, but the whole accelerator can be set up from any one of them, by the aid of three main interaction devices, the most important of which is the touch screen.

If you have got a general purpose console it is necessary to select what you want to control. This is normally done by having a tree-like structure where you make a first choice amongst a number of major systems, followed by a choice amongst the sub-systems of that system, and so on. This can be done by various methods: by typing in, by using light pens, by masses of pushbuttons and so on, but we developed for this purpose what we call the touch screen which gives you, in effect, an infinite number of push-buttons. This is shown in Fig. 12. The computer draws images of pushbuttons with legends on a cathode ray tube, and over this a sheet of glass onto which are deposited a series of fine grids, so arranged that when you put your finger on, it increases the local capacitance and triggers a circuit which informs the computer which button has been pressed. So having made your first choice at the most general level, which we call the Trunk, the buttons change to give a further choice the next layer down. One big advantage of this system is that, because they are under program control, you only show as many buttons as are relevant at any level of the tree, so you can guide the operator down the way he should go. Then having got down to the level where you want to actually operate on something, you get a dis-

4-13

play of some sort. Here the second interaction device comes into play - the rolling ball, or tracker ball as it is sometimes called. This is usually used to control a cursor on the display, to select the particular item it is required to operate on, or the particular part of a function to modify. If a value is to be varied, the third interaction device, a knob driving an incremental encoder, is automatically joined to the appropriate variable. Displays can show tables of variables, but most frequently they show mimic diagrams, to simplify the task of the operator.

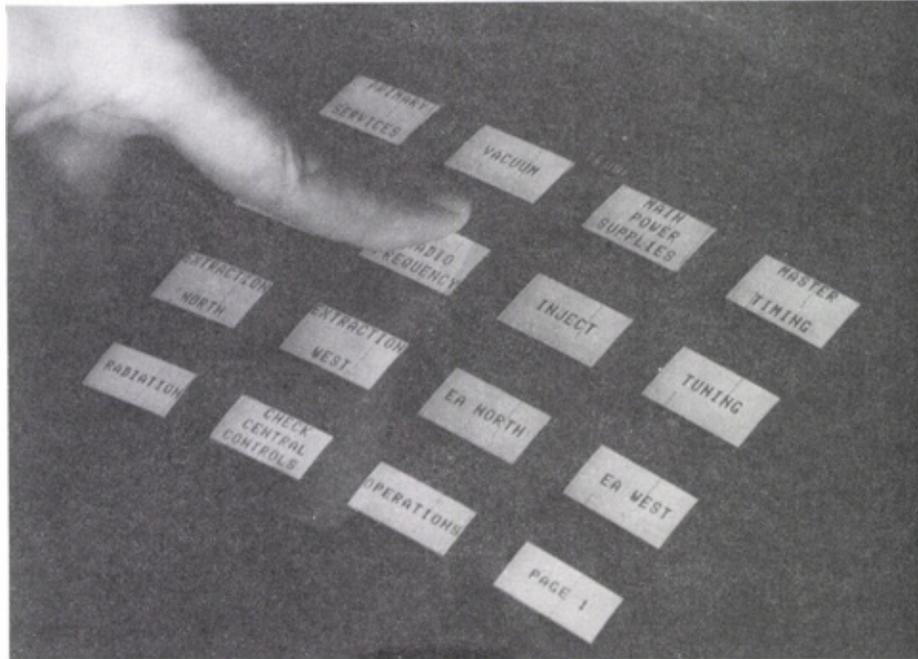


Fig. 12 Touch button panel

All displays are in television raster form and both black-and-white and colour monitors are used. For the colour monitors a convention is used. Figures in white are actual values measured from the equipment. Anything which is under the operator's control is in green and if he tries to set anything outside the allowed range, it turns to red. Transient values, where a change has been demanded but the equipment has not settled down to the new value, are in purple.

A very important part of the equipment is the alarm system, and since we have completely general purpose consoles, it is necessary to have a separate dedicated display for alarms, because the system does not know which console the operator is sitting in front of, but it is also a good thing to keep them separate. Surveillance programs run in most of the computers, and we find that it is necessary to have some consequential analysis incorporated into these programs, to avoid the alarm system being swamped with unnecessary messages.

A machine like this provides a high level of radiation while it is running. No one must go into the tunnels when the machine is on, and the computers are used to assist the operators in controlling the access of people. For safety's sake, all the doors are hard wired so that if anybody breaks in when they should not, it shuts the whole machine down, but there are also facilities for letting people in under controlled access when the machine is switched off. Fig. 13 gives a general view of the control room. The four control positions are arranged in pairs in the two semi-circular consoles. The reason for having four positions is that, although the whole accelerator can be controlled from one, when you are setting up the accelerator you want to set up parts in parallel rather than in series, and also one need to set up the beam lines to the experimental areas. We have found that four is about the right number for this. When the machine is running normally, one, or at the most two, are used but the other two are then available for program development or some other investigation work.

4-14

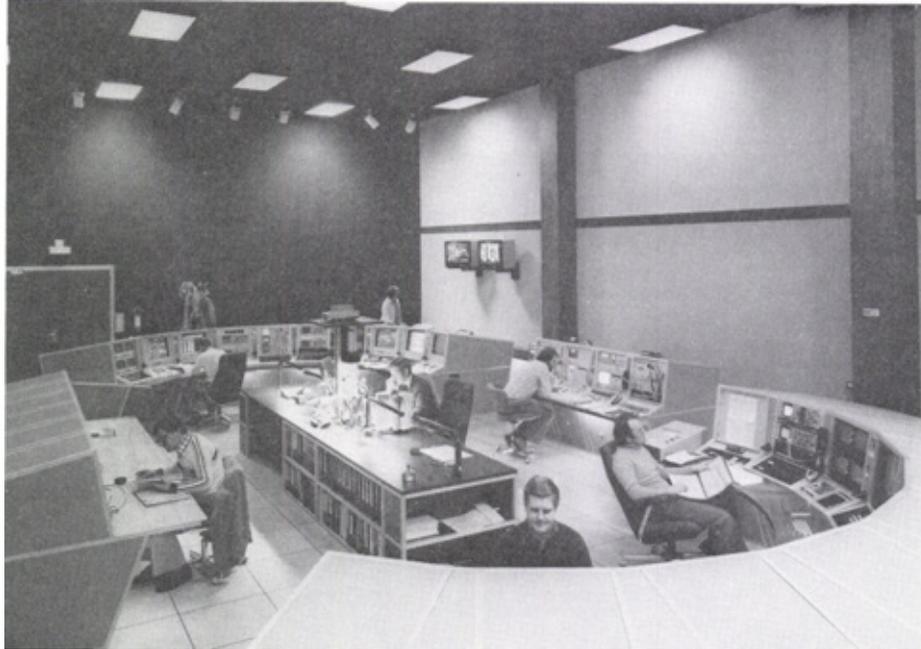


Fig. 13 General view of the control room

Before I finish I would like to add a word on new developments. One of the difficulties of doing physics with a high-energy machine is that if you shoot protons, for example, into a hydrogen target where proton hits on proton, quite a lot of energy of the incoming proton is just transferred to the target proton and knocks it on, rather like when you hit a billiard ball with another one. They both go on together. There is great interest in producing exotic particles which do not occur freely in nature or have very short lives, and for that one needs to release a lot of energy which then usually reappears as pairs of these strange particles. If, instead of having a moving particle hitting a fixed target where the two go on together, we could have two particles that came in opposite directions and collide, then the amount of energy given up would be very much greater, especially if those particles are a particle and its antiparticle, in which case they annihilate each other, forming a pure burst of energy from which these strange particles can be formed.

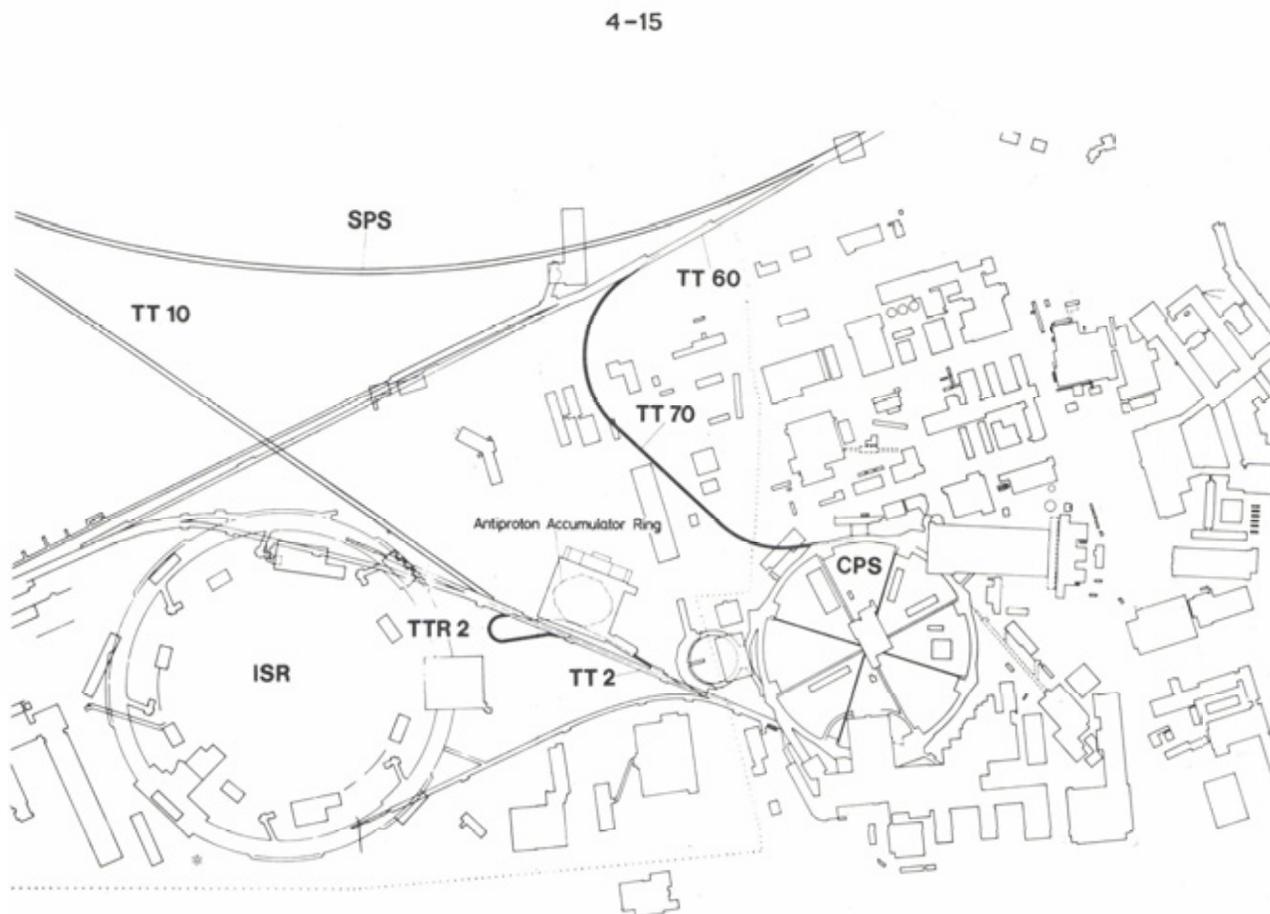


Fig 14 New beam transfer tunnels for antiproton injection

So what we are doing at CERN is to generate large quantities of antiprotons, store them and inject them into the SPS to accelerate them in the opposite direction to the protons. Fig. 14 shows the layout. We are proposing to take protons from the 25 GeV machine (CPS) onto a target in a new building marked "Antiproton Accumulator Ring". A certain number of antiprotons are formed as a result of the collisions of the protons with the heavy molecules in the target. We then collect these antiprotons by means of a magnetic channel and direct them into the accumulator ring. Since the numbers of these antiprotons are small and since their angular divergence can be quite large, we need to collect a large number of bursts and to reduce their angular divergence, in order to make them into a fine enough beam to go into the SPS. In order to do this, one of our staff invented a method which is called stochastic cooling, which enables one to progressively reduce the transverse momentum of a beam of circulating particles. After cooling, the antiprotons are sent back to the CPS, through the return loop TTR2 and TT2, to be accelerated up to 25 GeV, and injected into the SPS via the new tunnel TT70, to circulate in the opposite direction to the protons which will have previously been injected into the SPS. Two new underground experimental areas are to be built to house the apparatus to detect the results of the proton-antiproton collisions.

4-16

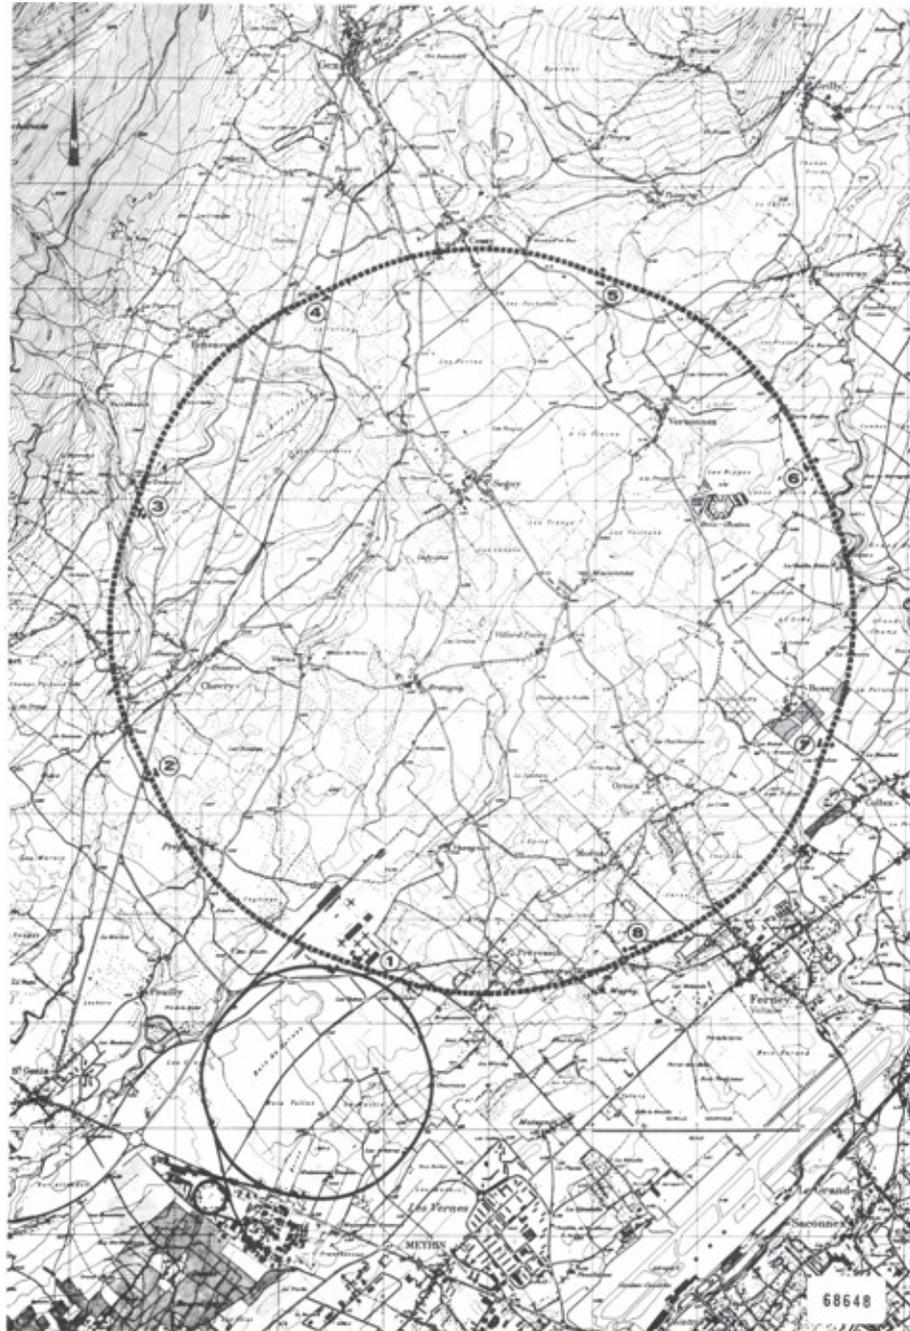


Fig. 15 Possible layout of large electron-positron colliding beam accelerator at CERN

Finally, an international body, the European Committee for Future Accelerators, is looking into the problem of what would be the best sort of accelerator to build in Europe as the next one. They have agreed that, from the physics that may be most important then, it would be best to have a colliding beam machine for electrons and their antiparticles, positrons. A problem here is that when you bend high-energy electrons from a straight path, they radiate power, which has to be made up by the RF system. Fig. 15 shows the layout of such a machine if built at CERN. On this scale, the original CPS is hardly visible at the bottom. The SPS ring, 2.2 km diameter, can be seen above it, and the new ring, which would have to be over 7 km in diameter to give the required centre-of-mass energy of about 200 GeV, fills most of the picture. Detailed design studies are now being made, and, if the approval of the Member States can be obtained, construction could start about 1982.