LRP 518/95 July 1996

THE CONTROL OF TCV PLASMAS

J.B. Lister, F. Hofmann, J.-M. Moret,
F. Bühlmann, M.J. Dutch, D. Fasel, A. Favre,
P.-F. Isoz, B. Marletaz, Ph. Marmillod,
Y. Martin, A. Perez & D.J. Ward

submitted for publication to
Fusion Technology

The Control of TCV Plasmas

J.B. Lister, F. Hofmann, J-M. Moret, F. Bühlmann, M.J. Dutch, D. Fasel, A. Favre, P-F. Isoz, B. Marletaz, Ph. Marmillod, Y. Martin, A. Perez and D.J. Ward

Centre de Recherches en Physique des Plasmas
Association EURATOM - Confédération Suisse
Ecole Polytechnique Fédérale de Lausanne
1015 Lausanne, Switzerland

(received

ABSTRACT

The general control of tokamak plasmas has evolved considerably over the last few years with an increase in the plasma pulse length, an increase in the control of additional heating and fuelling and an increase in the degree to which the shape of the plasma can be varied. The TCV tokamak is specifically designed to explore the operational benefits of plasma shaping over a wide variety of plasma shapes. Consequently, considerable attention has been given to the control of the poloidal field coil currents which impose the desired shape. This paper deals with all aspects of the control of TCV plasmas, from the diagnostic measurements to the power supplies, via particular control algorithms and overall supervision.

Submitted by :
Dr. J.B. Lister, CRPP-EPFL, 1015 LAUSANNE, Switzerland
Fax: ++41-21-693-5176
E-mail: LISTER @ ELTCA1.EPFL.CH

I. INTRODUCTION

The general questions of tokamak control impinge on many fields and for this reason, this paper sets out to discuss all aspects of the control of TCV plasmas. We describe the choices, the methodology and the experience gained, from the diagnostic measurements to the feedback control outputs via feedback control algorithms, from the grid supply to the field coils and from the control-desk to the operation of the experimental systems. Most of these topics are of course interlinked.

TCV (Tokamak à Configuration Variable) was constructed to explore the detailed effects of shape on plasma performance. The design of TCV is up-down symmetric, Fig. 1, and its constructional parameters are : R = 0.88m, a = 0.24m, height = 1.44m, B_ϕ = 1.43T. A total flux swing of 3.4 V sec is provided by two ohmic heating coils. To be freely capable of producing widely different configurations, TCV is equipped with a set of 16 independently powered poloidal shaping coils, eight in a vertical stack on each side of the vacuum vessel, shown in detail in Fig. 2. The vessel itself has an almost rectangular aperture with a height/width ratio of 3.0. It has an irregular shape and varying thickness with the poloidal distribution of electrical conductivity concentrated in the top and bottom plates. The electromagnetic properties of the vacuum vessel are most usefully described by a set of eigenmodes. The slowest up-down symmetric and anti-symmetric eigenmodes have decay time-constants of 13.4 and 8.1 milliseconds respectively. Details of the construction of the TCV tokamak can be found in [1].

The shapes produced in TCV since the first plasmas obtained in November 1992 already cover a wide range of current, elongation and triangularity, and include limited and diverted discharges with upper and lower Single-Nulls and Double-Nulls. Plasma confinement modes include both ELMy and ELM-free Diverted ohmic H-modes, as well as Limited ohmic H- and L-Modes. The achieved plasma currents extend up to 810 kA, the elongations to 2.05 and the triangularity varies from - 0.7 to 0.8 [2]. Figure 3 shows a set of equilibria achieved during the first 2 years of TCV operation, illustrating the flexibility already obtained. These equilibria are reconstructed from the magnetic measurements using the inverse equilibrium solver LIUQE [3]. Subsequent operation of TCV will require additional high-power ECRH, up to 4.5 MW. The control of additionally heated plasmas will be carried out using the equipment and techniques described in this paper but is not the object of any detailed comments.

The flexibility of TCV has led us to develop specialised solutions for the feedback control of the TCV plasmas, both for the hardware, discussed in Section II and for different control algorithms, discussed in Sections VI and VII. Figure 2 shows the diagnostic measurements available for plasma shape control, namely a set of poloidal field pick-up coils inside the vessel and a set of poloidal flux loops placed on the outer surface of the vacuum vessel. These magnetic diagnostics are discussed in detail in Section III. The control of the plasma shape requires high-power rectifiers and a turbo-generator, described in

Section IV. The whole of the TCV plant requires supervisory control, timing and protection, briefly mentioned in Section V. The control of the unstable vertical movement of TCV plasmas has elicited considerable attention, both experimental results and their modelling are reviewed in Section VIII. Breakdown of plasmas with 6-10 Volts per turn has required reasonably structured stray fields, briefly mentioned in Section IX. Section X includes some recent results on bulk plasma control, limited, without additional heating, to the control of the plasma density. Section XI concludes with some final remarks. All quantities in this paper are in MKSA units unless stated otherwise.

II. CONTROL IMPLEMENTATION ON TCV

II.A General Concepts

The complexity of the TCV poloidal field system, with 18 separately powered supplies, together with the flexibility of operation imposed by the TCV experimental program make a generalised approach to plasma control essential. The hardware, described in Section II.B was designed to be independent of any particular control algorithms. The only restriction was that the control algorithms had to be locally linear. In order that this not be a restriction on ultimate flexibility, a piece-wise linear system with over 1000 possible segments was decided upon. In practice, the ability to change the local linearity during the pulse has never been necessary, except for specific changes to the feedback gains to cut the vertical feedback, Section VIII.C, or for breakdown. This is at first sight surprising, as the relationship between the plasma equilibrium and the control currents and diagnostic measurements are intrinsically non-linear. Approaches using non-linear observers were developed for implementation on TCV [4,5] but have not yet been operated.

To understand this change of approach, we consider splitting the control problem into 4 parts:

- In the first part, we define a set of control parameters which we wish to fix precisely during the discharge. It is necessary to define a set of parameters which can be independently determined by the available actuators in the presence of finite signal noise and plasma noise, requiring that the locally linearised mapping between control parameters and coil currents must be well conditioned. In addition, we have the choice of an underconstrained number of control parameters, fewer than the available actuators, an exactly constrained or an over constrained number. In the first case we required additional constraints somewhere in the system. In the latter case we must take care not to impose conflicting integral feedback terms and allow the weighting of the controller matrix to resolve the conflicts between the incompatible control parameters.

- In the second part, we must decide what the pre-programmed reference

values should be for these control parameters.

- In a third part, we must conceive a controller capable of making changes to the currents in the poloidal field coils such that the errors between the reference signals and the observers be as close to zero as possible.

- In the fourth part, we must construct the signals to be applied to the actuators, in our case the voltages to be applied to the shaping coils and to the gas valve, in order to achieve the zero error in an acceptable time.

Such a split is implicit in the design of any control loop. The non-linearity of the system to be controlled can be incorporated in several ways. We could define control parameters, such as elongation and triangularity, which are intuitive to pre-program. However the observers to estimate these variables are non-linearly related to the detected signals from the tokamak and we therefore need to design non-linear observers. Neural Networks [4] or Function Parametrisation provide useful ways of implementing such non-linear mappings. However in this case the controller itself will also have to be non-linear. TCV provides a simple demonstration of this due to its extreme flexibility. Plasmas at the bottom of the vessel and at the top of the vessel have to have opposite signs in some of the controller terms for imposing a given change in triangularity. A particular coil may need to "pull" or "push" to increase triangularity, depending on the vertical position of the plasma relative to that coil. Any linear controller would therefore inevitably have to be switched as a function of the equilibrium parameters if we were to use such non-linear observers. However if we choose a set of observers which are linear combinations of the poloidal fields and fluxes, then we find in practice that the controller can be "generic". This is simpler, but of course we pay the price that the observers are not particularly intuitive. On the other hand if the observers are linear, then any arbitrary combination of these observers can be used to achieve an identical solution to the control problem; the observers and controller can be arbitrarily post- and pre-multiplied by any invertible matrix. Two approaches using linear observers and constant linear controllers have been implemented on TCV and are described in detail in Sections VI and VII.

The timescale of TCV operation, requiring us to stabilise plasmas with growth rates of the unstable vertical position up to 3000 s^{-1}, imposes an additional speed constraint on the control algorithms. We wanted to incorporate the control of fast internal vertical position control coils, without having to redesign or add to the initial system.

Yet another constraint was that the control system should be apparently simple for implementing simple control algorithms, such as were used when obtaining the first TCV plasmas. A trivial algorithm for creating and controlling circular plasmas was easily implemented in spite of the apparent complexity of the overall design.

During the conception of TCV, there was considerable debate on the controller algorithms to be implemented. The advantages of the more intuitive

generalised, i.e. fully matricial and not diagonal, Proportional-Integral-Derivative (PID) controller were compared with state control for which tools exist for developing optimal controllers. Although state control algorithms were simulated for flux control inside the TCV vessel, all plasmas have so far been produced using the generalised PID structure. Pragmatically, the generalised PID is working quite well up to the operational limits of TCV, which are currently due to the power supply switching times.

A final constraint was a desire to collaborate on the Plasma Control Systems (PCS) for the C-MOD and TCV tokamaks, under construction at the same time. It was therefore necessary that all hardware and software implementation be independent of the configuration of the machine. This additional constraint has led to a data-driven approach to the implementation of the plasma control software which has in fact been a driving force in generalising the control system hardware and software, to the benefit of both groups. By data-driven we mean that the definition of the control system functionality is contained in the data without implicit assumptions on the hardware to be controlled or the hardware used to perform the control. TCV and C-MOD only differ in the contents of their data structures. Both the software and the hardware on the two tokamaks are inter-operable to the extent that they are directly importable.

Although the Plasma Control System is mostly linked to the control of the poloidal field system for plasma shaping, the same hardware and software are used for the control of the density and the toroidal field. This will be extended, without structural modifications, to the control of the additional heating ECRH power, to the steering of the ECRH mirrors and ultimately to controlling the plasma beta under feedback.

II.B Hardware Implementation

II.B.a. Design Considerations

Figure 4 shows a schematic of the PCS hardware. 128 raw input signals must be converted to a reduced number of control parameters which are compared with reference inputs. The error signals must then be fed to a dynamic controller which will feed in turn the 18 poloidal power supplies, gas valve, fast vertical control supply and so on. Without even discussing the control strategies, we see that there will be many parallel calculations, mostly using matrix arithmetic.

Three approaches to effecting these calculations were considered initially:
- purely digital,
- a slow digital part and a fast analog part ,
- purely analog under digital control.

A purely digital system would have had a problem of speed if we had wished to perform all the control operations with it. It also would produce an

inherent cycle-time delay which induces undesired phase-lags in the controller. Its advantage is that it can perform any calculation we should desire, by "simply" re-programming. This point is readdressed in Section II.D.

The part-digital part-analog system has two disadvantages. Firstly, it requires the development of both approaches. Secondly, it is not obvious, for highly shaped plasmas, that we can simply adopt a single fast displacement mode to control the intrinsically unstable movement, as is frequently done for weakly shaped plasmas.

The purely analog system under digital control is a more complex design, and is limited to a fixed calculation, namely matrix multiplication.

We decided to opt for the third solution, restricting the calculations to matrix operations performed by DAC-multipliers and have developed the Hybrid Matrix Multiplier, described in the remainder of this Section. The full system is used to control both the C-MOD tokamak and the TCV tokamak.

The system is under computer control for loading the matrix multiplication weights. We decided to use the BITBUS fieldbus to communicate between the microVAX computers and the Hybrid Matrix Multiplier between shots, as this is the standard field bus for TCV, Section V.A.b. The management of the complete system is under MDS+, Section II.C.

II.B.b System Modularity

The A-Matrix (Fig. 4) has sizes of (96 inputs x 16 outputs) for C-MOD and (128 x 24) for TCV. The controller contains matrices of at least (16 x 16) and (24 x 24) respectively. A suitable modularity was chosen as follows, illustrated in Fig. 5. There is a Base Module of size (16 x 24). Base Modules are connected 3 per full 19" crate to make a (48 x 24) matrix. Crates can be connected to make (96 x 24), (144 x 24) matrices etc. Each Base Module consists of a group of up to 6 Multiplier cards, each performing a (16 x 4) multiplication. Inserting fewer than 6 cards, we can produce smaller matrices, giving (96 x 16) using 6 Base Modules with 4 cards each, as an example.

Since the final system has many thousands of matrix weights, considerable attention was paid to the interconnection cabling. The Multiplier cards communicate via a crate backplane, as do the Base Modules. All analog Input/Output signals are connected by standard 37-pole 'D'-connectors.

II.B.c Hybrid Matrix Multiplier Construction

Full technical details of the Hybrid Matrix Multiplier hardware can be found in [6].

The Base-Cell is designed around the AD (Analog Devices) 664 4-quadrant quad DAC loaded with 4 12-bit digital gains. If an analog input signal is applied to the reference input, the chip provides a simple digitally controlled gain. The

main drawback of this chip is a high frequency breakthrough at zero gain, leading to errors exceeding an LSB. We found that the whole batch of DACs could be adequately modelled and compensated for using fixed components.

The (16 x 4) modularity requires 16 quad DACs on a card with their outputs summed. The summing amplifier has a gain of 0.25 to reduce the probability of overflow. The 16 inputs and 4 outputs are obtained from the analog backplane and sent to a summation bus on the same backplane respectively. The digital input is loaded from a latch given a strobe signal. The latch is loaded by software in an INTEL 8044 microcontroller onboard each card, controlled by BITBUS. Each Multiplier Card also has its own memory which can contain up to 490 or 1900 matrices each of size (16 x 4 words). The Multiplier Cards are therefore autonomous but can be digitally controlled to set up a bank of possible matrices which must then be loaded into the DACs themselves in real-time.

The Base Module comprises up to 6 cards feeding up to 6 groups of 4 analog summing lines. A summing card in the Base Module sums the outputs from up to 3 Base Modules, and also adds in an "external" summing input for chaining crates or for inputting reference signals, perturbation signals or pre-programming signals. Where several Base Modules are connected in parallel in a crate, the additional modules have 16 inputs, but no outputs. A 19" Crate houses analog and digital power supplies, fan-cooling, analog and digital backplanes and the BITBUS Input/Output connectors.

The BITBUS fieldbus is adequate for loading the weights between pulses, but is too slow for the real-time control of these matrices. We therefore incorporated a parallel address bus (Private Bus) which is distributed to all cards in one matrix. The Private Bus defines the number of the matrix to be loaded, and each microcontroller then loads its latches. When all are ready, a strobe causes all matrix weights to switch simultaneously. The Private Bus is itself controlled by a separate BITBUS controller which can run simple real-time control algorithms in the future. The Private Bus also allows us the freedom to perform "intelligent" switching by a more powerful computer such as a PC, should the need arise. Provision of a separate external bus is a useful feature for future flexibility.

II.B.d Performance

The digital performance is defined by the matrix update time, measured at 1.3 msec. This is close to the TCV thyristor firing rate and is more than adequate for both TCV and C-MOD tokamaks.

The analog performance is limited by several factors, both DC and AC. The DC errors are:

- output offset, less than 1 LSB and correctable,
- zero gain breakthrough, up to 30 mV if all inputs are at +10V
- zero input, gain = 1.0 gives a non-zero output of a few mV.

The AC errors are not limited by the DAC bandwidth, the multiplier cards

being filtered at 50 kHz to eliminate switching transients as well as amplifier noise. The AC limitation is a capacitive breakthrough remaining even after compensation, when the phase also starts to shift.

A final performance figure is the time to load new matrices between discharges. This is limited by the BITBUS throughput, and we measure a time of 7 msec per matrix per multiplier card. Since each card can store many hundreds of matrices, we do not incur a large loading requirement between most discharges.

The full system performance corresponds to the order of 5200 digital multiply-and-add operations per cycle, trivially parallelisable. A bandwidth of 1 kHz would therefore entail 30×10^6 such operations per second with an input/output rate of 10^6 per second. Although this digital performance was not available when the TCV PCS was conceived, cheap parallel processors such as the SNAP [7], offering 2.5 GFLOPS and a high input/output rate would provide a simply implemented alternative capable of running any tokamak. This available computing power would also allow us to perform inverse equilibrium calculations in real-time as part of the feedback loop, especially for ITER with fewer control parameters and actuators. At present, a TCV inverse equilibrium calculation requires 100 MFlop.

II.C Waveform Generation

TCV requires for a large number of analog channels for pre-programming the control parameter reference waveforms, power supply voltage and current reference waveforms and waveforms for injecting perturbations into the control system for studying its dynamic performance. We have therefore implemented 128 such channels in 2 units of 64 channels with 1 msec resolution 12-bit analog waveform generation using a variant of the matrix multiplier card. Interpolation between neighbouring points is linear. An important feature of the generation of real-time waveforms is that we can choose between waveforms which are pre-programmed in absolute values, or waveforms which are pre-programmed proportional to the plasma current. The PCG algorithm, Section VI, uses the latter method whereas the MGAMS method, Section VII, currently uses the former method. The incorporation of the waveform generators into the control system is shown in Fig. 4 and further technical details can be found in [6].

II.D Software Implementation

The software for PCS had several design requirements:

- All data to create a discharge must be archived with the data for the discharge itself
- Any discharge must be replayable using the archived data
- There should be no user-specific data at the lower levels of the data structures, enabling the implementation of any algorithm on top of the

lower layers.

- It should be interfaced simply to the MDS+ data system used for the storage of all TCV data, of which more details are found in [8].

The matrices and the waveform generators are well structured devices. We define the matrices in MKSA units, together with the names of the inputs needed, the names of the outputs created and the hardware scales of the outputs (e.g. 50 kA/Volt). The controller time-constants are defined in seconds. These definitions are generated by algorithm-specific procedures written in an interactive language. Screen-inspection and screen-editing of the MKSA matrices is provided. The waveforms for the control variable references, power supply pre-programming and perturbation injection are calculated by procedures and stored in MKSA units. These waveforms can be specified in absolute units or proportional to the plasma current, according to a user-defined property of the waveform itself. The waveforms also have upper and lower limit setting as additional properties. The waveforms themselves can be screen-edited as an option to, or after, procedural creation.

These creation and editing actions cause data to be stored under MDS+ in pre-defined structures, ready for use in a plasma discharge. To prepare for a discharge, the MKSA matrices must be scaled according to the current calibration of the diagnostic signals, which can be varied from discharge to discharge, and according to the calibration of the actuator demand signals. This procedure leaves matrices which are simple analog gains for processing the real-time analog data. This format of data is downloaded during the discharge preparation cycle. The waveforms are scaled implicitly using the matrix output hardware scales and the proportional/absolute flag. During the discharge preparation the feedback control channels which are active must be switched on. The mode of the power supply is also set, having been specified by the control algorithm software.

Separating the MKSA specifications from the device loading has gained us two advantages. Firstly, the user has no need to worry about how the feedback is implemented in hardware, providing data which are valid for analog or digital implementation. This level can be defined as that which contains all the information to create the discharge. Secondly, this clean separation allows us to replay discharges irrespective of the hardware evolution since the original discharge.

This part of PCS has proven to be well-conceived, partly since the collaboration between TCV and C-MOD forced us in the direction of a solution which is independent of the details of the tokamak construction. There have been few modifications in the light of experience and there are only a few lines of software corresponding to the evolution of the system.

On top of this general functionality layer, we have implemented a menu-driven layer for the creation of discharge scenarios for testing the PCS hardware, testing the power supplies, running fiducial discharges, exercising the power supplies, calibrating the magnetics diagnostic, and so on. These can be

added to at will to reduce the operational workload. The PCS hardware is automatically tested overnight to avoid morning delays, for example.

This philosophy makes no assumption about the current state of the tokamak. An optional incremental philosophy would create a contextual problem when replaying discharges. Although making any modification to the plasma discharge provokes the execution of about 1 minute of software and hardware re-preparation, we consider that the contextual independence is essential. An illustration of the success of this approach is that we have been able to run two concurrent physics programs with the discharges interleaved with no additional difficulty.

III. MAGNETIC DIAGNOSTICS

III.A Introduction

The unique flexibility of TCV has important implications on the required precision and scope of the magnetic measurement set used for both real time control of the plasma shape and discharge parameters, and the off-line reconstruction of the magnetic topology. A detailed description of the magnetic measurements available in TCV is given here, covering sensor geometry and location, their fabrication and calibration, the electronics providing the required signal processing, the operation of the diagnostic and its continuous monitoring, which taken together guarantee the quality of the measurements. A substantial section is dedicated to the effort made to quantify the precision of the measurements and to derive corrections which optimize the consistency of the entire measurement set. We conclude with an objective evaluation of the system, listing both positive and negative aspects and areas for improvement.

III.B Layout

Magnetic probes measuring the component of the poloidal field tangential to the vacuum vessel were mounted inside the vessel at 38 almost evenly spaced positions in 4 poloidal cross-sections (equally spaced toroidally). Their exact position and orientation are indicated in Fig. 2. They are numbered from the inside mid-plane, clockwise. The signals from two diametrically opposite sectors are acquired and used for the equilibrium reconstruction, averaged to cancel any toroidal asymmetries and also fed in real time to the plasma control system. Two toroidal arrays located on the inner and outer equatorial midplanes and consisting of 8 and 16 probes respectively, were installed for MHD mode analysis.

Poloidal flux loops are mounted on the outside of the vacuum vessel. Their position with respect to each of the 38 magnetic field probes was chosen to facilitate flux extrapolation towards the plasma boundary, by pairing magnetic probes and flux loops (Fig. 2). In addition, all shaping and ohmic transformer

coils (except coil A, which is inaccessible), were equipped with a poloidal flux loop (Fig. 2).

III.C DC Current transducers

The DC current transducers are based on the Hall effect. The magnetic field around the conductor is annulled by an additional winding around the Hall effect detectors. The current in the additional winding is measured using a precision resistance chosen to give a resistive drop of 10 Volts for the nominal maximum current. Error sources are the knowledge of this resistance and the quality of the geometrical arrangement of the current-carrying conductors and the additional winding. These current transducers are placed in the rectifier assemblies themselves, to be part of the deliverable equipment, whereas ideally they would be placed in a region where the bus-bar geometry is simpler. The toroidal supply presents the additional complication of providing 2 parallel sources of current, Section IV, and the current transducer outputs are summed in the rectifier control electronics. The precision of these measurements is considered to be of the order of 0.25% of the full range, dominated by the constructional accuracy of the additional winding.

III.D Poloidal flux loops

Vessel flux loops are made of one turn of 1 mm diameter THERMOCOAX®, mineral insulated coaxial wire, compatible with the 400°C baking capability of the vacuum vessel. Coil flux loops are made of one turn of a conducting adhesive tape placed directly on the coil.

Only 28 of the 38 vessel flux loops are perfect, constant-radius windings since some of them must bypass the ports. These deviations were kept short and distributed equally either side of the ideal path where possible. Because of the presence of large rectangular ports on the top and bottom of the vessel, there are 6 incomplete, or partial flux loops, which measure the flux difference between neighbouring loops rather than the full poloidal flux.

Only the coil loops and a reduced set of 10 complete vessel loops are used to measure the full poloidal flux. The signals are dominated by the ohmic transformer flux and in order to maintain a good dynamic range in the amplifying electronics and data acquisition, the inner midplane vessel loop signal is subtracted from all vessel loops, with the exception of the 6 partial loops.

From the constructional data, the exact area of each loop segment can be evaluated and from the assumption that the normal component of the magnetic field is constant between two adjacent loops, coefficients can be derived which allow us to compute the flux measured by each loop as a linear combination of ideal fluxes as they would have been measured by perfect loops. These coefficients are contained in a square matrix T_{rf}. A better estimate of these

matrices based on real measurements is developed in Section III.H.

III.E Magnetic probes

The design of the magnetic probes was subject to 5 technical constraints:

1) They should be bakeable to 400°C.
2) The maximum thickness of the protective tiles was limited to 24 mm for a variety of reasons (wall stabilisation, maximum plasma volume, etc). The minimum tile thickness estimated from thermal load calculations was 12 mm, leaving 12 mm in which to house the magnetic probes.
3) The minimum required signal level implied a total probe area of a few hundred cm^2. (The expected signal level is treated in the following Section).
4) The signal cables of a 38 probe array has to fit in 4 signal feedthroughs. Thus the corresponding 20 cable torch had to be small enough to run behind the tiles.
5) The frequency response of the probes should be adequate (a few times 10 kHz) to allow the same probe to be used for MHD activity studies. The probe frequency characteristics are described in Section III.F.

The selected solution is a moulded ceramic core on which two layers are wound, each of 20 turns. The wire is 0.5 mm diameter THERMOCOAX® with a copper central conductor and an Inconel shield. The wire is continuous from the probe to outside the vessel, so that no connections are required in vacuum and the vacuum-tight shield does not need to be broken.

The magnetic field probes were designed to give measurements with a precision of a fraction of 1%. To obtain this overall precision, each element must be calibrated with a precision of the order of 1‰. Since the probe effective area could not be guaranteed to lie within these limits, the probe areas had to be calibrated individually and were found to vary from 346 cm^2 to 366 cm^2.

The measured frequency response up to 16 kHz cannot be reduced to a simple first order transfer function. Therefore, the equivalent probe model was chosen to include the cable capacitance and the shielding. The transfer function between magnetic field and probe voltage is then given by:

$$H_p = \frac{U_b}{-sB} = \frac{A_p}{1+(\frac{L_s}{R_s}+R_p C_p)+s^2 R_p C_p(\frac{L_p}{R_p}+\frac{L_s}{R_s})}$$

where A_p is the probe area, L_p, R_p and L_S, R_S are the self-inductance and resistance of the probe and its shielding respectively and C_p is the total cable capacitance. The measured data points are consequently fitted to a transfer

function of the form: $H_p = \frac{A_p}{(1+s\tau_1+s^2\tau_2^2)}$, to extract the best estimate of the probe area A_p, the probe cut-off frequency $1/(2\pi\tau_1)$ and the second characteristic frequency $1/(2\pi\tau_2)$. Typical values for these frequencies are 100 kHz and 300 kHz respectively. The statistical error on the estimation of the probe area by this fitting is typically 0.05%.

III.F Electronic Signal processing

Figure 6 shows the general signal processing block diagram for each type of signal. Signals from the LEM modules measuring the coil currents are simply amplified. The flux loop signals are either amplified to be used as loop voltages or integrated to obtain the poloidal flux. Magnetic probe signals are preamplified near the machine, before the long transmission line. They are also either amplified to be used as the time derivative of the field or integrated to obtain the field itself. Some magnetic probe signals are also attenuated to study fast events such as disruptions. Each signal is low-pass filtered before being acquired or fed to the control system. Each of the elements in the chain is now briefly described.

The signal fed to the amplifier passes through a first-order low pass filter (f_C = 80 kHz) which is later compensated, and a programmable amplifier with gains of $^1/_4$, $^1/_2$, 1 or 2. It is amplified by a second stage with a gain of 1 or 16, giving an overall gain of $^1/_4$ to 32. The amplifiers have a common mode rejection of 40 V, automatic offset compensation and saturation detection that can be read after a tokamak discharge.

The integrator signal is low-pass filtered (first order, f_C = 3.5 kHz) and amplified (gain = $^1/_4$, $^1/_2$, 1 or 2) before being actively integrated with a time constant of 200 ms or 12.5 ms, giving an overall integration time constant from 6.25 ms to 800 ms. The feedback circuit of the integrator stage compensates the low-pass input filter and the overall transfer function is that of a perfect integrator. The integrators have the same offset compensation, saturation detection and common mode rejection characteristics as the amplifiers. The integrators are not gated as conventionally done. The output drift compensation is turned off during the plasma pulse, allowing the circuit to integrate the input. The toroidal field pulse is included in this active integration window.

The magnetic probe signal amplification needs particular care because of the large frequency range and wide variety of signals that they detect. The low frequency signal components that come from the coil currents or the plasma current are of the order of 10 mV. Mirnov oscillations with a typical amplitude of a few percent of the poloidal field at 10 kHz give a 10 V signal. In addition, fast MHD events have been observed with spikes up to 300 V. Therefore, in order to prevent saturation of the preamplifier, but maintain a reasonable signal level, it was necessary to precede the preamplifier by a first-order passive filter

with a cut-off pole at 100 Hz and a zero at 3.5 kHz, giving an attenuation factor of 35 for the high frequency signals. The transfer function of this filter is combined with the input stage filter of the integrator and the integrating stage itself to give a global transfer function like that of a perfect integrator, as seen in Fig. 7. Differences in zero-pole cancellation produce no error in steady state conditions and a relative error in the signal at most equal to the relative error in the zero-pole pair divided by the ratio between the integration time constant and that of the implied pair. This ratio is typically 20.

We use second-order Chebychev, 0.1dB ripple, low-pass filter cells with a programmable cut-off frequency from 1 kHz to 25 kHz. Anti-aliasing filters consist of two cells set to half the sampling frequency. The group delay is ~0.7 samples at low frequency, increasing to almost 2 samples at half the sampling frequency. The signals used for real time shape and position control are processed by a single filter cell to minimize the phase lag.

III.G Operation

All preamplifier, amplifier and integrator gains, as well as the filter cut-off frequencies are remotely controlled via BITBUS. The same system monitors and archives possible saturations in the amplifying chain so that machine operators are warned of inadequate gain settings and invalid acquired data is flagged. This represents an overall set of about 2000 variables to be controlled and monitored.

The whole detection chain of each signal, from the probe, cabling, electronics to the acquisition and the real time control system, is checked every morning with a specially designed 2 second test. This is done by consecutively powering each ohmic transformer, shaping coil and the toroidal field coil set. The recorded magnetic signals are shown in Fig. 8. The expected poloidal flux or magnetic field for each loop or probe (based on the coil current) is compared with the measured signal in order to detect any faults in the diagnostic. The coil current evolutions are slow enough that vessel eddy currents can be neglected.

III.H Consistency and precision

The degree of consistency in the magnetic measurements was investigated with experiments in which each coil was in turn powered by a constant current, thus providing 19 independent measurements for each magnetic signal. The theoretical values for poloidal flux and magnetic field, determined from the constructional dimensions and coil current, I_a, are $M_{fa}(R_f,Z_f,R_a,Z_a)I_a$ and $B_{ma}(R_m,Z_m,\alpha_m,R_a,Z_a)I_a$ respectively. M_{fa} is the mutual inductance between the flux loop located at (R_f,Z_f) and the coil at (R_a,Z_a) and B_{ma} is the magnetic field per unit current measured by the probe at (R_m,Z_m) with orientation α_m. In this computation, the coils were split into individual turns (1 to 3 cm in size)

and the magnetic probe spatial extension was taken into account. Comparison of the theoretical values with the measured flux and field ψ_f and B_m indicates discrepancies of the order of 2% and 5% respectively (Fig. 9), which clearly exceed the expected, and tolerable, errors.

The explanation for these differences is to be found mainly in small calibration and positional errors. A relative calibration error ΔG_a in a coil current yields a flux and field error of $\Delta G_a M_{fa} I_a$ and $\Delta G_a B_{ma} I_a$, whilst a calibration error in the flux ΔG_f and field ΔG_m produces errors of $\Delta G_f M_{fa} I_a$ and $\Delta G_m B_{ma} I_a$. The differences produced by misplacement of coils or sensors take, for example in the case of a flux loop radial displacement ΔR_f, the form $\Delta R_f \partial_{R_f} M_{fa}(R_f,Z_f,R_a,Z_a) I_a$.

All these corrections can be adjusted to reproduce the observed differences. This may be achieved by minimising the residual errors in the flux and field:

$$\Delta\psi_f = \psi_f - M_{fa} I_a - (\Delta G_a M_{fa} I_a + \Delta G_f M_{fa} I_a + \Delta R_f \partial_{R_f} M_{fa} I_a + \ldots)$$
$$\Delta B_m = B_m - B_{ma} I_a - (\Delta G_a B_{ma} I_a + \Delta G_m B_{ma} I_a + \Delta R_m \partial_{R_m} B_{ma} I_a + \ldots)$$

by means of the cost function $J = \Sigma_{f,k}\, w_f^2\, \Delta\psi_{fk}^2 + \Sigma_{m,k}\, w_m^2\, \Delta B_{mk}^2$

where w_f and w_m are appropriate weights and the index k runs over the experiments. The deduced corrections remain small, a few % in the calibrations and a few mm in the positions. They exceeded, however, the expected tolerances and this can be understood by inspecting the covariance matrix (Fig. 10). This shows strong correlations, one of the most obvious being that a correction in a current calibration is compensated by an opposite correction in the flux and field measurements.

This effect can be diminished by choosing a cost function that favours small corrections, namely $J' = J + \Sigma_a\, w_{Ga}^2\, \Delta G_a^2 + \ldots + \Sigma_f\, w_{Rf}^2\, \Delta R_f^2 + \ldots$ where, the weights were chosen to be the inverse of the expected errors in the corrected parameters. All the corrections deduced with this cost function are now limited to a tolerable amplitude (Fig. 11), with only a slight increase in the flux and field residuals (0.5→0.8 mWb and 0.3→0.4 mT). The reduction in the correction correlations is significant (Fig. 10).

The linearity of the measurements was verified by repeating the experiments at half current. Only the coil current measurements show a small nonlinearity, which may be due to the presence of ferromagnetic elements in the LEM modules.

The original idea of this consistency study was to verify the position of each loop and probe using the fact that, as seen in Fig. 12, the isoflux surfaces from all the coils intersect with very different angles at each loop position. Thus the determination of the exact position by magnetic triangulation should be possible. An expanded view of the crossing region (Fig. 12) shows how a

geometric resolution of the problem would be difficult. The formal approach used here proved to be much more suitable and in addition allowed us to introduce corrections other than geometrical effects, such as calibration errors. For illustrative purposes, the corrected flux surfaces are also plotted in Fig. 12 to show the extent to which the crossing area has been reduced.

All flux loops that make detours around the ports were excluded from the above exercise, because obviously these detours would impair an exact determination of the loop position. The data recorded for these loops can however be used to experimentally derive the coefficients used to combine them into an ideal flux loop set (Section III.D). This is simply done by searching for the coefficients that give the combinations of the theoretical ideal loop fluxes as close as possible to the measured fluxes for all data points. The resulting coefficients, Fig. 13, are not far from a crude estimate based on the constructional data, but are certainly more precise.

The return bar and connections between toroidal field coils produce a poloidal magnetic field. The associated flux is detected by those flux loops that make a complete turn (Fig. 14). For the loops that make detours around the ports (open circles), either their complicated path and/or the fact that they may not average the toroidal field ripple to zero certainly produces a parasitic sensitivity to the toroidal magnetic field itself that must be compensated. The picture is less clear for the magnetic field measurements because the parasitic sensitivity to the toroidal field is much larger than the associated poloidal field. This is due to misalignment of the probe windings and to the separation at the vacuum feedthroughs of the otherwise twisted-pair cable. This however remains a small fraction of the effective field, ~ 1%.

III.I Plasma current estimator

Since no continuous Rogowski coil has been installed around the vacuum vessel, the total plasma current must be estimated with an appropriate combination of the poloidal magnetic field probes. The disadvantage of such a discrete measurement is that it may not be precise and is sensitive to currents flowing outside the integration path, especially the coil currents and vessel eddy currents.

A simple first choice for the linear combination coefficients can be made by defining a polygonal integration contour given by the probe location and performing a trapezoidal approximation of the integral. The error in the estimated enclosed current using this approximation is localised near the integration contour. In practice however, no significant plasma current flows in this region and the total plasma current estimator is excellent.

Sensitivity to coil currents is significant, especially for those coils near vertices of the polygonal integration contour, namely F1, F2, F7 and F8. This effect is however easily measured and compensated. The experimentally deduced compensation coefficients agree well with the constructional estimates

and lie in the range -0.2 to +0.8 A/A.

The influence of vessel eddy currents on the trapezoidal plasma current estimator has been computed theoretically for each up-down symmetric vessel eigenmode. This influence is small for low order eigenmodes, but peaks around mode numbers 38, 76 and ~114, for which the eigenmode current distribution coincides with the probe array discretisation, giving a resonant contribution. Fortunately these highly spatially structured current distributions are difficult to couple to and their excitation amplitude remains small and does not significantly perturb the plasma current estimation. The residual current error at breakdown is typically ~ 3 kA or 0.3% full range.

III.J Discussion

We choose to conclude this Section on magnetic measurements with an evaluation of the entire system, listing both positive features (⊕) that would be retained if TCV were to be rebuilt and those areas in which there is room for improvement (⊖).

⊕ The number and position of both the magnetic probe and flux loop measurements has been found to be adequate even though the extreme flexibility of TCV poses severe requirements on the system. We have not encountered any difficulties that would be simply resolved by an expanded set of magnetic measurements.

⊖ This positive statement is soured to some extent by the partial flux loops on the top and bottom of the vessel. Of course detours around ports are unavoidable, but the partial flux loops exacerbate the difficulties. In order to estimate the flux in the region covered by these loops, it is necessary to interpret exactly the signal of each partial loop and to sum them up to the point of interest, thus accumulating errors. Full turn loops even with large detours around the biggest ports would suffer much less from this problem.

⊖ Flux loops wound directly on the coils have not yet been of any use either for real time control or for equilibrium reconstruction. The problem is that the fast variation of the flux near the coils makes the measurement very sensitive to the exact loop position and to the conductor distribution within the coils.

⊖ The specifications on the precision of the coil current sensors were not sufficiently severe. The consequences of this lack of precision and linearity are problems when using this data, for example, in the equilibrium reconstruction or to map the magnetic field during breakdown. The solution in both cases has been to allow a certain uncertainty in this data and to use the large redundancy in the whole magnetic data set to reduce this uncertainty.

⊖ Major problems have been encountered with the somewhat fragile wire used to wind the magnetic probes. Firstly, the electrical isolation of the wire is guaranteed to only 250 V and as previously mentioned, very fast MHD events can give rise to voltage spikes exceeding 300 V. Secondly, the mechanical strength of this wire is incompatible with the severe constraints of a tokamak environment. Wires have been broken during the installation process and short circuits have appeared between the central conductor and the shield, which are

attributed to crushed wire in the bindings. To correct these problems, more robust replacement probes with a 1 mm diameter wire and a reduced area are under construction.

⊖ The magnetic probes have only been calibrated along their axis. This calibration should be extended to three spatial directions. This is a minor point that does not compromise the measurements, but the consistency analysis and the associated determination of the diagnostic quality would be easier if the parasitic sensitivity to a transverse field was known.

⊕ The precision and reliability of the amplifying electronics have proven to be excellent in all respects. No systematic failure has yet appeared that would jeopardise the use of the system.

⊕ A high level of consistency and precision in the magnetic data set has been reached: typically 0.5 mWb and 0.5 mT for the flux and magnetic field measurements respectively. Although a preliminary analysis indicated higher errors than expected, it has been possible to reduce these errors to an acceptable level by determining small corrections in the system calibration and in the coil, flux loop and magnetic probe locations, compatible with the known parameters of the system. The formal analysis method developed for this purpose yields a very precise knowledge of the magnetic measurement characteristics and a high degree of confidence in the data.

IV. POWER SUPPLY SYSTEMS

IV.A General Network and Outline Design

This Section describes the electrical network designed to supply the 19 coils and additional heating of the TCV tokamak, shown in Fig. 15. The power utility line has a distribution voltage of 20 kV, rated power of 40 MVA and short-circuit power of 130 MVA and is too weak to feed the TCV and a motor generator is therefore needed to deliver the peak power and energy. Because of the special requirements for the poloidal coil supplies, such as high frequency response and not to be limited by the characteristics of the power network, the TCV power supply system is designed with some innovative solutions. As a result, the specifications of the motor generator present some particular features, described in detail in Section IV.C. The 19 power rectifiers are described in detail in Section IV.B.

In order to achieve the maximum flexibility during the operation of TCV, it was decided to feed the individual coils with their own power supplies, bi-directional in current and voltage for the shaping and ohmic coils, and bi-directional only in voltage for the toroidal coil. This choice does not lead to the minimum possible installed power but is operationally simpler and presupposes no preferred operational mode or plasma configuration. The peak power demand normally occurs at the end of the discharge with maximum shaping currents and maximum current in the ohmic coils. Although the design of our previous TCA tokamak used switched resistors to provide the maximum

required voltage at breakdown by dissipating the ohmic coil current, the TCV tokamak uses full voltage control during all phases of the tokamak pulse. With this in-built flexibility, switching between configurations does not require any intervention on the power system hardware.

Two power sources are available. The motor generator is used as the principal power source during regular operation. In addition, the TCV network can use a test transformer directly connected to the grid via a circuit breaker.

Five pairs of incoming feeder isolators allow us to (dis)connect one of these two power sources to the five following groups of power supplies: toroidal, two ohmic, shaping high-field side and shaping low-field side. The toroidal and the two ohmic busbars each have one feeder isolator. The two shaping busbars each have eight feeder isolators. Three-phase cables connect the feeder isolators to the transformer primaries and subsequently the transformer secondaries to the rectifiers.

Each power supply consists of 12-pulse series or parallel thyristor rectifiers. In order to obtain the necessary phase shift of 30 degrees between the two groups of secondary voltages, the primary windings are of the "elongated delta" type. The active parts of each transformer are identical, only the primary connections differ to obtain the ±15 degrees of phase shift, so asymmetries due to manufacturing tolerances are reduced, compared with the usual coupling "delta-Y" and "delta-delta".

DC coaxial cables in parallel then connect the rectifier DC output busbars to the tokamak coils, except for the toroidal coil which is fed through copper plate busbars. The same type of cable is used to connect the coils to crowbars which protect both the tokamak coils and the rectifiers from overvoltages typically originating from plasma disruptions. The crowbars include several thyristors in anti-parallel, depending on the rectifier design current, in parallel with the coil and are triggered locally following over-voltage detection or remotely by a command signal from the rectifier electronics.

IV.B Choice of Rectifier Power Supplies and their Performance

Network switching thyristor rectifiers supplied with a 50 Hz voltage have their cut-off frequency below 30 Hz for full amplitude voltage variations, which is lower than that necessary given the vessel current decay timescale for controlling the unstable vertical plasma motion. Small variations can have a greater effective bandwidth. On the other hand, power supplies like transistors, choppers or GTO thyristors, which have much higher cut-off frequencies, are very expensive for the required installed power. Therefore, the best technical and economic compromise was considered to be the use of network switching thyristors supplied with as high an AC voltage as possible.

All poloidal and toroidal power supplies on TCV are based on thyristor semi-conductors which do not need sophisticated driving hardware and are fast

enough for our application. The generator frequency was chosen to be as high as possible in order to minimise the dead time of the rectifiers. For the same reason, 12-pulse rectifiers are used throughout.

Our first requirement was to use only one type of phase control thyristor for all the rectifiers, with a medium Si plate size (50 mm) and a short turn off time resulting from the use of a higher frequency. These thyristors are assembled in a compact 6-pulse bridge, 3900A-1800V for pulses limited to 4 seconds every 5 minutes, used as the basic module for the assembled rectifiers and as spares in the case of thyristor failure. These bridges are water cooled and include fuses placed on the AC input phases which protect the equipment against thyristor faults.

Each 6-pulse bridge is fed by its own three-phase cable in order to distribute the power losses equally between the thyristors in parallel and to minimise the impedance between the transformers and the rectifiers. Table 4.1 summarises the characteristics of each rectifier type.

The toroidal coil rectifier consists of two assemblies, each one including ten 6-pulse bridges, connected in parallel by two air cored reactors. This power supply alone works only in 2 quadrants. The ohmic coil rectifiers consist of two assemblies connected in series, each one including eight 6-pulse bridges. The same is true for the shaping coil rectifiers with two assemblies, each one including two 6-pulse bridges.

All of the 18 poloidal coil power supplies work in 4 quadrants without a circulating current for several reasons. The turns ratio between the ohmic and shaping coils is not large enough to produce an overvoltage during zero crossing of the current when the circuit is open. The anti-parallel connection of both 2 quadrant systems is possible without decoupling coils. The transformers are also considerably simplified as described above. With provision for a circulating current, at least two separate secondaries would be necessary to feed a 6-pulse assembly. The control electronics are simpler and permit a fast direct control of the output voltage.

The design of the control electronics had to take into consideration several differences between this specific application and normal practice:

- the initial high working frequency
- the decreasing frequency (120 Hz to 96 Hz) during the pulse with a maximum assumed deceleration of 25 Hz/s.
- important voltage notches created by the thyristor switching in our case in which the generator only feeds power rectifiers.
- the firing algorithms have to drive the power equipment when fed at 50 Hz and at half the nominal voltage, during testing.

Because of the decreasing frequency, it was necessary to adapt the maximum firing angle (invertor condition) during the plasma discharge. Special requirements on the absolute phase angle accuracy and the constant spacing between firing pulses are needed to minimise the harmonics on the output

voltage, which could disturb the plasma control.

Three different internal regulation modes allow us to control these rectifiers remotely during tokamak operation. These modes are software selected and separately configurable for each power supply. Figure 16 shows three block diagrams corresponding to the three modes:

a) Current feedback is the usual method of driving such rectifiers. The current is controlled by the internal controller and limited at the same time. This mode is used for the control of the toroidal field supply and also during acceptance and maintenance tests.

b) Open loop voltage driving is used when there is no pre-programmed current value, but the plasma control system provides voltage demand signals for each power supply. In this case it is necessary to add a current limitation to avoid overcurrent. This mode is used in the MGAMS method, Section VII.

c) Hybrid mode is the combination of the two above methods. It allows us to pre-program a current demand signal and to override this pre-programmed current with external voltage demand signals. In this case the internal current control cannot have integral feedback which would create a conflict between the PCS and the internal power supply controller. This mode is used in the PCG control algorithm, Section VI.

Current work includes experiments to measure the effective transfer function of these industrial power supplies during normal operation. More technical details on the rectifier power supplies can be found in [9].

IV.C Motor Generator and Performance

Figure 17 shows the nominal total apparent, active and reactive power to be delivered by the generator. Hundreds of Hz and MVA might seem to present important technical problems for a motor generator. However, a frequency of 120 Hz can be delivered by a turbo-generator running at 3600 rpm with a 4 pole rotor. Because of their high pulse power, 50 MVA, the switching of the toroidal or ohmic rectifiers can generate large voltage transients. In order to reduce the interference between the rectifiers, the generator saturated subtransient short circuit power must be higher than 1000 MVA at maximum frequency. This leads to very large short circuit currents, particularly at the minimum allowed frequency.

The tokamak plasma is initiated by applying a loop voltage of up to 10 V/turn, by the maximum voltage inversion which can be generated with the ohmic rectifiers. This inversion produces a significant load change for the generator, seen in Fig. 17 at t = 1.4 secs which must not lead to excessive voltage fluctuations. During the entire pulse, the generator voltage must be kept within +/- 2% of its nominal value.

The nominal pulse shown in Fig. 17 corresponds to a spent energy of 100 MJ. By limiting the minimum working frequency to 80 % of the maximum, the stored energy of the generator must not be less than 300 MJ. This value is

usually available in the standard turbo-generators in the 50 to 100 MVA continuous power range. Therefore, the use of a flywheel was not necessary for the basic TCV performance, although an additional inertia providing an extra 150 MJ of stored energy will be required for the additional heating of TCV plasmas, with the correspondingly increased plasma discharge pulse length.

These requirements have been fulfilled by the TCV turbo-generator, derived from standard ABB air cooled turbo-generators. Its design is mainly influenced by the high short-circuit power requested and the very low voltage fluctuations which are allowed. Complying with these two requirements, which are interdependent, leads to an active volume which is largely sufficient to comply with the nominal required stored energy without an additional flywheel, at least for the TCV basic ohmic performance. Its characteristics are summarized in Table 4.2.

The whole generator plant is controlled and supervised by a fully electronic programmable system of a type currently used by ABB in power generation plants. The machine can be controlled from a local console and directly from the TCV control system through an RS232 interface. The plant status is available on the TCV control console. Operation of the turbo-generator is automatic during TCV operation, with no personnel required.

A frequency convertor starts the machine up to its maximum speed within 12 minutes and speeds it up from 80 % to 100 % in 3 minutes. During the speed-up phases, the convertor works at maximum and constant active power in order to reduce the reactive power consumption on the 20 kV mains. This convertor, supplied by two transformers of 1.2 MVA each, delivers 2500 V and 512 A at 120 Hz.

Special care must be taken to prevent excessive voltage perturbations on the local 50 Hz network, to conform with new rules applied in Switzerland, similar to those in force elsewhere. Despite the use of 12-pulse rectifiers for the excitation system and the starting convertor, the harmonics they can generate exceed the maximum values allowed. Consequently, a harmonic filter is connected to the 20 kV input to the generator plant and consists of four branches, filtering the 5th, 7th, 11th, 13th and 25th harmonics. Due to the poor short circuit power of the local network the voltage fluctuations generated by the excitation rectifier during the tokamak pulse, 6 to 7 MVA, are only just within the limits imposed.

More technical details on the design and construction of the TCV turbo-generator can be found in [10].

IV.D Fast Internal Coils

The vertical stabilisation of the most highly elongated TCV plasmas planned is not possible with the main poloidal power supply system and the external poloidal field coils just described, due to the relatively slow switching

time of these thyristor supplies compared with the growth rates of the plasmas which we wish to control. Growth rates up to 1000 s^{-1} have already been obtained in controlled discharges, but we wish to go as far as 3000 s^{-1} in the future. Consequently, a Fast Power Supply (FPS) will supply a low impedance Fast Internal Coil set (FIC) located inside the vacuum chamber to produce the fast radial magnetic field variations necessary to slow down the plasma vertical movement to within the controllability range of the slower power supplies. A pulsed 2000 A / 560 V IGBT switching mode power supply has been constructed and is under test for this purpose. The number of FIC turns and their location were fixed taking into account the vessel geometry and available power supply technology. The FIC comprises two toroidal three-turn coils located in the outer corners at the top and bottom of the vacuum vessel, Fig. 1. The windings are made of non-insulated copper bars, fixed with ceramic rings and protected from direct exposure to the plasma by graphite tiles. The turns are divided into quarter-turns for ease of installation and to leave open an option for the control of non-axisymmetric perturbations with four independent power supplies.

Considering that the slow thyristor power supplies can take over the plasma vertical control within a few milliseconds of a control excursion, a non-linear (switching mode) power supply was chosen and specified in terms of its voltage reaction time and current controllability rather than in terms of bandwidth and accuracy, according to Table 4.3.

The chosen structure of the FPS consists of two main elements, the DC power source which produces a regulated DC voltage and the Fast Output Stage which applies the fast voltage pulses to the load, Fig. 18. The DC power source comprises a fully controllable thyristor rectifier normally supplied by the TCV motor-generator. A DC chopper is added for better DC voltage control. The fast output stage comprises an H-Bridge IGBT invertor followed by small chokes to limit the output current di/dt. A thyristor bi-directional crowbar is added to protect the system in case of a major plasma disruption. The H-Bridge invertor is connected to the FIC with standard low voltage cables. The main parameters of the Fast Output Stage are listed in Table 4.4. More technical details of the FPS can be found in [11].

IV.E Operational Experience

The power requirements calculated during the TCV design phase were already shown in Fig. 17, while those resulting from a real TCV discharge (D-shaped, 700 kA, 1.5 sec. total, 0.8 sec. flattop) are shown in Fig. 19. Comparing these two sets of curves highlights the following main differences:

- The time chosen to ramp up the ohmic pre-magnetisation coil currents before the plasma breakdown is shorter in practice than in the design requirements. This choice reflects the use of the available voltage on these coils while at the same time reducing the total dissipated energy.
- In the design requirements, 20 MW were considered for additional plasma heating during the 0.5 second plasma current flattop. This power is not yet

delivered in the TCV discharge shown in Fig. 19. Despite this difference, the active power needed for the real discharge is much lower than planned in the design phase. As a consequence, the reactive power is higher for the real TCV discharge. The reason is partly due to the worst case analysis of the coil resistance when defining the power requirement during the plateau. The manufacturers managed to construct the toroidal and poloidal coils well within the imposed tolerances.

- The actual plasma duration is already longer than originally expected, without exceeding the nominal TCV flux swing (i.e. the nominal ohmic coil current). This is also due to a worst case specification of the volt-seconds requirement, reduced by the vessel conditioning techniques applied in TCV.
- The actual coil current ramp down time is shorter than planned. This is again a choice of using the available coil voltages.

The shorter ramp-up and ramp-down times as well as the lower active power are the result of one single fact: the nominal power of the rectifiers is much higher than required for the nominal TCV discharge. The maximum voltages that can be applied to the coils are therefore significantly higher than the voltages requested for the real discharge of Fig. 19. The actual coil currents being very similar to those which were planned, calculated directly from the plasma equilibria, the consequence for the rectifier AC network is a reduction of the active power and an increase of the reactive power, thus a reduction of the power factor.

The available voltage margin was intended to be used for plasma stabilisation and to provide a fast response during the plasma breakdown. The latter point is important in that a considerable fraction of the available volt-seconds can be lost if the plasma control is inadequate during the initial phase of the plasma current ramp-up. The power transients which are required during plasma breakdown do not differ significantly between the design values and the real discharges. Figure 19 shows that large active power and reactive power slew rates in excess of 300 MW/second and 250 MVAR/second are achieved as expected, an attractive property of the design of the TCV turbo-generator. On the other hand, the feedback regulating power which is available in the TCV power supply system does not seem to be required for the present plasma configurations.

Assessing the installed power required for feedback regulation is a delicate one during design, due to the questions of capital investment, but is also not simple to assess during operation. This regulating power is perhaps really necessary and used but is not visible in the real total power requirements because it appears during very short transients which have not been investigated up to now. Another more physical explanation could be that this power is just being extracted from one coil and delivered to another coil without a net effect on the total power consumed. This is realistic in that the total poloidal stored energy is not modified significantly by the feedback. In order to shed some light on this question, we have repeated a real TCV discharge with the power supplies in feedforward control to provide the same coil current evolution

without any plasma-generated control voltages. The discharge chosen was a particular one with repeated L-H transitions and ELMs, corresponding to a large feedback requirement to maintain the equilibrium, superimposed onto vertical height modulations which do not appear in the pre-programming. This type of discharge is further discussed in Section X. The traces showing the reactive and resistive powers for the pre-programmed discharge and the real discharge are shown in Fig. 20. Before and after the plasma current pulse, the two discharges are essentially identical. During breakdown there is a 20 MVAR difference in the reactive power requirement whereas during the plasma current itself the reactive and resistive powers are virtually identical. it therefore appears rather difficult to detect the required feedback power, even in a "noisy" discharge, due to the dominantly recycled power.

V. SUPERVISORY CONTROL AND PROTECTION

V.A The TCV Overall Control System

V.A.a Component Parts

The TCV tokamak requires many computerised functions: plant control, plasma control, data acquisition and data analysis. We have combined these functions into the Système Intégré de Données (SID) reflecting the fact that the data must be integrated as first priority, rather than just the hardware. In this way, the data which are generated as one part of SID must be simply available to all other parts. The overall philosophy was laid out in 1987, following which many of the fine details have been modified. The constraints of available manpower encouraged us to buy in as many components as possible. In fact the late 1980's saw a vast array of useable products on the market, compared with the late 1970's.

In spite of this intention to buy in as much as possible, two major developments were considered necessary. Firstly we completed the construction of the Plasma Control System hardware, Section II.B, and secondly we developed a concept of modular BITBUS controllers for TCV for which more complete technical details are available in [12].

The full SID is shown in schematic form in Fig. 21. A cluster of microVAX computers was chosen for its compatibility with available software. Two microVAXes are equipped with an open Q-BUS, for control and acquisition input/output. Other microVAXes and more recently DEC AXP computers are available for plant control, visualisation and data analysis. Within the TCV building, 14 ETHERNET segments and an FDDI ring, electrically isolated by fiber-optics, provide communication for all computers. We have deliberately chosen not to use different platforms for plant control and data acquisition, usually corresponding to the physics/engineering dichotomy, firstly since for an advanced tokamak with the flexibility of TCV such a separation

seemed artificial and secondly such a separation renders the free circulation of data more difficult. Although we have suffered from this imposed complexity, we have not regretted this unusual decision.

The timescales of tokamak operation separate nicely into slow reaction time, using an asynchronous network with a response of ~ 0.5 sec which is adequate for the plant control, Sections V.A.b and V.A.c, and fast, synchronous timing with a jitter of < 1 μsec, hardwired, used for plasma discharge control, Section V.A.e.

Although not part of the actual control system, all signals are monitored by the TCV data acquisition which is the MDS+ system [8]. The major source of data is CAMAC, with the exception of video data, handled separately.

V.A.b Plant Controllers

Plant control requires 4 components: plant control hardware, a field bus, a data store and visualisation software. In TCV the plant control hardware is relatively heterogeneous. The fieldbus selected was BITBUS which has the following advantages:

- rigidly defined protocol, immune RS485
- fast enough (375kBaud)
- multi-drop to reduce cabling costs (250 nodes/bus)
- automatic transmission checking and recovery
- protocol implemented in the INTEL 8044 microcontrollers
- INTEL 8044 is multi-task with up to 7 downloaded tasks
- VAX and PC interfaces available
- some modules are available commercially
- the basic node is cheap (~FRS 500.-)

So far, we have had to develop several I/O cards, including:

- A 24-line PIO card with 3 timer counters for pulse generation.
- A 16 output 12-bit DAC card with internal and external reference voltages, providing voltage generation and programmable gain amplifiers.
- A 16 input 12-bit ADC card providing cheap and simple data acquisition up to 20 kHz for one channel.
- A 12 channel 16-bit counter-timer.
- An RS232 UART to control any RS232 device.
- The amplifier, integrator and filter cards necessary for the magnetics diagnostic.

We have installed 2 BITBUS field buses, not dedicated to any particular equipment, and 170 modular controllers, including the Hybrid Multiplier. We are at present exploring the use of an Ethernet interface to BITBUS. Seven years of experience with this fieldbus have convinced us of its choice. More details of the plant hardware can be found in [12].

V.A.c Plant Control Software and Philosophy

The data store and plant visualisation are provided by the commercial VSYSTEM control system toolbox package [13]. This allows us to create live databases across the DECNET network and to display the plant state on X-terminals for both monitoring and active control at the control desk which comprises 5 colour terminals. Using this commercial package has reduced the programming effort to a fraction of a person's time. We currently have around 8600 data items (channels) in the TCV control database for controlling the whole of the TCV plant.

The plant data, covering the control of TCV from personnel access to the acquisition timing is a hierarchy-free flat structure, resident in memory in several microVAX computers. Access to this data is transparent whether it is for visualisation, modification of plant data or automation of parts of the plant. This flat structure can be considered as a software bus across the computer network in real time and constitutes a part of the data acquired with each plasma discharge to retain a complete snap-shot of the machine conditions. Connection between this live database and the plant hardware is carried out with generic operations on BITBUS whose closed protocol allows us to perform this communication based on information which is a property of the data channels themselves. There is therefore no device-specific code for device communication on TCV, which was one of the original design aims.

Using this structure, automation of the different sections of the plant has been carried out independently of any knowledge of the communication with the devices, another original design aim. Small pieces of readable code perform read/write operations on the live database which then has the responsibility of communication with the device. Not mixing up the device communication with the automation was yet another design aim which has been fully realised. More details on the plant control can be found in [12].

V.A.d Discharge Sequencing

By discharge sequencing we refer to all the synchronising activity necessary to run plasma discharges. The fast part, Section V.A.e, 20 seconds around the plasma discharge itself, is under fire-and-forget control with the exception of hardwired security and the operation of the feedback control loops. The sequencing therefore carries out the activities listed in Table 5.1 in turn, waiting for completion before continuing, assuming that the tokamak has been set up correctly by the operational software described in Section V.B. The full sequence takes about 2 minutes on TCV, limited by the speed-up of the turbo-generator and the initialisation of the data acquisition structures and devices for the next pulse.

Table 5.1 does not illustrate a state machine. Although individual systems, pumps, motor generator, power supplies, run with internal state machines, we have preferred a pure sequence-driven pulse cycle. Since the goal of the sequencer is trivially "all requested systems set up as requested", we saw no

need for the top level control being a formal state machine. The complexity of a TCV state machine would be enormous, the software requirement huge and the benefit zero, since the simple approach is adequate. One of the reasons for this conclusion is the very small degree of inter-dependence between subsystems combined with the cyclic nature of a tokamak. Equipment is controlled by this Cycle Sequencer by transition requested and the lower level control loops signal their state allowing the Cycle Sequencer to poll for completion.

An important point to note is that we have not predefined any relationship between the computer hardware and the work to be carried out. The work initiated by the Cycle Sequencer executes instructions to the live database, unaware of the physical location of the data in the database, treating it as a true software bus. The databases then perform actions on the control hardware independently of their purpose, using predefined data structures. As an example, the power supplies are controlled from one microVAX computer as far as the live database is concerned. The supplies and the generator are controlled by 4 BITBUS controllers which in turn control 20 local PLCs on the hardware itself. The essential point is that these relationships are not explicitly defined anywhere in the structure of the plant control - they are irrelevant at all except local levels.

V.A.e Plasma Control Timing and Discharge Synchronisation

Once we accept that the ETHERNET and fieldbus cannot guarantee the jitter necessary for plasma discharge control, we require a separate timing bus. This timing system is distributed throughout the TCV building and provides synchronisation for 8 seconds before the pulse and 5 seconds after the pulse. Outside this "real-time" window the discharge sequencer already described provides the necessary control.

We decided that many simple cheap modules fed from a master clock were required and that remote-to-remote timing communication would be limited to a few (~5) channels. Our consequently simple fan-out timing system comprises a system-clock running at 1 MHz providing overall synchronisation. The master clock unit sends out a variable frequency clock train which the remote counters count down to provide local absolute timing. This clock provides a resolution during the plasma pulse of 50 μsec with a precision of better than 1 μsec. Before and after the plasma pulse the resolution is 2 msec with the same precision. The variation in the electronic delays, compounded with the different cable lengths in the TCV building have been measured. The point-of-delivery delay compared with the master clock generation ranges from 260 nsec to 790 nsec. The BITBUS 12-channel timing cards are programmable to provide not only a leading edge when triggered by the counter, but can also generate square waves for gating the power supplies and integrators or for providing strobes.

V.B TCV Operational Software

In the previous section we have described the hardware and software

capable of executing the work necessary to perform tokamak discharges. The only remaining level is that closest to the experimentalist who has to define the parameters of the discharge to be run. We treat this as an offline task, although performed in "live" operational time and this task is carried out as a sequential set of procedures illustrated in Table 5.2. An important property of this approach is that the preexistant state of the tokamak is considered to be unknown. Therefore all actions are carried out in order to produce the discharge conditions required, even if this is often mostly redundant. The only exception to this is the request to repeat the last discharge without modification.

Both MDS+ and VSYSTEM have been interfaced to the IDL and MATLAB symbolic data manipulation packages. The preparation of the TCV discharge is generated by IDL code which is extremely simple to modify on the fly. Of course, this itself leads to problems of software control and inevitably errors, mentioned in Section V.C.

V.C Protection

The dominant philosophy of protection for TCV is that the operating devices, software and communication are as fallible as human operation and communication. We therefore assume that demands will inevitably be sent to equipment which would endanger it. All equipment is therefore responsible at the local level for its own protection. This relaxes the requirements on the quality control of the computer-level software and hardware.

One particular action cannot be delegated to the individual devices and that is the electromagnetic protection of the tokamak structure itself, due to the forces between the different coil systems. TCV can be seriously damaged by operation of the individual power supplies within their individual limits. Such operation can not correspond to normal plasma operation and we have constructed a "last resort" safety system which monitors all the poloidal coil currents, calculates the dominant forces on the coils and causes an alarm if these calculated forces exceed predefined limits. This type of alarm is fanned in over the TCV network and causes an immediate (15 μsec) trip of the power source breaker. This trip has never in fact been used in real operational conditions, yet, but allows us to maintain the philosophy we have established.

Since current in earth-loops can provoke errors in the diagnostic signals as well as in the protection logic, considerable attention was dedicated to detecting earth-loops during tokamak operation. Each branch of the "star" configuration of the earth layout is equipped with an injection of a continuous different audio frequency voltage. The current at these frequencies picked up on the other branches indicates the matrix of impedances between the different branches. These values are monitored regularly and if any sudden change should be detected, the next pulse is inhibited. The system is suppressed during the plasma pulse to avoid polluting the magnetic diagnostic.

A disruption detector, sampling the plasma current feedback error and its

derivative during the pulse, cuts all the feedback loops after a disruption. This is due to the possible presence of errors in the magnetic signals during the fast transients following a disruption and inhibits the Plasma Control System from sending inconsistent voltage demand signals to the poloidal field coil power supplies.

The only other safety action which is not delegated to the normal plant operation is personnel access safety. A hard-wired interlock system and access counter inhibits the closure of the two main series power supply contact breakers if all personnel have not left the TCV zone.

Although not part of the plasma-related protection, the high short circuit power of the generator at the minimal and unusual frequency of 90 Hz, higher than 1500 MVA, requires a short-circuit current limiter explosive fuse to protect the generator and the 10 kV network up to the primary of the rectifier transformers from the consequences of a short circuit. A mechanical breaker permits the interruption of electrical faults occurring on the load side of the secondaries of the rectifier transformers and can interrupt fault currents up to twice the generator nominal current.

VI. PCG CONTROL ALGORITHM

VI.A Outline of the Approach

The Poloidal Command Generation (PCG) algorithm [14] is one of a class of algorithms which relies on the existence of a database of achieved or modelled tokamak equilibria rather than a priori relationships between the diagnostic signals and the equilibrium parameters to be controlled.

In the PCG approach several plasma shape parameters, with the exception of the vertical position and the plasma current, are controlled using a set of heuristic combinations of the poloidal flux and poloidal field measurements. Initial TCV operation which started with circular plasmas was extended to shaped plasmas using these heuristic control variables and enabled us to cover a significant range of shaping operation, namely $20<I_p(kA)<500$, $-0.6<\delta<0.55$, $0.8<\kappa<1.74$, $-0.15<z(m)<0.4$, $2<q_a<22$ for almost up-down symmetric plasmas. The reference waveforms for these first discharges were programmed using an intuitive understanding of the relationship between the shapes and the fluxes and fields. This approach is feasible with a small number of controlled parameters but was tiresome, motivating the automation of this generation of the reference waveforms corresponding to the desired equilibrium parameters.

A TCV achieved equilibria database was generated from this initial "manual" tokamak operation, producing a knowledge database of the relationship between the reconstructed equilibrium parameters and the poloidal field and flux measurements. Each equilibrium possesses a complete vector of

diagnostic measurements. Since the heuristic control variables are themselves linear functions of the same poloidal fields and fluxes, each reconstructed equilibrium possesses a vector of values of the controlled variables, whatever their definition should be. In order to generalise these point relationships in the Input:Output space of Equilibrium-Parameters:Control-Parameters, a non-linear mapping is required, since the plasma equilibrium is itself non-linearly related to the poloidal fields and fluxes. We have thereby succeeded in removing the non-linearity from the observers and the controller, leaving it in the design of the reference waveforms and in the choice of the control parameters. The chosen mapping was the Multi-Layer Perceptron type of Neural Network [4,5] with a configuration of, for example, 5 input equilibrium parameters, 7 hidden-layer neurons and 5 output reference waveforms. The Neural Network therefore maps the desired equilibrium parameters onto the required reference waveforms according to the chosen set of shape observers.

These control parameters, which are linearly related to the poloidal fields and fluxes, are empirically found to be almost linearly related to the currents in the poloidal field coils, over the database used. This unsurprising simplification allows us to generate pre-programmed current waveforms according to this linear mapping, generated by a Singular Value Decomposition of the database examples of the relationship Control-Parameters:Poloidal-Currents. Applying a model of TCV which neglects both the plasma current and the vacuum vessel currents generates approximate pre-programmed voltage demand signals for the power supplies, operating in the Hybrid Mode, Section IV.B. Since the necessary precision of the pre-programmed current and voltage waveforms is only that which is enough to reduce significantly the demand on the feedback controller, the simplifications used in their generation are found to be adequate operationally.

The use of the power supplies in the Hybrid Mode is important for the PCG algorithm as we are producing discharges using 18 poloidal field coil supplies while only feeding back on typically 5 control variables. The current pre-programming of the power supplies provides 18 additional constraints which cause the coil currents to be uniquely determined. The 18 poloidal currents which satisfy the 5 integral feedback constraints are those which have the least weighted "error" in the power supply current waveforms. Such underconstrained control was easily implemented. We obtain an additional advantage in that we can open the feedback control loops, typically done immediately on detecting a plasma current disruption, while forcing the coil currents to remain close to the pre-programmed waveforms. Using only voltage control, the currents become unconstrained when opening the feedback loops, limited only by the power supply limit settings.

Having achieved the programming of the references and the programming demand signals to the power supplies, all that remains is to provide a feedback controller to correct for the errors in the controlled variables. Section VI.B describes the heuristic control parameters used and the feedback algorithm implemented.

VI.B PCG Observers and the Feedback Algorithm

The observers used for initial operation were inboard-outboard flux balance for radial control and up-down flux balance for vertical control to obtain roughly centred plasmas. Addition of local poloidal field measurements allowed more accurate centering. In order to produce triangularity on the inboard and outboard sides of the plasma, collinearity measurements of the inboard and outboard flux surfaces were approximated by up-down symmetric combinations of flux imbalances between loops close to the mid-plane and loops above and below the plasma. These intuitive combinations gave good control of the plasma triangularity and their combination controlled the elongation.

We have implemented several plasma vertical position observers including simple up-down flux-loop differences, up-down extrapolated flux-loop differences and poloidal field moments. PCG currently uses an observer of the current-weighted vertical plasma position, $z(t)I_p(t)$, hereafter referred to as zI_p. We use the current-weighted plasma vertical position zI_p since it is linearly related to any applied poloidal flux variation. We approximate the desired Shafranov contour integral by the expression:

$$\mu_o\, zI_p = \sum_{i=1}^{38} z_i\, B_{pol_i}\, \delta l_i + \sum_{i=2}^{38} \Delta\psi_i\, \beta_i \sim \int\{z\, B_{//} + R\log(R/R_o) B_{\perp}\}dl \qquad (6.1)$$

in which z_i and δl_i are the vertical position and the spacing between the 38 poloidal magnetic field measurements, B_{pol_i}, and $\Delta\psi_i$ are 37 flux-loop differences with respect to the poloidal flux at the vessel mid-plane on the inner side. β_i is a set of coefficients chosen to produce minimum contributions of the individual poloidal field coil currents to the zI_p observer in the absence of any plasma current. zI_p is compared in real time with a pre-programmed reference waveform $z_{ref}(t)$ multiplied in real time by the measured value of $I_p(t)$. This form of reference programming is preferred to programming a pure reference value of zI_p directly, since the latter produces an apparent vertical position error in the presence of a plasma current error, causing problems in the control of the plasma vertical position.

The plasma current is calculated from the poloidal magnetic field probes, using a trapezoidal approximation to the contour integral, as described in Section III.I.

Having created the observers for shape, current and position parameters, we must correct the poloidal coil currents for the errors in all of these control parameters and subsequently determine the timescales on which these corrections must be applied. This is done differently for the shape parameters and the plasma current or vertical position. The observers for the plasma current or the vertical position are designed to have no response in the absence of plasma current. That is to say, we need the plasma response explicitly in order to correct for an error. For the shape parameters, the distribution of the

correction currents among the different shaping coils, δI_{pol}, is estimated as the inverse of the observer sensitivity matrix. Our choice in PCG is to derive the observer sensitivity matrix neglecting both the plasma current and the vessel current in which case the vector of correction currents becomes:

$$\delta \mathbf{I}_{pol} = \left(\mathbb{A} \cdot \frac{\partial [\mathbf{B}_{pol}, \Delta\psi]}{\partial \mathbf{I}_{pol}} \right)^{-1} \cdot \varepsilon_{CP} \qquad (6.2)$$

where $\mathbb{A}$ is the matrix creating the heuristic shape parameters from the input signals, [$\mathbf{B}_{pol}$ and $\Delta\psi$]. (Bold type indicates a vector of signals throughout and open-type indicates a matrix). ε_{CP} is the set of errors in the heuristic control parameters. The individual measurement sensitivity vectors with respect to the poloidal currents are obtained from the constructional model of TCV rather than from measurements, although the agreement with measurements is adequate, as already discussed.

This simple method produces a vector of poloidal current corrections which is independent of the plasma equilibrium and which only depends on the definition of the particular linear shape parameters to be controlled. Although applying this vector of corrections would not lead to precisely the shape parameter change required, the angle between the plasma-less calculation of the correction vector and a calculation including the full plasma response is small. Since this correction is only relevant inside the control loop, its imprecision does not create any steady-state imprecision, given the integral control used. Although this procedure leads to strict decoupling between the separate shape control parameters only in the absence of plasma, the decoupling actually obtained in the presence of a plasma is more than adequate. This method of plasma-less response controller has recently also been implemented on C-MOD [15] and is routinely used in TSC simulations for ITER [16].

The feedback controller for the dynamic correction of the poloidal shaping currents is of the form :

$$\delta \dot{\mathbf{I}}_{pol} = \left[\mathbb{P} \cdot \mathrm{diag}\left(\frac{1}{\tau_P}\right) + \mathbb{I} \cdot \mathrm{diag}\left(\frac{1}{s\,\tau_I \tau_P}\right) + \mathbb{D} \cdot \mathrm{diag}\left(s\,\frac{\tau_D}{\tau_P}\right) \right] \cdot \varepsilon_{CP} \qquad (6.3)$$

For all the shape parameters excluding the vertical position, the matrices $\mathbb{P}$, $\mathbb{I}$, and $\mathbb{D}$ are the same, derived as in Equation (6.2) and their feedback time-constants are the same. Diag(x) refers to a matrix with the diagonal having values of the vector x. Each control parameter therefore has its own dynamics defined. Those control parameters which are coupled, such as the shape control, have to have the same values of these time constants. Each τ_P defines the timescale of the proportional correction of the feedback error, τ_D defines the relative importance of the derivative feedback and τ_I defines the timescale on which integral errors are cancelled. For the shape parameters, the values of τ_P, τ_I and τ_D are typically 2, 40 and 1 msec.

For the control of the vertical position, such a simple method is

inapplicable since the observer defined in Equation (6.1) is specifically constructed to be independent of the control coil currents in the absence of plasma, contrary to the case of the shape parameters which allowed the use of Equation (6.2). We calculate the necessary radial field correction assuming a typical vertical field curvature and apply this correction using many poloidal field coils, distributing the correction current requests between them intuitively on the basis of previous modelling [17]. Proportional and derivative feedback are applied to the coil-pairs E2-E7, E3-E6, E4-E5, F2-F7,F3-F6 and integral feedback is applied to the coil-pairs F1-F8, F2-F7, E1-E8 and E2-E7. We use weak integral feedback, $\tau_I \sim 100$ msec, to ensure that a zero position offset is obtained with finite proportional gain, $\tau_P \sim 20$ msec. Derivative feedback, $\tau_D \sim 1$ msec, produces vertical stabilisation and the damping of vertical position control oscillations. Integral feedback also obviates the requirement for compensation of the coil resistances.

For the control of the plasma current, we assume a nominal value for the plasma inductance and consider the mutual inductance between the plasma column and the poloidal shaping and ohmic coils to be approximated by their mutual inductance with the poloidal flux loop at the inboard mid-plane.

VI.C Results

The PCG algorithm was used for the creation of several of the discharges illustrated in Fig. 3 (#5913,4211,6479). Discharges were obtained with elongations up to 1.76, plasma current up to 510 kA, both signs of triangularity, Double-Null diverted discharges in L- and H-Mode. It is important to note that all of the these conditions were obtained with the same controller algorithm, the same observers and the same pre-programming method. It is clear that restricting the range of the equilibria to be controlled would make the problem easier. Application to TCV is therefore a stringent test for any new algorithms.

Figure 22 shows the superimposed reference and obtained waveforms for the 5 poloidal configuration variables for a simple discharge, #6410. The plasma current recovers the pre-programmed waveform at 38 msec into the discharge, after which it remains within 2 kA of the reference until the end of the discharge, when a small error of 4 kA precedes a loss of current at 1.03 seconds at the end of the current ramp-down. The vertical position is maintained to within 3 mm during the plateau and within 1 cm with the exception of the initial breakdown. The PCG controller keeps the three shape control variables, plotted normalised to I_p, close to their reference waveforms except at the very beginning and end of the discharge, when some of the poloidal power supplies are inhibited at zero current. The lower curves show the last closed flux surface elongation and triangularity for these discharges.

In the second discharge, #4239, a feedback controller with no integral gain was unable to correct for pre-programming uncertainties and both I_p and the inboard collinearity had errors at the highest values of I_p. The outboard collinearity was unable to follow a pre-programmed excursion early in the

discharge, again due to power supply blocking at zero current. These two discharges illustrate the ability of the PCG controller to provide correct feedback even in a primitive version with no integral feedback (# 4239) and how the reference waveforms are correctly fedback using the full integral feedback control (# 6410) defined in the previous section.

Figure 23 shows the comparison between the physics parameters drawn by the operator (lines), namely I_p, elongation, triangularity and mean radius, mapped by the Neural Network to the heuristic variables, and the actual values for these parameters reconstructed after the discharge (crosses). The elongation was maintained close to the desired value even during an inversion of the triangularity during the discharge. The mean radius reference was significantly moved at the end of the discharge and successfully tracked.

The Neural Network pre-programming gave poor results when extrapolating outside the database, as to be expected, normally failing to extend the variable very far beyond the populated region. The PCG database is now being updated on the basis of new discharges. The question of the optimal choice of the control variables to be used is also being addressed. To test the resiliance of the method to the choice of control parameters, a new set of parameters, simply the values of the poloidal flux differences at several points around the vessel, was implemented, new reference waveforms were generated, a new controller was automatically calculated as described in Section VI.B and a virtually identical discharge was obtained.

In order to test the decoupling properties of the PCG controller algorithm, square pulse perturbations were injected into each of the control parameter signals, to measure the interaction of such a perturbation with the other control parameters. Such a perturbation should lead to an initial step in the control parameter perturbed but this response should be nulled on the timescale of the integral feedback, even during the application of the perturbation. This behaviour is seen in Fig. 24. The decoupling of the z-position and the plasma current are excellent, as expected. The decoupling between the triangularity parameters is less good, also as expected, but the residual coupling is acceptable for routine operation.

Using pre-programmed currents leads to several advantages. The controller gains can be reduced. The coil currents are controlled following a disruption in which case the feedback is cut off automatically and the currents revert to their pro-programmed values. A pre-shot check on their values allows us to avoid current saturation.

VII. MGAMS CONTROL ALGORITHM

VII.A The Method

TCV is a rather unusual tokamak since it allows the creation of an extreme variety of plasma shapes and magnetic configurations. The MGAMS shape and position control algorithm [18] has been developed specifically for TCV with the aims of versatility, accuracy and ease of operation. The basic idea is that the operator only specifies the plasma shape and plasma current evolution and everything else is done by the MGAMS software. This concept has proven very efficient in TCV since it allows us to create completely new plasma shapes in a single discharge, without using any information from previous discharges.

As mentioned in the introduction, plasma shape control in tokamaks involves basically four steps: (a) Identification of the plasma shape in real time, (b) Comparison of the real plasma shape with a pre-programmed shape, (c) Evaluation of coil current corrections such as to bring the the real shape as close as possible to the pre-programmed shape, (d) Evaluation of coil voltages which are necessary to produce the desired coil current corrections. In the following paragraphs, we will show how each of these four steps is implemented in the MGAMS algorithm.

VII.A.a Shape Identification

In MGAMS, the plasma current distribution is reconstructed, in real time, in the form of a finite element matrix [19]. The current elements are fixed in space, and their amplitudes are determined such as to produce the best fit to the magnetic measurements (magnetic field probes, flux loops, coil currents and vessel currents). It should be noted that in MGAMS, it is essential to include the effects of vessel currents in the shape identification procedure, since the vacuum vessel has very low toroidal resistance (55 $\mu\Omega$). Once the plasma current distribution is known, shape parameters can be derived in various ways. The standard method in MGAMS consists of defining a number of moments of the plasma current distribution,

$\Sigma I_i (Z_i-Z_o)$	vertical position
$\Sigma I_i (R_i-R_o)$	radial position
$\Sigma I_i (Z_i-Z_o)^2$	elongation
$\Sigma I_i (Z_i-Z_o)^2 (R_o-R_i)$	triangularity

where I_i is the total current in the i-th element, (R_i, Z_i) are the coordinates of the current centroid of the i-th element and (R_o, Z_o) are the coordinates of an arbitrary reference position, usually taken as the position of the magnetic axis of the final equilibrium of the discharge scenario under study.

Alternatively, the plasma shape can be defined in terms of flux errors at the plasma boundary [18] or, equivalently, in terms of gaps between plasma

surface and vessel wall.

VII.A.b Reference shape

The time evolution of the reference plasma shape must be specified by the operator, either in analytic form,

$$R(t) = R_0 + a\cos(\theta + \delta\sin\theta + \lambda\sin2\theta)$$
$$Z(t) = Z_0 + a\,\kappa\sin\theta$$

where θ is the poloidal angle and the parameters R_0, Z_0, a, κ, δ, λ are functions of time, or in the form of discrete boundary points, $R_i(t)$, $Z_i(t)$. The number of boundary points specified at each time slice is typically between 4 and 20. In addition, coordinates of X-points and separatrix strike points can be prescribed as functions of time. Using this information, MGAMS starts the discharge preparation by computing a number of free-boundary equilibria [20]. This gives us a first estimate of the coil currents as functions of time. From these equilibria we also compute the various moments of the plasma current distribution, as defined in Section VII.A.a. They serve as reference waveforms for those shape parameters which are feedback controlled. The final step in the discharge preparation is the calculation of feedforward coil voltages. This is achieved by solving the circuit equations for the complete tokamak discharge, including plasma and vessel currents.

VII.A.c Coil current corrections

Coil current corrections are computed in real time by the TCV Hybrid Matrix Multiplier, Section II. In MGAMS, these corrections are usually applied in the form of moments (vertical field, radial field, quadrupole field, hexapole field, etc.). Each moment is defined by a particular coil current distribution and, ideally, controls only one shape parameter. The moments are constructed as an orthogonal set. The amplitude of each moment, in real time, depends on the deviation of the measured shape parameter from its pre-programmed value. If the plasma shape is expressed in terms of flux errors at the pre-programmed boundary, the coil current corrections are evaluated in such a way as to minimize the flux errors [18].

VII.A.d Coil voltages

The computation of instantaneous coil voltages which are required to produce the desired coil current corrections is basically straightforward, but complicated by the perturbing effect of the vessel currents.

VII.A.e Vessel currents

Compensation of vessel currents is important mainly during the breakdown and current rampup phases. In MGAMS, this is achieved by performing a dynamic simulation of the entire discharge, based on a sequence

of pre-calculated plasma equilibria. In this simulation, the coil currents are computed such as to obtain the required equilibrium fields inside the vessel, taking into account induced vessel currents.

VII.B User Interface

A graphical user interface has been developed for MGAMS, based on MATLAB software. The operator first specifies all input parameters. The plasma current must be given as a function of time (rampup phase, flattop and rampdown phase). The shape evolution must be specified, as explained in Section VII.A.b. The rampdown phase is usually assumed to be the exact inverse of the rampup phase. However, in many cases, it is advantageous to program the rampdown phase slightly differently in order to avoid disruptions during rampdown. Finally the proportional, differential and integral gains have to be specified for the various feedback loops (plasma current, vertical position, radial position, plasma density, shape parameters, etc.).

In MGAMS, all integral gains are assumed zero, and differential gain is only used for vertical position control. Proportional gains are chosen such as to avoid excessive coil voltages. The optimum ratio between differential and proportional gain for vertical position control depends on the open loop growth rate of the unstable plasma. The choice of these gains is not critical except in cases where the inverse growth rate becomes comparable with the response time of the power amplifiers.

Once the operator has set the various parameters described above, the MGAMS discharge preparation is launched. Coil currents and voltages are computed as functions of time, and the control matrices are generated.

VII.C Results

The MGAMS algorithm has allowed the creation of many different plasma shapes in TCV. An example of the equilibrium reconstructions for a high current plasma evolution is shown in Fig. 25. The plasma current is ramped up to 810 kA in 0.43 sec and the maximum elongation is 2.0. The accuracy of the algorithm can be judged from Figs. 26 and 27. where we show pre-programmed and measured shape parameters for two typical discharges. Fig. 26 shows elongation and triangularity as functions of time for a limiter discharge and Fig. 27. shows the X-point coordinates of a single-null divertor discharge. Note that the deviation of the measured X-point position with respect to the pre-programmed position is typically less than 1 cm. The accuracy of shape control has also been tested by making use of other diagnostics, such as soft X-ray and visible imaging cameras. It is found that, for vertically elongated plasmas, the radial accuracy is generally better than 5 mm and the vertical accuracy is of the order of 1 cm. Fig. 28. shows a number of plasma shapes that were produced in TCV, using MGAMS. The sequence is ordered chronologically, starting in September 1993 and ending in October 1994. Each image is taken from a

different discharge, and in each case, the shape is shown at the time when the plasma current has reached its maximum value. Some of the shapes shown here were generated for testing the algorithm, but in most cases, the particular shape was a necessary condition for reaching a certain physics goal.

VII.D Conclusions

MGAMS has proven to be an extremely versatile tool. It can generate any arbitrary plasma shape, subject to obvious physical constraints, without reference to a data base of precalculated equilibria or previous discharges. The shape accuracy depends on the precision of the magnetic measurements and is typically between 5 mm and 1 cm. Thanks to a graphical user interface, the algorithm is very easy to use. The original goal, that the operator only specifies the evolution of plasma current and shape and the algorithm does the rest, has been fully attained.

VIII. VERTICAL CONTROL EXPERIMENTS

VIII.A Introduction

Vertically elongated tokamak plasmas achieve higher performance due to an increased plasma current for given major and minor plasma radii and toroidal magnetic field. However they require active stabilisation of the inherently unstable movement of the vertical position as well as a more complex poloidal field coil system. All elongated tokamaks have achieved this stabilisation within certain operational limits and there is an extensive literature on the modelling of vertical position control. In large devices, the design of the feedback control of the vertical position must take into particular account three aspects. Firstly, the power supplies required for poloidal field control have a high bandwidth and large inductive power, so reducing their performance requirements by appropriately designing the vertical position control would be economically important. Secondly, in the event of a loss of vertical position control, the currents induced in the vessel by the feedback control and by the plasma motion give rise to huge vertical forces in the presence of the externally applied radial fields, so reducing the magnitude of the induced vessel currents and the applied radial fields is mechanically important. Thirdly, the AC currents induced in the vessel or the surrounding structures by the flux variations associated with the vertical feedback control are significant and can give rise to unwanted or unnecessary ohmic dissipation in those elements, so reducing the eddy currents for a given feedback capability could allow greater vertical instability growth rates to be handled for a given dissipation.

The control of the vertical position is a sensitive task in plasma shape control, since it alone is unstable. It can usefully be separated from the general problems of shape control for two reasons. Firstly, the frequency bandwidth and the required feedback voltages are imposed by the growth rate of the

vertical position instability and by the variation in the plasma vertical position which must be controlled, rather than by a choice of operational time-scales. Shape and current control can always be slowed down to stay within the limits of the poloidal field supplies, whereas the vertical position control cannot. In present devices, shape control is usually performed on a slower time-scale than the time-constants of the vessel current distribution. Secondly, the vertical position control is only weakly dependent on the plasma equilibrium model but depends on the spatial structure of the radial field produced by up-down anti-symmetric shaping current distributions.

The vertical movement of the plasma can be studied experimentally in several ways. The simplest is to open the control loop and measure the ensuing unstable vertical motion. This has the disadvantage of leading to a disruption as well as non-linearity problems as the excursion increases, but yields a result which is simple to interpret. This type of experiment on TCV is described in Section VIII.C. Stimulating the closed feedback loop by injecting perturbation signals into the loop is an alternative method also employed on TCV to determine the relative merits of the different poloidal field coils for providing vertical stabilisation, described in Section VIII.B. Both experimental techniques are compared with predictions of a Rigid Current Displacement Model and the NOVA-W deformable plasma model.

VIII.B Closed-Loop experiments

Stimulating the closed feedback loop by injecting perturbation signals into the loop is one method of studying the vertical position control. In this method perturbations are directly summed into the voltage demand signals for the poloidal field coil supplies, shown as the inputs Q in a simplified block diagram, Fig. 29. In this case we measure the response of the loop to a perturbation, not its response to a change in the reference signal. Comparing the responses of different coils in this manner in fact gives the same result as comparing the open loop responses. Injecting perturbations into the closed feedback loop indicates which coils are able to provoke a significant vertical movement for different frequencies. Those coils which provoke the smallest movement at high frequency are likely to be least useful for controlling the unstable vertical position.

We have carried out such perturbation excitation measurements for all eight up-down anti-symmetric pairs, such as E8-E1, F5-F4, Fig. 2. Two types of perturbation have been applied, square-wave injection and continuous random injection; the experiments are described separately. Since both methods excite the same points in the feedback loop, the experiments are essentially the same, in that the same information could be extracted from both. In practice, the spectral content of the continuous excitation is fairly flat out to 250 Hz and ought to provide more information for a given excursion of the controlled parameter. These experiments and their analysis are described in more detail in [21].

VIII.B.a Square-wave injection

In the first such experiment, we successively excited each up-down anti-symmetric coil-pair within a single tokamak discharge, applying single 25 msec. pulses of opposite signs to the upper and lower coils simultaneously. Figure 30 shows the time-history of the voltage perturbations (Q), the demand signals and perturbations summed (U, Fig. 29), the coil currents and the plasma position response (zI_p, Fig. 29). The asymptotic response of the feedback loop to these perturbations is to annul them in the presence of integral feedback gain, since only a step in a reference input should lead to a step in a controlled parameter. However, the time-scale of the perturbations is such that the feedback loop does not have time to annul them completely, although the tendency is visible in the figure. This is not a problem for us, since we are essentially looking for responses on the timescale of the vessel current decay and faster.

Up-down anti-symmetric excitation provides a zero net poloidal flux perturbation and does not couple to the total plasma current. Its measured effect on all shape parameters other than the vertical position was negligible for the reasons invoked in the introduction, confirming the simplicity of this experiment.

The differences between the responses of the different coil-pairs are striking, considering that the coil-pairs were excited by the same perturbation with similar volt-seconds injected. The responses for the coil-pairs F8-F1 and E8-E1 are particularly damped, indicating that these coils will not be able to provide fast recuperation of jumps in the vertical position, nor will they be able to stabilise the most unstable plasmas. The coil-pairs E6-E3 and E5-E4 show significant overshoot in the response, indicating a faster penetration of the radial control flux. The coils which provide both the fastest and largest experimental responses are those already used in the vertical position feedback loop itself.

VIII.B.b Random Excitation

In the second experiment, a random binary sequence continuous excitation was injected into the same points in the feedback loop as in the case of the square-wave injection (inputs Q). In this method, a random time-sequence of either +U or -U volts was applied to the two coils in each coil-pair, in opposition, for a duration of 700 msec. The time-step of the random sequence was 1 msec, interpolation between points being linear. This procedure was repeated one coil-pair at a time over a series of eight identical stationary discharges with the same random excitation sequence used for all coil-pairs. Coupling to the current and all shape parameters other than the vertical position was measured to be negligible, as for the square-wave excitation.

Figure 31 shows the time-response of the current-weighted plasma vertical position which is linearly related to the applied flux perturbation. Its standard deviation, $\sigma(zI_p)$, is also indicated for each coil-pair. Before and after the perturbation injection there is a residual $\sigma(zI_p)$ of about 60 A.m. A similar

experiment was carried out with simultaneous injection into all coil-pairs, but this time injecting non-correlated random sequences into each coil-pair. $\sigma(zI_p)$ was measured to be 1897 A.m, close to the $\sigma(zI_p)$ calculated from the separate excitations, 1832 A.m, illustrating that within 700 msec we can achieve effective decorrelation of the different noise waveforms injected into the individual coils. Increasing and decreasing the amplitude of these simultaneous injections by a factor of two verified the linearity of the response. The different amplitudes of the responses are clear, but the nature of the response, overshoot or damped, is not visible in the raw data as it was for the square-wave injection.

$\sigma(zI_p)$ varied widely from coil-pair to coil-pair over a range 165 - 1213 A.m, tabulated in Column 1 of Table 8.1 and shown in Fig. 32(a) with the coil-pairs plotted ordered around the vessel. The response is even slightly overstated due to the 60 A.m residual variation without excitation. The smallest response was obtained for the coil-pair F8-F1. The outer stack coil-pair F7-F2 and the extreme inner stack coil-pair E8-E1 are also fairly weakly coupled to the plasma vertical position. The coil-pairs which provided the largest $\sigma(zI_p)$ response, E6-E3, E5-E4 and F6-F3 also provided the largest and fastest pulse-response in the first experiment.

Column 5 of Table 8.1 and Fig. 32(b) summarises the power for each individual coil-pair, needed to excite a unit vertical position movement (1 A.m) with the voltage perturbation flat as a function of the excitation frequency. This quantity is evaluated combining the standard deviation of the coil-pair currents and voltages: $\sigma(U)\sigma(I_{pol}) / \sigma^2(zI_p)$. The variation between the reactive power used for the in-board and out-board coils is less than the variation in their $\sigma(zI_p)$ responses for fixed applied voltages due to their different inductances, but the hierarchy already established is maintained.

We have analysed the poloidal distribution of the vessel currents from the poloidal distribution of the loop voltage measured using the 38 voltage loops placed on the outer surface of the vessel, combined with a constructional estimate of the poloidal distribution of the vacuum vessel resistivity. Figure 33 shows the measured poloidal distribution of these currents, illustrating their closeness to the driven coils as to be expected. Table 8.1 column 6 shows the values of the ohmic power dissipated in the whole vessel, Fig. 32(c), normalised for $\sigma(zI_p) = 1$ A.m. The ohmic power dissipated in the vacuum vessel for a given vertical plasma displacement varies by a factor of 10 between the F1-F8 combination, the worst, and the E6-E3 combination, the best. Changing the coils used for the vertical feedback control therefore has a factor of 10 effect on the power dissipated in the vacuum vessel, or has a factor of $\sqrt{10}$ on the amplitude of the excursion which can be controlled given limited feedback power.

VIII.B.c Resume of the Experimental Information

These two experiments carried out on TCV have shown that there is a significant difference between the reaction of the plasma vertical position to voltage excitation of the different up-down anti-symmetric poloidal field coil-

pairs. Some coil-pairs provide little high frequency response, others provoke a higher amplitude response at higher frequencies. In order to understand the underlying mechanism determining this difference between coil-pairs, we inspect the predictions of a model of the vertical plasma movement based on simple considerations of the movement of a rigid but distributed plasma current.

VIII.B.d Comparison with the Rigid Current Displacement Model

A simple rigid current displacement model describing the interaction between a centred filamentary plasma with a single active anti-symmetric poloidal coil in the presence of a single vacuum vessel eigenmode was applied earlier to the DIII-D tokamak [17]. This model remains attractive by its simplicity and its similarity to experimental observations, as well as to its ability to predict good vertical position feedback controllers [22]. However we have had to extend this model in several directions for application to TCV.

Firstly, the design of the TCV vacuum vessel together with the fact that the plasma current may not be centred vertically leads to time constants of the vacuum vessel current distribution eigenmodes which are poorly separated. As a result we take the slowest 20 eigenmodes of the vessel into consideration, including up-down symmetric and up-down anti-symmetric modes. All conducting structures in which toroidal current can be induced during vertical position control can be included into this eigenmode decomposition, although TCV does not need this complication.

Secondly, we wish to treat the system with all the coils active together rather than modelling separate coil-pairs on their own. This is done by explicitly including all poloidal field shaping coils in the circuit equations as eight anti-symmetric pairs.

Thirdly, in the face of the highly structured vacuum equilibrium field possible for TCV plasmas, we perform a volume integration of the forces on the plasma due to the radial fields from the vacuum vessel currents and the coil-pair currents. This does not modify the structure of the control problem and we find that it does not modify the estimate of the growth-rate, provided the radial position of the current filament is taken as the major-radius-weighted value of the current distribution, $\langle R_{fil} \rangle = \int j(r)\, R(r) d^2r \,/\, Ip$.

The feedback of the vertical position error $\varepsilon(zI_p)$ is added to the model as a vector of voltage corrections, Fig. 4: $U_{fb} = \mathbb{M}_{aa} \bullet (1/s\, G_i + G_p + sG_d) \bullet \varepsilon(zI_p) = (1/s\, C_i + C_p + sC_d) \bullet \varepsilon(zI_p)$, where G_i, G_p and G_d are the vectors of gain values actually used in the discharges to be modelled and $\mathbb{M}_{aa}$ is the mutual inductance matrix for the poloidal field coils. In order to model the integral feedback term, we add an extra variable, representing the time-integral of the error signal.

Finally we must make a first approximation for the power supply

response. The necessity for this is obvious from the almost square response of the coil-pair currents to the voltage input signals, Fig. 30, indicating an apparent resistive load for frequencies at which the coils themselves are mostly inductive. This correction is due to the control mode of the power supplies as they were used in these particular experiments. The external voltage demand input signal is internally in opposition to a proportional controller maintaining the coil current close to a pre-programmed current demand input signal. The weak gain in this internal power-supply controller corresponds exactly to an modifying the effective active coil resistance, $R_a \rightarrow R_a + K$ where K = 2x0.51 Ω and 2x3.24 Ω for the E- and F-coil-pairs respectively.

We presuppose that the only voltage demand inputs to the power supply pairs arise from the zI_p feedback controller C(s). In practice, the other feedback loops controlling the plasma shape are also active and would produce a control loop response if their errors were sensitive to the zI_p error. To first order there is no coupling for the up-down symmetric shape control used in these discharges, as already discussed.

With all of these modifications, the set of differential equations describing the plasma current vertical position, the vessel eigenmode currents and the coil-pair currents and voltages becomes :

$$\begin{bmatrix} s M_{aa} + R_a & s M_{ae} & s\frac{\partial M_{ap}(z)}{\partial z} + s C_d + C_p & C_i \\ s M_{ae}^T & s M_{ee} + R_e & s\frac{\partial M_{ep}(z)}{\partial z} & 0 \\ -s\frac{\partial M_{ap}(z)}{\partial z} & -s\frac{\partial M_{ep}(z)}{\partial z} & -s\sum\frac{\left(\frac{\partial M_{ep}(z)}{\partial z}\right)^2}{L_e}\frac{\alpha(z,I_{pol})}{\alpha_{crit}} & 0 \\ 0 & 0 & -1 & s \end{bmatrix} \begin{bmatrix} I_a \\ I_e \\ z\,I_p \\ \int z\,I_p\,dt \end{bmatrix} = \begin{bmatrix} Q \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

where I_e, M_{ee}, dM_{ep}/dz and Re are the vessel eigenmode currents, their mutual inductances, their mutual inductance with the vertical position and their resistances; I_a, M_{aa}, dM_{ap}/dz and R_a are the same for the active shaping coil-pairs; Le is the self-inductance of the vessel eigenmodes; Mae is the mutual inductance matrix between the vessel eigenmodes and the coil-pairs; the term α/α_{crit} defines the stability margin F_{stab}/F_{destab}, where α_{crit} defines the maximum value of field curvature for which the vertical position is stable in a perfectly conducting vessel. It corresponds to the limit as the determinant of the above equations approaches zero. The mutual inductances between the currents and the plasma vertical position was calculated by integrating over the plasma current distribution, avoiding the simplification of the rigid filament model. All

electromagnetic properties of TCV are derived from the geometry of TCV rather than from physical measurements.

The destabilising force due to the imposed field curvature, $\partial B_R/\partial z$, is given by $F_{destab} = \delta z \int j(r) \frac{\partial B_R(r)}{\partial z} d^3r$ where $B_R(r)$ is the radial equilibrium field and δz is a vertical displacement. The image currents induced in the vessel are given by the amplitudes of their eigenmodes $I_{ek} = \delta z \int j(r) \frac{\partial B_{Rek}(r)}{\partial z} d^3 r$ and give rise to a stabilising force $F_{stab} = \int j(r) \sum_k \frac{\partial B_{Rek}(r)}{\partial z} I_{ek} d^3 r$ where $B_{Rek}(r)$ is the radial field produced by a unit current in the kth vessel eigenmode. We prefer the form α/α_{crit} in the circuit equations to the previously used form involving the poloidal field curvature index since the former does not depend on the equilibrium vertical field which has no role in the vertical stabilisation.

The circuit equations were first used to model the response of the plasma vertical position to arbitrary and multiple input waveforms by time-integration. Eight input pulses were injected during a single simulation, exactly as in the experiment, Fig. 30, and the result is shown as a dotted-dashed line in Fig. 34, overlaying the experimental signal, shown as a solid line. The rather damped experimental responses of coil-pairs E1-E8 and F1-F8 are well reproduced. The fast experimental responses with significant overshoot, coil-pairs E6-E3 and E5-E4, are also well modelled. This excellent agreement implies that the underlying mechanisms determining the zI_p evolution are contained in the circuit equations, specifically, the large differences between the effectiveness of the various coil-pairs already seen in Fig. 31 are faithfully reproduced.

Secondly we evaluated the modelled time-response to the random injection in the same manner and calculated its standard deviation. However, the random injection has a fairly flat voltage frequency spectrum and for each coil-pair we therefore expect $\sigma(zI_p)$ to be given by: $\sigma^2(zI_p) \sim \sum_k H^2_{QiZ}(\omega_k)$, where H_{QiZ} is the transfer function between the coil i and the vertical position z. These two methods of evaluating the modelled response agreed. Figure 35 shows the values of the modelled response for the eight coil-pairs, adding a 60 A.m. residual amplitude, plotted against the experimental responses $\sigma(zI_p)$ taken from Fig. 31. The agreement is excellent.

The very simple experiments described in this section have demonstrated a large variation in the quality of the plasma vertical position response to different forms of voltage perturbation applied to up-down anti-symmetric coil-pairs in TCV. These different responses were compared with the vertical position feedback loop closed, avoiding the necessity of evaluating the less accessible open-loop responses. The results clearly show the extreme importance of the poloidal location of the vertical feedback control coils. We can reduce the reactive power, reduce the power dissipation leading to vessel

heating or reduce the image currents in the vessel by choosing the coil positions used for the control of the vertical position according to these experimental results.

The coil-pairs F1-F8 and E1-E8 provide a large radial field on axis for a given current, Column 4 of Table 8.1. However they also couple strongly to the slowest up-down antisymmetic eigenmode of the vacuum vessel, Column 5 of Table 8.1. The slow vessel eigenmodes, whose amplitudes peak at the top and bottom of the TCV vessel, effectively shield the radial magnetic field generated by these coil-pairs. The coil-pairs which are best coupled to the fast vertical motion are those in the inner stack close to the equatorial plane, providing relatively weak radial magnetic field on axis, namely E6-E3 and E5-E4. However, since they are also weakly coupled to the slow vacuum vessel eigenmodes, these coil-pairs allow faster action on the plasma position. The distance of the coil from the plasma axis, Table 8.1 Column 6, does not uniquely determine the amplitude of the response either. Work is currently underway to exploit the amplitude and phase information contained in the results from these experiments in order to improve the vertical position feedback controller.

VIII.C Open-Loop experiments

The open-loop growth rate of the vertical instability has been experimentally measured in a set of TCV discharges. The growth rate was measured by opening the vertical (and radial) feedback control loops with fixed coil currents and observing the ensuing exponential growth of the vertical displacement via the magnetic diagnostics. We have made two different estimates of the plasma vertical position, both based on magnetic signals. Firstly, we employ the realtime zI_p observer used for the feedback control of the plasma vertical position, which is a linear combination of the magnetic probe and flux loop measurements. The vertical position determined from the full equilibrium reconstruction code LIUQE, is used to obtain a second estimate of the vertical growth rate.

The feedback cut was performed in hardware by a pre-programmed switch to a new set of PCS control matrices (Figure 4), in which the elements corresponding to the vertical and radial PID feedback gains were set to zero. The matrix switch was timed to occur during the current plateau, after the target configuration had been fully established. Figure 36 shows an example of the unstable vertical motion which follows a vertical feedback cut during the current flat-top (at t = 0.47 s).

We limited the subsequent growth rate analysis to small initial displacements, less than 4 cm, for which the plasma current and cross-section do not change significantly. During this time, before t = 0.477 s in Fig. 36, the plasma motion is purely vertical and essentially rigid. This is also evident from the time evolution of the global plasma parameters which are shown in Fig. 37.

Low current (250 kA) limiter discharges were chosen in order to minimize the harmful effects of disruptions which invariably followed the feedback cut. Single-parameter scans were performed in order to investigate the dependence of the vertical growth rate on elongation, triangularity and vertical position of the magnetic axis. The series of limiter configurations which were produced for this study spans a range of edge elongations from 1.4 to 1.75, plasma triangularity from 0.1 to 0.65, and vertical positions from the midplane ($z = 0$) to the top of the vacuum vessel ($z = +0.29$ m). The plasma configurations which correspond to the extremes of each global parameter range are shown in Fig. 38.

Figure 39 shows the experimentally measured growth rates as a function of the global plasma parameters: elongation, κ, triangularity, δ and vertical position of the magnetic axis, z. The growth rates are calculated from the realtime zI_p observer and the equilibrium reconstruction code, LIUQE. The calculated growth rates are seen to show very good agreement, within the estimated relative uncertainties in each measurement.

The scan of vertical growth rate versus plasma elongation (Fig. 39(a)) shows that the growth rate increases very strongly with elongation. The growth rate of approximately 1000 s^{-1} measured at an elongation of 1.75 represents one of the most vertically unstable discharges produced in TCV to date. The dependence of the growth rate on triangularity (Fig. 39(b)) indicates that positive triangularity has a substantial stabilizing effect. Discharges with negative triangularity cannot be regularly produced at present in TCV, since the outer vacuum vessel wall is not equipped with carbon tiles to support these type of discharges. The series of similar discharges in which a plasma of elliptical cross-section is progressively shifted vertically, Figure 39(c) demonstrates the stabilizing influence of the thick conducting vacuum vessel top plates. The open-loop growth rate is reduced when the plasma is very close to the top wall.

The experimental growth rates shown in Fig. 39 have been compared with those obtained from the RCDM model described in the preceding section. The growth-rates predicted by the RCDM model are 10% in excess of the measured growth-rate and work is currently underway to understand this difference. We also compare the growth rate calculated with a fully deformable plasma code, NOWA-W [23]. This calculation slightly underestimated the growth rate. However the dependence of the growth-rate on the vertical position, elongation and triangularity are well reproduced by both models, indicating that the variation of the vertical instability can be well predicted on the basis of the pre-programmed poloidal coil currents.

IX. FIELD STRUCTURE AT BREAKDOWN

The TCV vacuum vessel is continuously welded and has a low toroidal resistance (45 $\mu\Omega$) for passive stabilisation the plasma vertical motion. Consequently large vessel currents, up to 200 kA, are induced by the loop

voltage during breakdown. They create a vertical magnetic field of the order of 10 to 20 mT that must be carefully compensated to achieve discharge breakdown. This compensation as well as the creation of a poloidal field quadrupole rely on accurate magnetic measurements and a precise model of the electromagnetic properties of the vacuum vessel. The available magnetic measurements consist of: (i) 38 poloidal flux loops wound on the outside of the vessel: (ii) 4 arrays of 38 magnetic probes measuring the tangential poloidal field inside the vacuum vessel; (iii) one flux loop on each of the 23 poloidal field coils; (iv) 20 current sensors for the poloidal and toroidal field coils. In addition to an accurate calibration, the consistency of all these measurements has been enhanced by small corrections in their calibration and in the sensor position, leading to an error better than 0.5 mWb and 0.5 mT for the flux and field values, and about 1 mm for their position [24]. The vacuum vessel is made of stainless steel, 15 mm thick in the cylindrical parts and 20 mm thick on the top and bottom. Port holes in the vessel make its electromagnetic modeling difficult. Instead of a complicated 3-D representation, a toroidally symmetric model with an experimentally determined spatial distribution of the effective resistivity was chosen. The vessel is divided in current filaments I_v whose circuit equations are

$$0 = R_{vv}\, I_v + M_{vv}\, dI_v/dt + M_{va}\, dI_a/dt \qquad (9.1)$$

where R_{vv} is the diagonal matrix of the filament resistances, I_a are the currents in the active coils and the M's are the mutual inductance matrices. These depend only on the geometry and are assumed to be known exactly. Trapezoidal current waveforms lasting much longer than the vessel time constant (13 ms) are injected in each active coil one at a time. During the ramps, I_v and dI_a/dt are constant and from eq. (9.1) one derives the vessel currents

$$I_v = -\, Y_{vv}\, M_{va}\, dI_a/dt, \qquad (9.2)$$

where $Y_{vv} = R_{vv}^{-1}$. The fluxes and fields measured by the magnetic sensors are

$$\psi_f = M_{fa}\, I_a + M_{fv}\, I_v$$
$$B_m = B_{ma}\, I_a + B_{mv}\, I_v \qquad (9.3)$$

The vessel filament resistances can be deduced from these measurements by solving the equations for Y_{vv} in the least squares sense

$$\psi_f - M_{fa}\, I_a = M_{fv}\, I_v = -\, M_{fv}\, Y_{vv}\, M_{va}\, dI_a/dt$$
$$B_m - B_{ma}\, I_a = B_{mv}\, I_v = -\, B_{mv}\, Y_{vv}\, M_{va}\, dI_a/dt \qquad (9.4)$$

The differences in the rhs are calculated using the flux and field values during the current flat top. In fact the available data allow the determination of only a limited number of filament resistances. Thus the latter are combined either in vessel segments, appropriate base functions or vessel eigenmodes by mean of a rectangular matrix T_{vw} and by replacing Y_{vv} by $T_{vw}\, Y_{ww}$ in eq.

(9.4). Figure 40 compares the measured vessel resistivity with a rough estimation obtained from the vessel geometry and its specific resistivity. The influence of the port holes is clear and leads to an average discrepancy of 20% which must be accounted for in programming the breakdown field.

The information required to determine the magnetic field topology inside the vacuum vessel prior to plasma breakdown are the currents in the active coils and in the vessel. These can be derived from the magnetic measurements by solving the set of equations 9.5 in a least square sense and with appropriate weights

$$\begin{aligned} \psi_f &= M_{fa}\, I_a + M_{fv}\, I_v \\ B_m &= B_{ma}\, I_a + B_{mv}\, I_v \\ J_a &= I_a \end{aligned} \qquad (9.5)$$

The last equation has been introduced to account for errors between the measured (J_a) and adjusted (I_a) coil currents. However this method has proven inadequate and gave a poloidal field null location that did not coincide with that of the breakdown seen on H_α detectors. Equation (9.1) was then also used to account for the time dependent vessel currents in a consistent way. This new set of equations (9.5+9.1) is then solved over a time period prior to breakdown, using a two point difference operator for the time derivatives. The difference in the vessel current distribution estimated by these two methods is large (Fig. 41). However, the latter gives an estimate of the field inside the vessel and has been extensively used to optimise the breakdown conditions. The drawback is the large size of the system to be solved, typically 6000 x 2000. To circumvent this, a Kalman filter was set up which works as a state estimator for the state space system whose states are (I_a,I_v), whose inputs are dI_a/dt and whose outputs are the (ψ_f,B_m,J_a). This approach gives comparable results at much reduced computation time. Because it is fast and it relies only on past samples, it could be implemented in real time as part of the future TCV digital control system.

Breakdown conditions are established in TCV by the MGAMS plasma control algorithm. As part of the shot preparation, MGAMS computes the coil currents as a function of time such that the poloidal field vanishes at the desired breakdown position, at time t=0. In addition, we prescribe a positive vertical field gradient, $dBz/dr > 0$, in order to ensure vertical stability of the initial plasma. These calculations involve a complete simulation of the experiment, including induced currents in the vacuum vessel and plasma current startup. The main difficulty with this approach is that the precision of the poloidal field configuration that can be achieved in the experiment depends on the accuracy of the tokamak model used in the simulation. In TCV, this problem is particularly severe since the vacuum vessel has very low toroidal resistance. Reliable breakdown at loop voltages of the order of 10V can only be achieved when the poloidal field errors are less than about 0.5mT. This accuracy can be obtained by adding small corrections to the prescribed values of Bz, Br, dBz/dr, dBr/dr, at the desired breakdown position, at time t=0. The corrections are deduced

from off-line vacuum field analysis, as described in the previous paragraph.. Plasma breakdown has been set up at various positions in the vessel: z = 0, ±0.23 and ±0.40 m, for both directions of loop voltage and of toroidal field. In all situations the vessel currents is similar to that plotted in Fig. 41. An example of the magnetic field configuration prior to breakdown is shown on Fig. 42 (left) for z = -0.23 m. The field amplitude is less than 1 mT in a region of more than 0.2 m in diameter. The example on the right of this figure is the field configuration for a simultaneous breakdown at z = ±0.40 m used for the creation of a doublet .

Up to now breakdown has always been obtained using a glowing filament in either the top or the bottom of the vessel. We assume that the UV emissivity of the filament is responsible for assisting the breakdown. In the future we intend to use ECRH assist.

X. BULK PLASMA CONTROL

Before the installation of the 4.5 MW of additional ECRH power under construction, bulk plasma control on TCV only concerns the control of the plasma density. Two points are mentioned, namely the feedback control of the plasma density by gas-puffing and the control of the density in particular conditions using configurational control.

Feedback control of the plasma density in TCV is illustrated in Fig. 43, showing the plasma density tracking the demand signal during the current and density rise, using a PID controller with relative time-constants $\tau_I/\tau_P = 10$, $\tau_D/\tau_P = 0.14$. The τ_P gain is not known in absolute terms. At 0.39 seconds a square pulse is imposed on the demand signal, seen in the second window of the figure. During the requested density rise the derivative term provides a very fast leading edge to the gas-valve response (shown together with the preprogrammed signal) and the response of the plasma density is excellent, being fast and non-oscillatory. When the density demand signal returns to its original plateau value, the derivative term completely closes the gas-valve, shut for demand voltages below 1.8 Volts. The plasma response however is very slow, corresponding to the recycling properties of the vessel. The variation of the vessel conditioning can be usefully monitored using this fiducial discharge which is normally run at the start of each day's operation. Modelling the plasma response to the density control actuator is therefore non-trivial. The recycling behaviour is a strong function of the plasma current and the value of the plasma density itself. It depends on the actual conditions of the vessel, varying strongly before and after boronisation, and will be most unsatisfactorily modelled by constant linear behaviour. The lower window shows a 5cm swing of the vertical position, used to verify the vertical control, which has no effect on the density control in this limited D-shaped plasma.

In ELM-free H-mode discharges, the plasma density generally rises uncontrollably, often accompanied by impurity accumulation. ELMs may enable

steady-state operation at constant plasma density and impurity levels through controlled energy and particle losses. The optimum ELM frequency is a compromise between the desired plasma density, impurity control and the degradation in confinement due to the ELMs. Previously reported methods of influencing the ELM frequency include varying the separatrix to limiter separation in ASDEX, ECRH near the separatrix or changes in the shape in DIII-D , heavy gas fuelling to cool the edge combined with off-axis ICRH in JET, Li pellet injection in C-MOD and external Resonant Magnetic Perturbations on COMPASS.

A single-null divertor configuration, in which there is only one X-point on the separatrix is referred to as SND-U or SND-L, with the former applying to a configuration in which this X-point is above the magnetic axis. In a perfectly balanced DND configuration, both X-points lie on the separatrix. Since in practice this is rarely the case, the DND-U and DND-L configurations are introduced by analogy with the SND-U and SND-L. In the DND configuration, we found that the presence of ELMs is correlated with the presence of a DND-U or DND-L configuration, so switching between these configurations should provide an ELM control mechanism.

Switching between DND-U and DND-L was achieved by modulating only the zI_p reference signal. Figure 44 shows the result of the experiment with a periodic vertical shift of $\pm$ 1.25 cm. The most striking feature of this figure is the synchronisation of the ELM-free and ELMy phases with the imposed modulation. The magnetic configurations corresponding to upper and lower displacements, inset in the figure, indicate that the ELMy phase corresponds to a DND-L, and the ELM-free phase to a DND-U configuration. The flux difference between the X-points follows the vertical position reference signal quite faithfully, but shows considerable perturbation during the ELMy phases. The plasma density remains bounded throughout the 1.5 second H-mode, rising during the ELM-free phases and decreasing in the ELMy phases. It is important to note that during this experiment the density reference signal, shown in the figure, caused the gas valve to open slightly when the density was too low. The soft X-ray intensity during the ELM-free phase rose faster than could be accounted for by a change in the plasma temperature and the increase in plasma density, implying an increase in impurity concentration. The decrease in the soft X-ray emission during the ELMy phases is consistent with a loss of the accumulated impurities. The ELMs are seen to stabilise the average impurity concentration, albeit at a level above that obtained during the L-mode. More detailed results on this technique are available in [25].

This example of density control using configurational toggling was implemented in the TCV Plasma Control System using the hardware and software already described, just by modifying some of the matrix elements in the generalised PID controller. The control system had not been constructed presupposing the separation of any of the control parameters. Had we designed the control system around our prior "knowledge" that the z-position has no subtle effect on the density evolution, then software redesign would have been necessary. We consider that it will be important in any future tokamak control

system that such prior "knowledge" should not be incorporated at all.

XI. DISCUSSION

The control of highly shaped plasmas in the TCV tokamak has required a complex poloidal field coil system and a correspondingly complete plasma control system. Early operation of TCV has shown that all the systems implemented are wholly adequate for their task up to the limit of the power supply cut-off frequency. Different algorithms implemented so far have shown that well-controlled plasmas can be obtained using a variety of methods capable of automating the generation of feedback controllers.

The generalisation of the PCS has clearly required considerable effort during the design and construction of TCV and a simpler more restrictive system could have been built. However, the inclusion of all feedback control actions in the one generic system without any preconceptions of the final structure of the control method, allowed us to implement the configurational control of the density of H-Mode plasmas, discussed in Section X, with no hardware or software modifications, simply with an addition of algorithmic data to the pre-defined structures. We consider that the operational flexibility obtained outweighs the additional effort. Although some of the matrix weights could have been eliminated, the operational ease is important. All functional features of the PCS have been usefully employed. The only under-exploited feature is the number of matrix switches in a single discharge, which has never exceeded three for TCV. This overdesign was simply a choice between available memory sizes and implied no extra effort.

Considerable effort has been put into the magnetic diagnostics; without adequate quality in these, no reasonable control of the plasma shape or reconstruction of its shape is possible. We are currently improving the control and pre-programming methods for different operational requirements as the experimental program evolves, using the flexibility of the top-level PCS software over the algorithm-independent layers.

In the near future we will be implementing the Fast Internal Coil set to stabilise more elongated plasmas with plasma currents in excess of the 800 kA achieved so far.

The simple structure of a distributed database has allowed us to implement the TCV plant control in a very simple way. This open system has allowed us to construct intelligent control layers in software up to an extremely simple and maintainable cycle sequencer which avoids the complexity of a formal state machine.

ACKNOWLEDGEMENTS

It is a pleasure to acknowledge the collaboration in many aspects of this work with colleagues at the MIT Plasma Fusion Center, particularly Martin Greenwald, Tom Fredian, Steve Horne, Ian Hutchinson, Josh Stillerman and Steve Wolfe. Our colleagues on the TCV team have had to be tolerant during part of the plasma control learning curve. This work was partly supported by the Fonds National Suisse de la Recherche Scientifique.

REFERENCES

1. F. Hofmann, S.C. Jardin, F.B. Marcus, A. Perez and A.D. Turnbull, "Equilibrium and axisymmetric stability of the proposed TCV tokamak", Proc. 14th Symp. on Fusion Technology, 687 Avignon, 1986 (1987)
2. F. Hofmann, J.B. Lister, M. Anton, S. Barry, R. Behn et al., "Creation and control of variably shaped plasmas in TCV", Plasma Phys. and Contr. Fusion, 36, B277 (1994)
   J.B. Lister, F. Hofmann, M. Anton, R. Behn, S. Bernel et al., "Variable configuration plasmas in TCV", 15th Int. IAEA Conf. on Plasma Phys. and Contr. Nucl. Fusion, A-5-1-2, Seville (1994)
3. F. Hofmann and G. Tonetti, "Tokamak equilibrium reconstruction using Faraday rotation measurements", Nuclear Fusion 28, 1871 (1988)
4. J.B. Lister and H. Schnurrenberger, "Fast non-linear extraction of plasma equilibrium parameters using a neural network mapping", Nuclear Fusion, 31, 1291 (1991)
5. J.B. Lister, H. Schnurrenberger and Ph. Marmillod, "Implementation of a multi-layer perceptron for a non-linear control problem", Lausanne Report LRP 398/90 (1990)
6. P.F. Isoz, J.B. Lister and Ph. Marmillod, "A hybrid matrix multiplier for control of the TCV tokamak", Proc. 16th Symp. on Fusion Technology, 1264 Oxford, 1990 (1991)
7. R.W. Means, HNC Software Inc., private communication
8. G. Flor, G. Manduchi, T.W. Fredian, J.A. Stillerman and K.A. Klare, "MDS-PLUS. A comprehensive data acquisition and analysis system", Proc. 16th Symp on Fusion Technology,1272 London, 1990 (1991)
9. D. Fasel, G. Depreville, A. Favre, J-D. Pahud, A. Perez and F. Puchar, "19 rectifiers to supply the coils of the TCV tokamak", Proc. 16th Symp. on Fusion Technology, 1492 London, 1990 (1991)
10. A. Perez, I.M. Canay, J-J. Morf, J-D. Pahud, R. Seysen and J-J. Simond, "A 220 MVA turbo-alternator for the TCV tokamak power supplies", Proc. 15th Symp on Fusion Technology, 1449 Utrecht, 1988 (1989)
11. A. Favre, A. Elkjaer, D. Fasel, J-M. Moret, A. Perez and E. Steinmann, "Fast power supply for vertical stabilisation of TCV tokamak plasmas", Conf. on Electrical Power Technology in European Physics Research, CERN Geneva, 1995, 43E A-15 (1995)
12. J.B. Lister, Ph. Marmillod, P-F. Isoz, B. Marletaz and I.E. Piacentini, "Distributed control of the TCV tokamak and modular BITBUS nodes", Proc. 16th Symp. on Fusion Technology, 1268 London, 1990 (1991)
13. VSYSTEM, commercial product of Vista Control Systems, Los Alamos.
14. J.B. Lister, Y. Martin and J-M. Moret, "The dynamics of shaping control in TCV", 21st EPS Conf. on Contr. Fusion and Plasma Phys., Montpellier (1994)
15. I. Hutchinson, S. Horne, G. Tinios, S.M. Wolfe and R.S. Granetz, "Plasma shape control: A general approach and its application ot Alcator C-MOD", MIT Report, PFC/JA-95-4 (1995)
16. C. Kessel and S. Jardin, private communication
17. E.A. Lazarus, J.B. Lister and G.H. Nielson, "Control of the vertical instability in tokamaks", Nuclear Fusion, 30, 111 (1990)

18. F. Hofmann and S.C. Jardin, "Plasma shape and position control in highly elongated tokamaks", Nuclear Fusion 30, 2013 (1990)
19. F. Hofmann and G. Tonetti, "Fast identification of plasma boundary and X-points in elongated tokamaks", Nuclear Fusion 28, 519 (1988)
20. F. Hofmann, "FBT - A free boundary tokamak equilibrium code for highly elongated and shaped plasmas", Computer Physics Communications 48, 207(1988)
21. J.B. Lister, Y. Martin and J-M. Moret, "On locating the poloidal field coils for tokamak vertical position control", Lausanne Report, LRP 533/95 (1995), accepted for publication in Nuclear Fusion
22. J.B. Lister, E.A. Lazarus, A.G. Kellman, J-M. Moret, J.R. Ferron et al., "Experimental study of the vertical stability of high decay index plasmas in the DIII-D tokamak", Nuclear Fusion, 30, 2349 (1990)
23. D.J. Ward, S.J. Jardin and C.Z. Chen, "Calculations of axisymmetric stability of tokamak plasmas with active and passive feedback", Journal of Comp. Physics 104, 221 (1993)
24. J.-M. Moret, F. Buhlmann, F. Hofmann and G. Tonetti, "Consistency and precision of the TCV magnetic measurements", 22nd EPS Conf. on Cont. Fusion and Plasma Phys., Bournemouth, 1995, p. 417, vol. 19B part I. (1995)
25. M.J. Dutch, F. Hofmann, B.P. Duval, M. Corboz, A. Hirt et al., "ELM control during double-null ohmic H-modes in TCV", Nuclear Fusion, 35, 650 (1995)

TABLES

	U dio (V)	I d (kA)	n
1 x TOR	626	78	2 x 10
2 x OH	1400	31	2 x 8
8 x E	651	7.7	2 x 2
8 x F	1250	7.7	2 x 2

n = number of 6-pulse bridges

TABLE 4.1 *Characteristics of the 12-phase Power Rectifiers.*

Type	Turbo, 4 pole, air cooled
Duty cycle	3 seconds/5 minutes
Nominal power	220 MVA
Nominal voltage	10 kV
Voltage fluctuation	+/- 2 %, for the nominal pulse
Short circuit power	1420 MVA (120 Hz, saturated)
Initial speed	3600 rpm
Frequency range	120 Hz - 96 Hz
Max. stored energy	340 MJ
Extractable energy	122 MJ (120 Hz - 96 Hz)
Total weight	146 tons

TABLE 4.2 *Delivered Turbo-generator characteristics.*

Input Voltage	U_{ac}	10 kV - 120 Hz (5.2 kV - 50 Hz)
Output no load Voltage	U_{out}	± 566 V (= √2 • 400 V)
NB. of Output Voltage Steps		3 (- 566 V / 0 / + 566 V)
Max. Output Current	I_{out}	± 2000 A (2 s / 5 min.)
Voltage Max. Turn-on Time	t_{on}	5 µs
Voltage Max. Turn-off Time	t_{off}	5 µs
Voltage Min. On-time	$t_{on\ min}$	5 µs
Voltage Min. Off-time	$t_{off1\ on2}$	5 µs
Current Max. Ripple and controllability	$I_{\sigma N}$ $I_{\sigma \ll}$	200 A p-p for 2000 A 40 A p-p for 50 A
Max. Current Max. Rise Time from 0 to 2 kA	$t_{I\ rise}$	250 µs (70 mΩ - 50 µH load)

TABLE 4.3 *Fast Power Supply Parameters*

H-Bridge: Output Voltage	U_{out}	± 566 V
		(- 566V / 0 / + 566 V)
Output Current	I_{out}	± 2000 A (2 s / 5 min.)
		± 500 A (continuously)
Switching Frequency: typical	f_e	10 kHz
Max		30 kHz
Crowbar: Serial Resistance	R_{cr}	50 mΩ
Threshold Voltage	U_{trig}	± 750 V
Electrodynamical Withstand	$I_{disrupt}$	25 kA
Thermal Withstand	I^2t	10 MA^2s / 5 min.
Output Filter: Capacitance	C_f	0.3 µF
Resistance	R_f	10 Ω
Inductance	L_f	2 • 1 µH
(H-Bridge Output)		

TABLE 4.4 *Fast Output Stage Parameters*

START of CYCLE
Wait for the pilot to request the next discharge
Initialise the MDS+ Data Acquisition
Initialise the Hybrid Matrix data
Initialise the Waveform Generators
Force a refresh of the live database on all input devices
Check that all requested sub-systems are ready
Check that all automated systems are ready
Set the plant status interrogation to medium speed
Give a start command to all automated systems (e.g. Motor Generator)
Set up the Plasma Control System hardware
Wait until the Data Acquisition has completed initialisation
Initialise the Plasma Control System hardware
Wait until all of the automated systems are ready
Load the pulse sequencing timers
Force a refresh of the live database on all input devices
Check that all requested autonomous or automated systems are ready
Inform pilot that the discharge is ready to fire
Allow pilot 15 seconds to initiate the tokamak pulse
Force a final refresh of the live database on all input devices
Check that all requested autonomous or automated systems are ready
Fire the discharge sequence timer
Wait until the end of the timer
Execute specific close-down actions
Start the MDS+ Data Acquisition
Wait for the Data Acquisition to complete
Execute post-cycle actions
END of CYCLE

TABLE 5.1 *Sequence of Actions in the TCV Cycle Sequencer. All actions must complete before continuing. These actions comprise the whole TCV sequencer, illustrating the simplicity resulting from the design choice.*

Define the observers to be used in the discharge

Apply a feedback controller algorithm

Generate the plasma control reference waveforms

Open or close the feedback gates and set the feedback gains

Select the power supply mode of operation

Generate the pre-programmed waveforms

Define the time-windows for the power supplies

Define the time-windows for the data acquisition

Select the sub-systems necessary for this discharge

TABLE 5.2 *Procedural actions to be carried out in sequence during the preparation of a TCV discharge.*

	1	2	3	4	5	6
Coil-Pair	$\sigma(zI_p)$	B_R T*10^6	k-mut	1./sep m^{-1}	VA VA	P_{oh} W
E5-E4	945	32	0.11	2.6	0.21	0.005
E6-E3	1213	57	0.25	2.1	0.11	0.004
E7-E2	612	46	0.37	1.6	0.42	0.008
E8-E1	274	32	0.43	1.3	2.0	0.038
F8-F1	165	59	0.45	1.1	3.1	0.032
F7-F2	255	71	0.47	1.3	1.3	0.017
F6-F3	543	78	0.32	1.9	0.29	0.005
F5-F4	393	52	0.19	2.2	1.0	0.007

TABLE 8.1 *Experimental results of the efficiency of the different poloidal field coil pairs during continuous random injection experiments for all coil-pairs: 1) $\sigma(zI_p)$ for +-200 Volt excitation on each coil-pair; 2) Radial field on axis per unit coil-pair current; 3) Coupling coefficient between the coil-pair and the slowest anti-symmetric eigenmode of the vessel; 4) Inverse separation of the coil currents from the magnetic axis; 5) Reactive power in the coil-pair for a unit RMS displacement; 6) Power dissipated in the vacuum vessel for unit $\sigma(zI_p)$.*

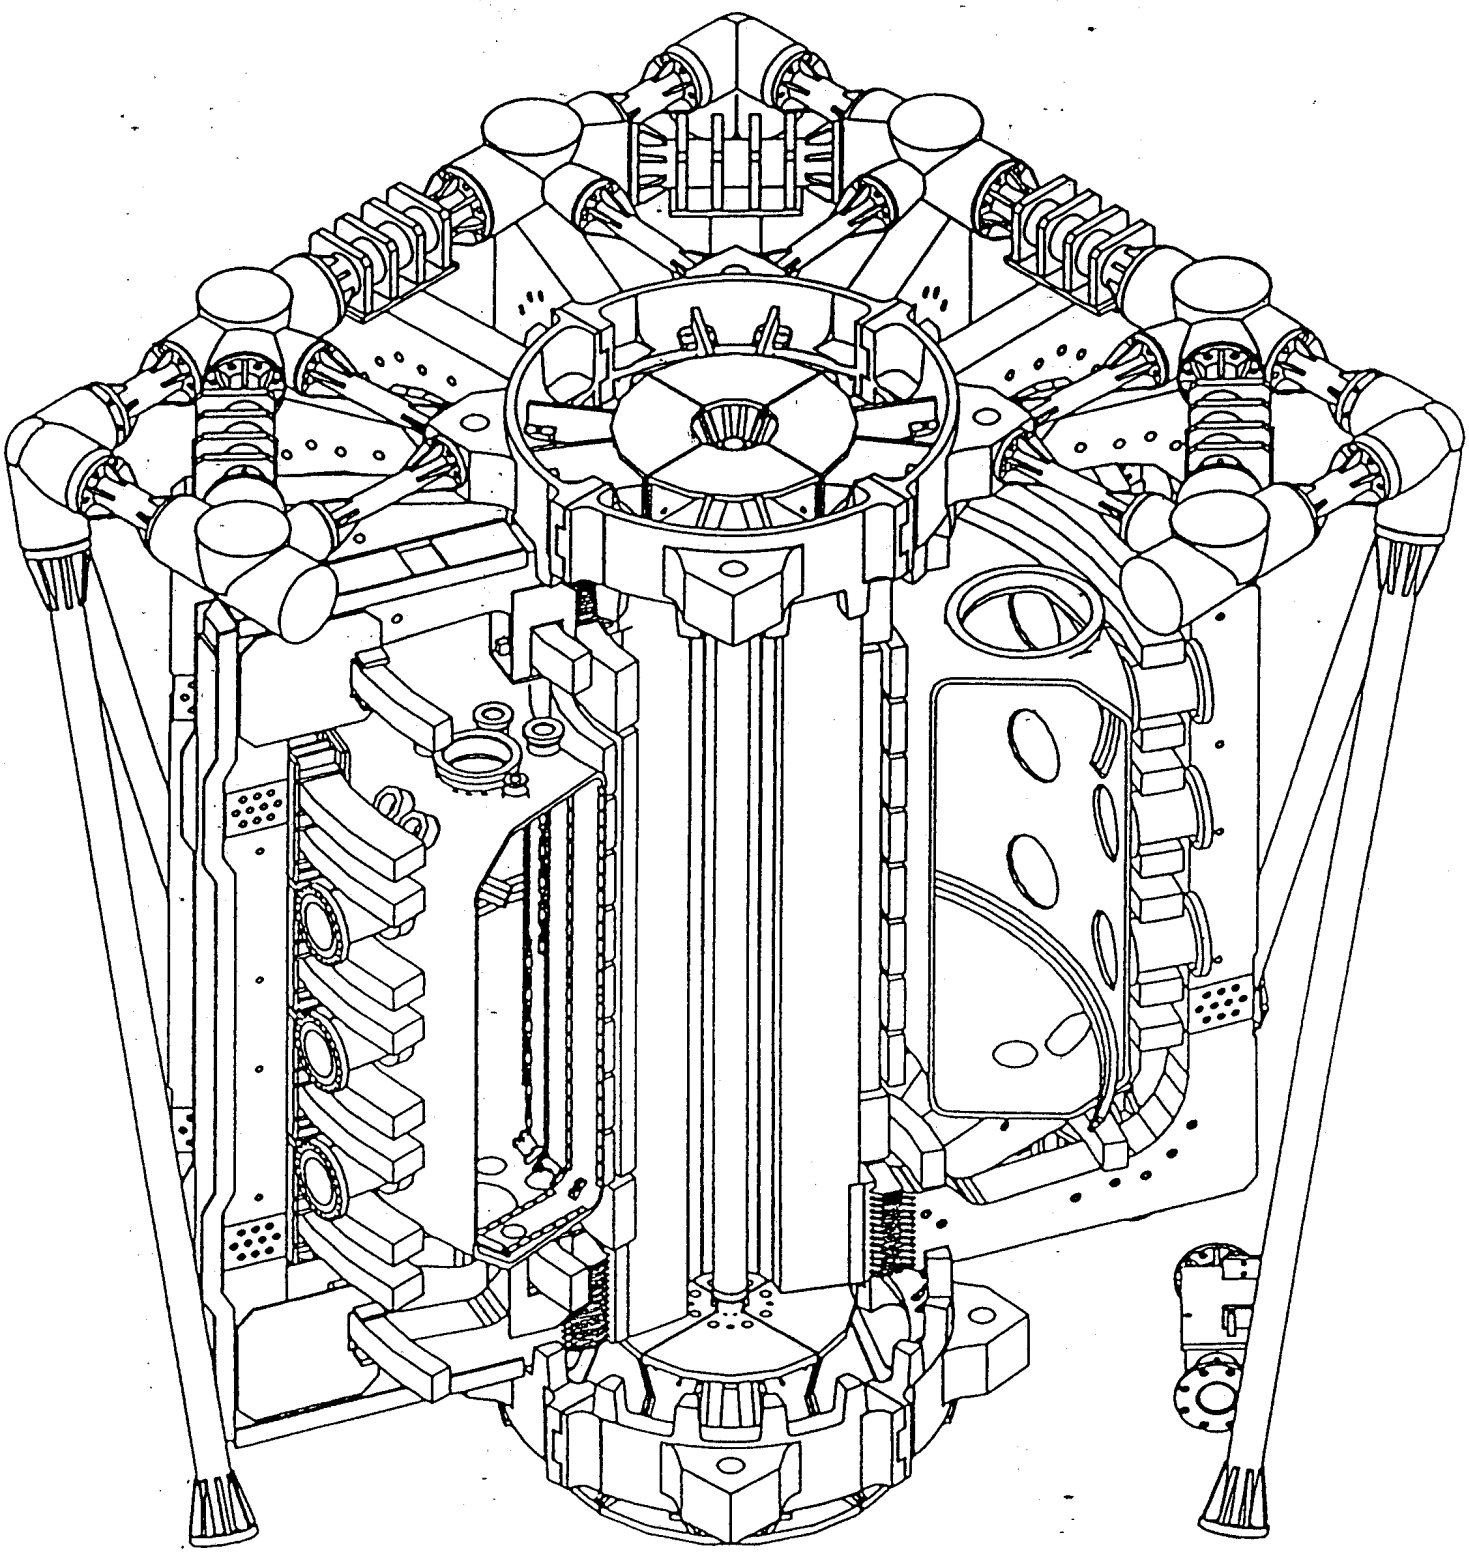

FIGURE 1 *The construction of the TCV tokamak, shown as a cut-away drawing.*

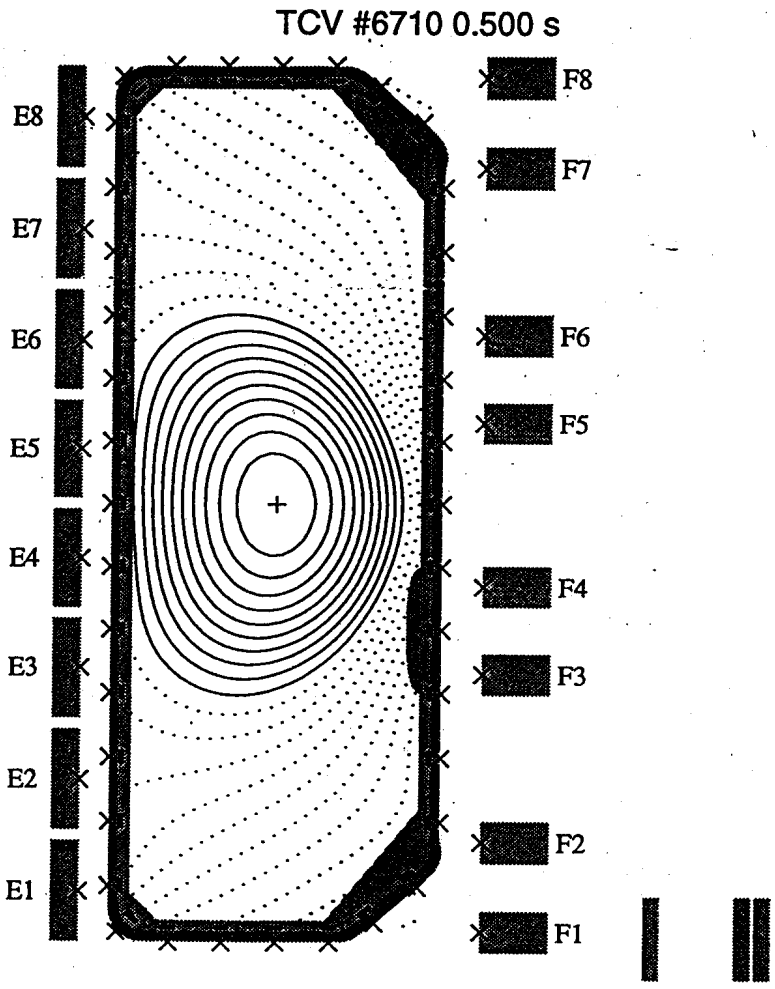

FIGURE 2 *The TCV shaping system, showing the poloidal field coil labelling, the position of the flux-loops (marked x on the vessel) and field probes (marked - inside the tiles). The plasma shown (# 6710) is a weakly-shaped limited plasma with the following parameters : B_ϕ = 1.43 T, I_p = 200 kA, δ = 0.21, κ = 1.45, q_a = 4.6.*

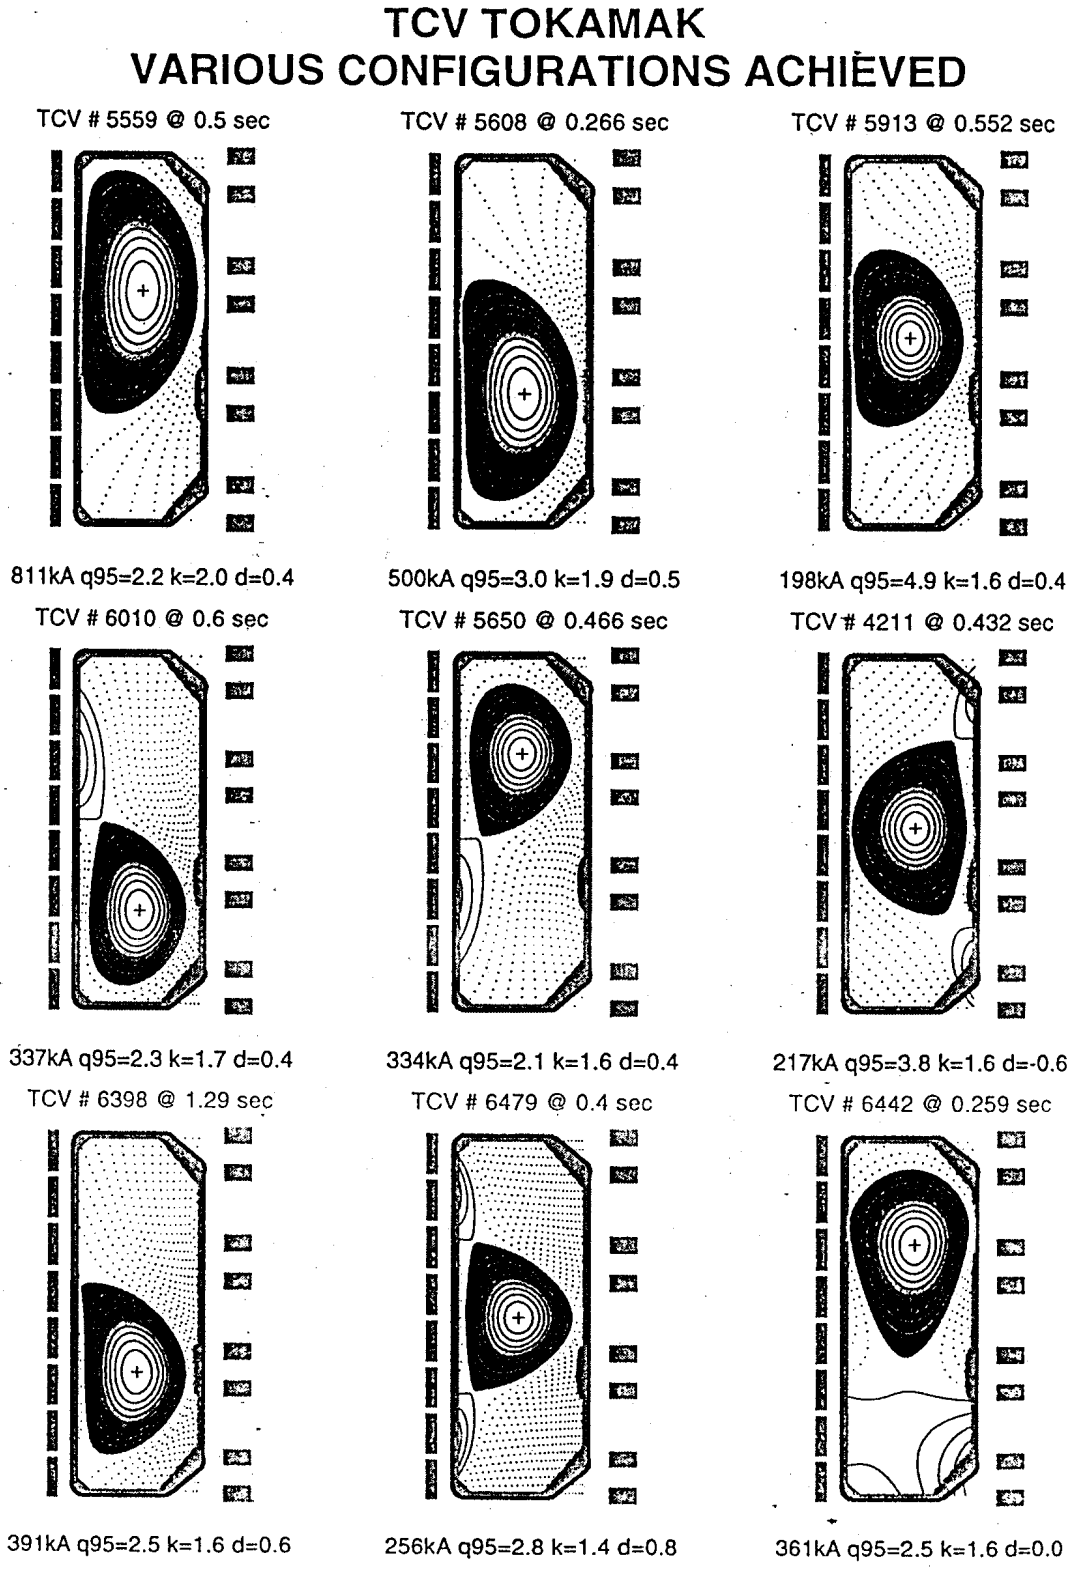

FIGURE 3 *Equilibria obtained on the TCV tokamak during the initial ohmic phase of operation.*

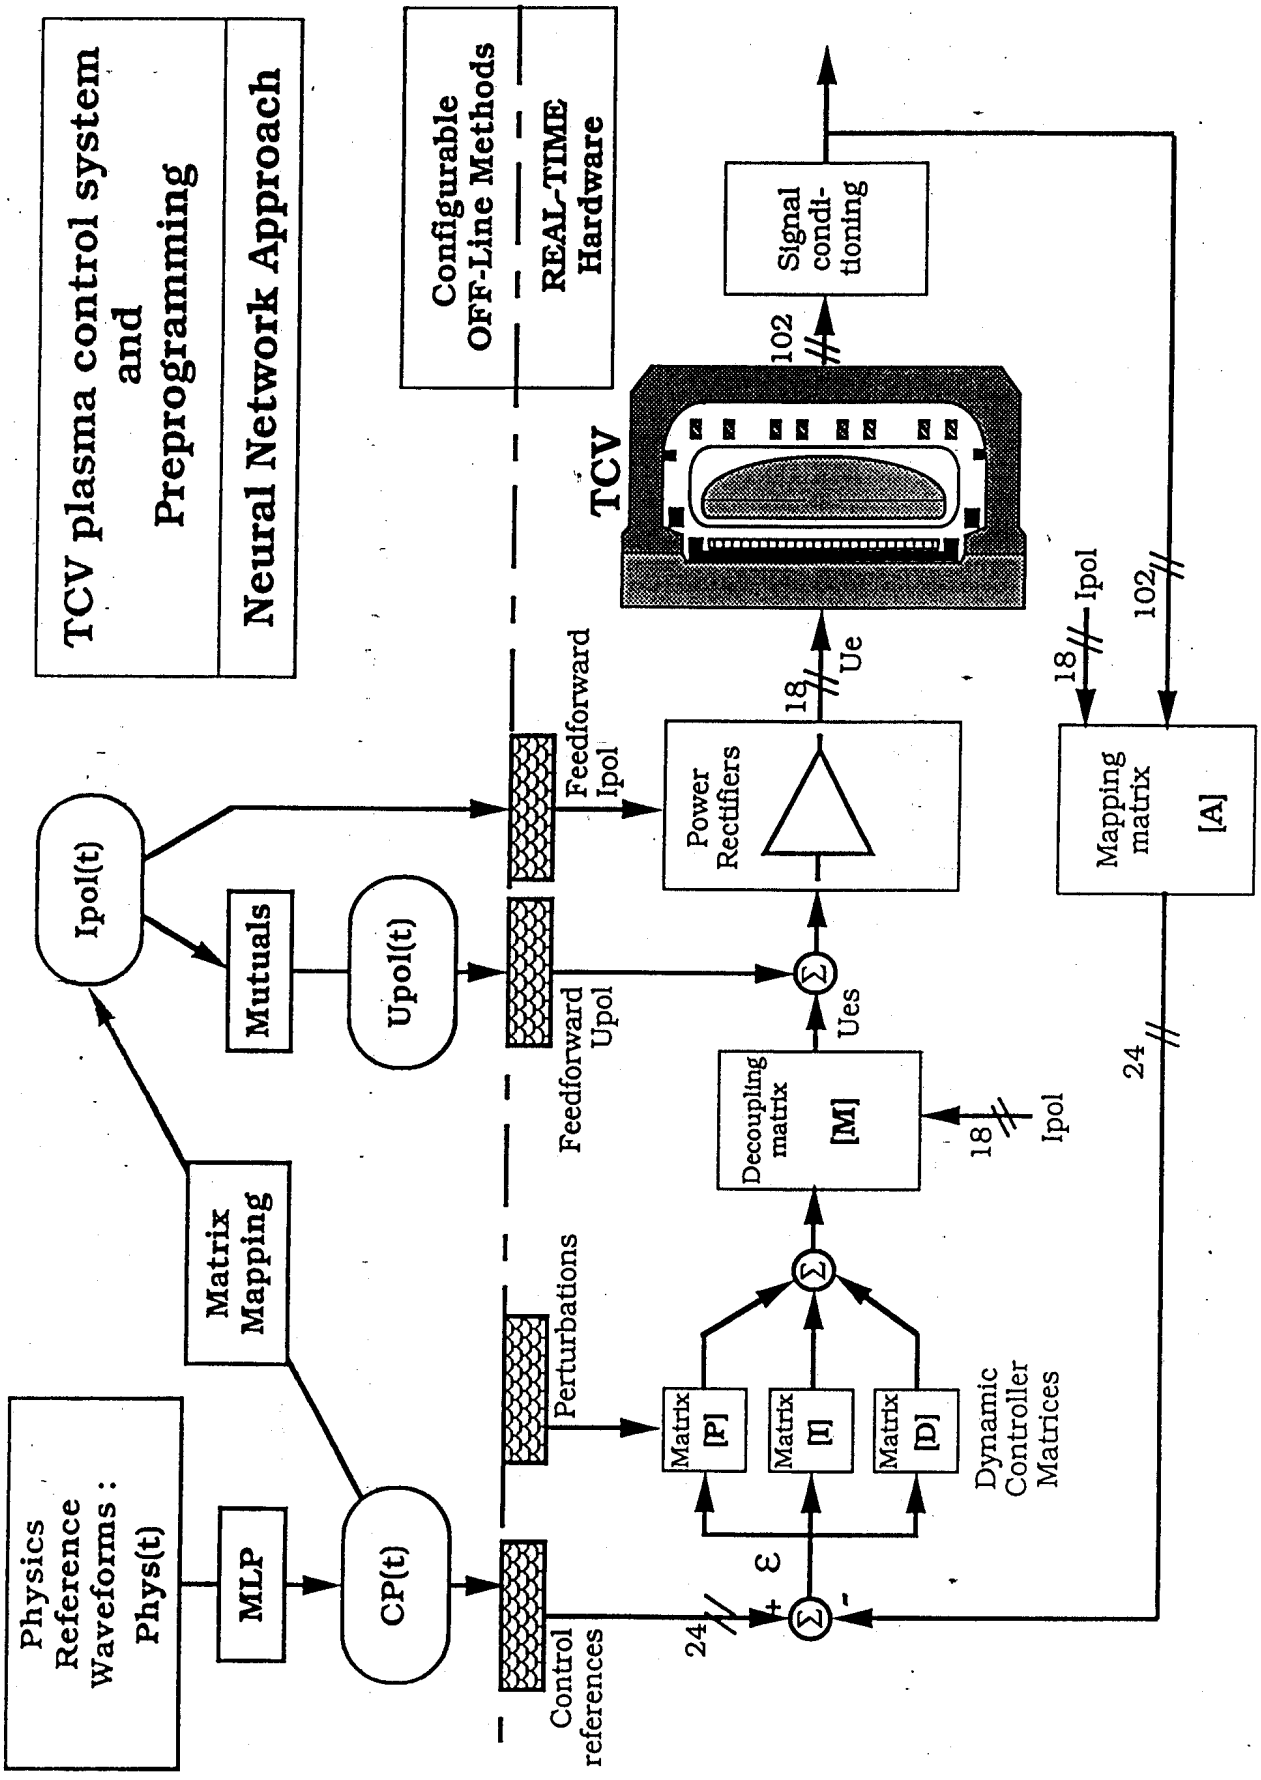

FIGURE 4 *Schematic of the PCS Control Loop showing the relationship between the hardware elements, namely the waveform generators, matrix multipliers and gates and the software which is responsible for the generation of the data which are automatically downloaded into the hardware.*

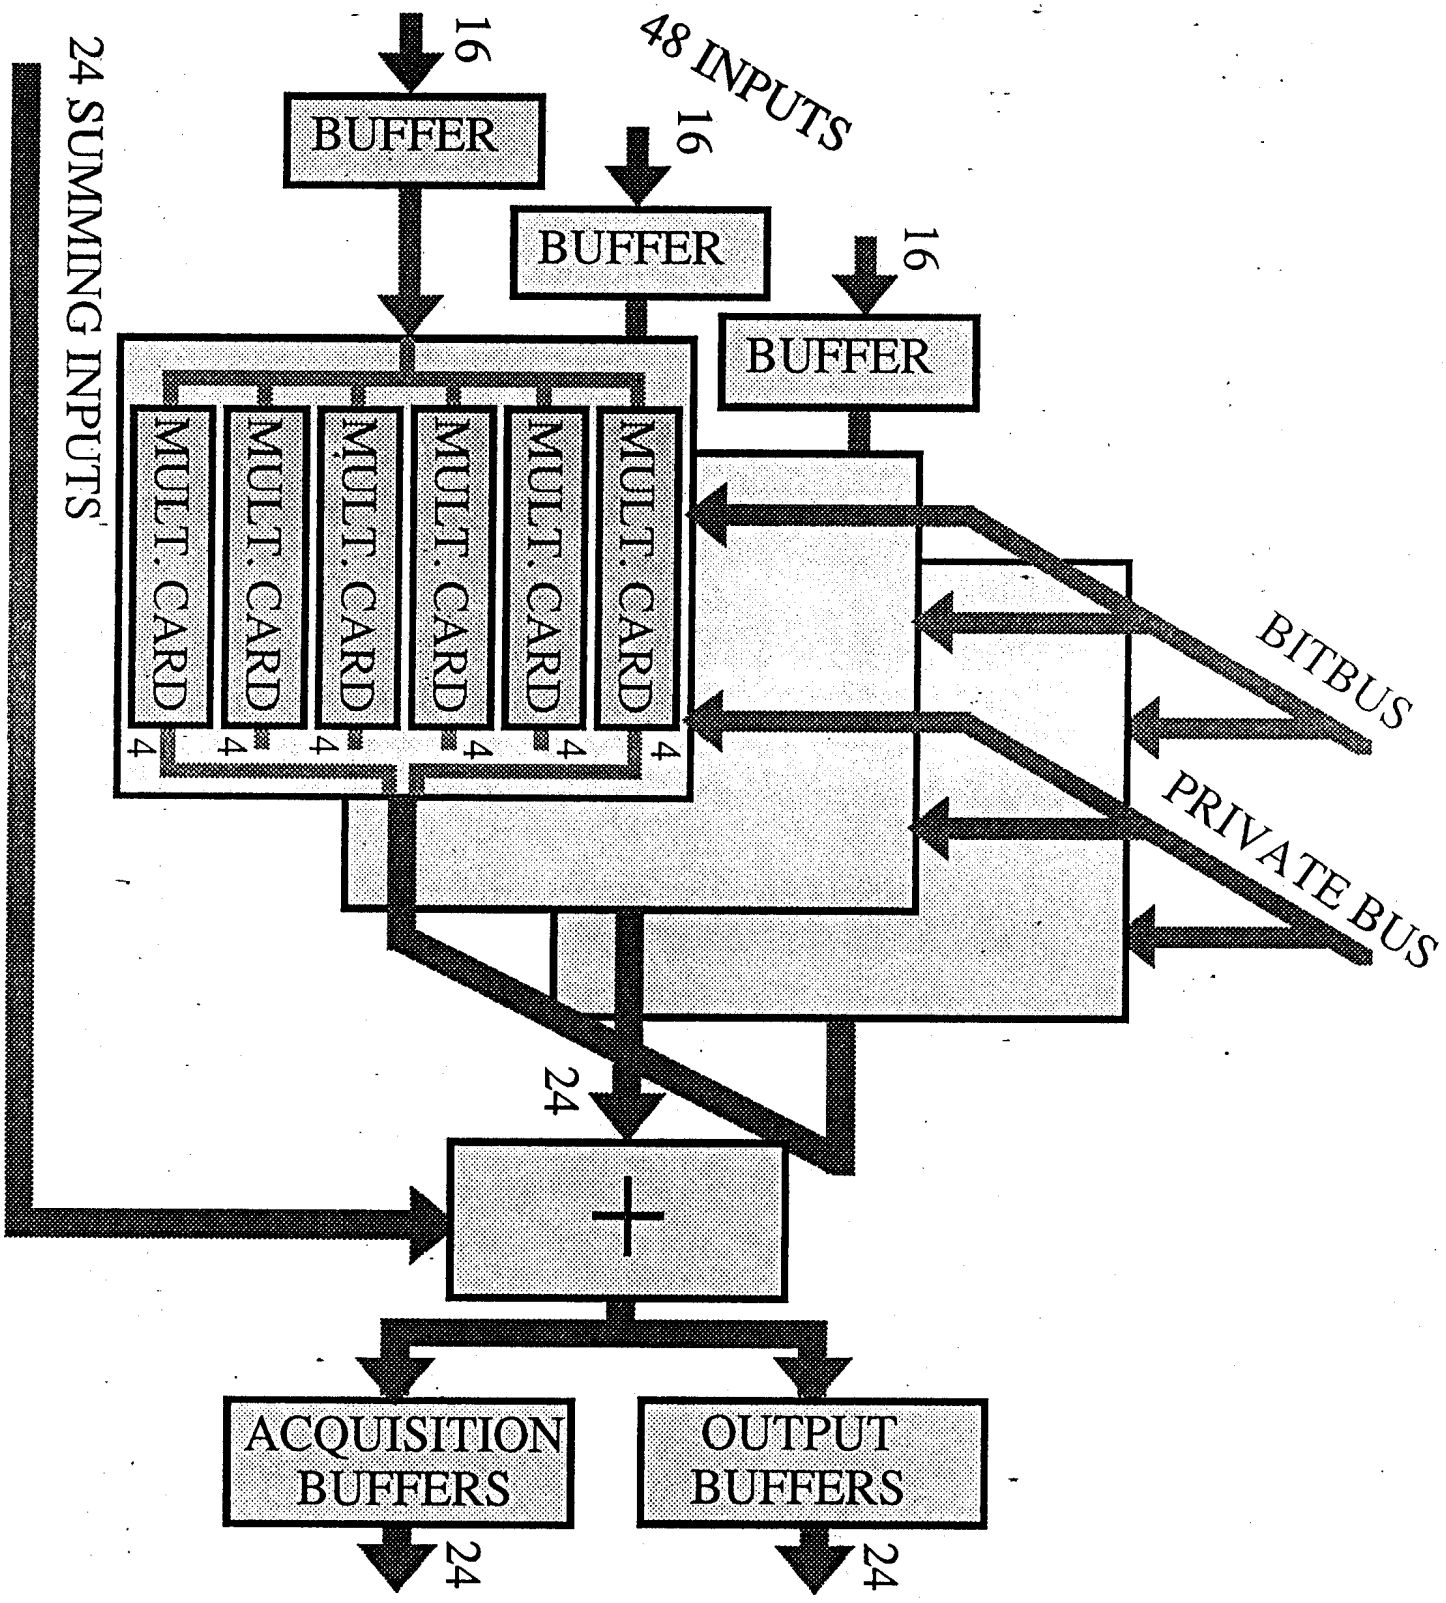

FIGURE 5 *Schematic of the hybrid multiplier modularity.*

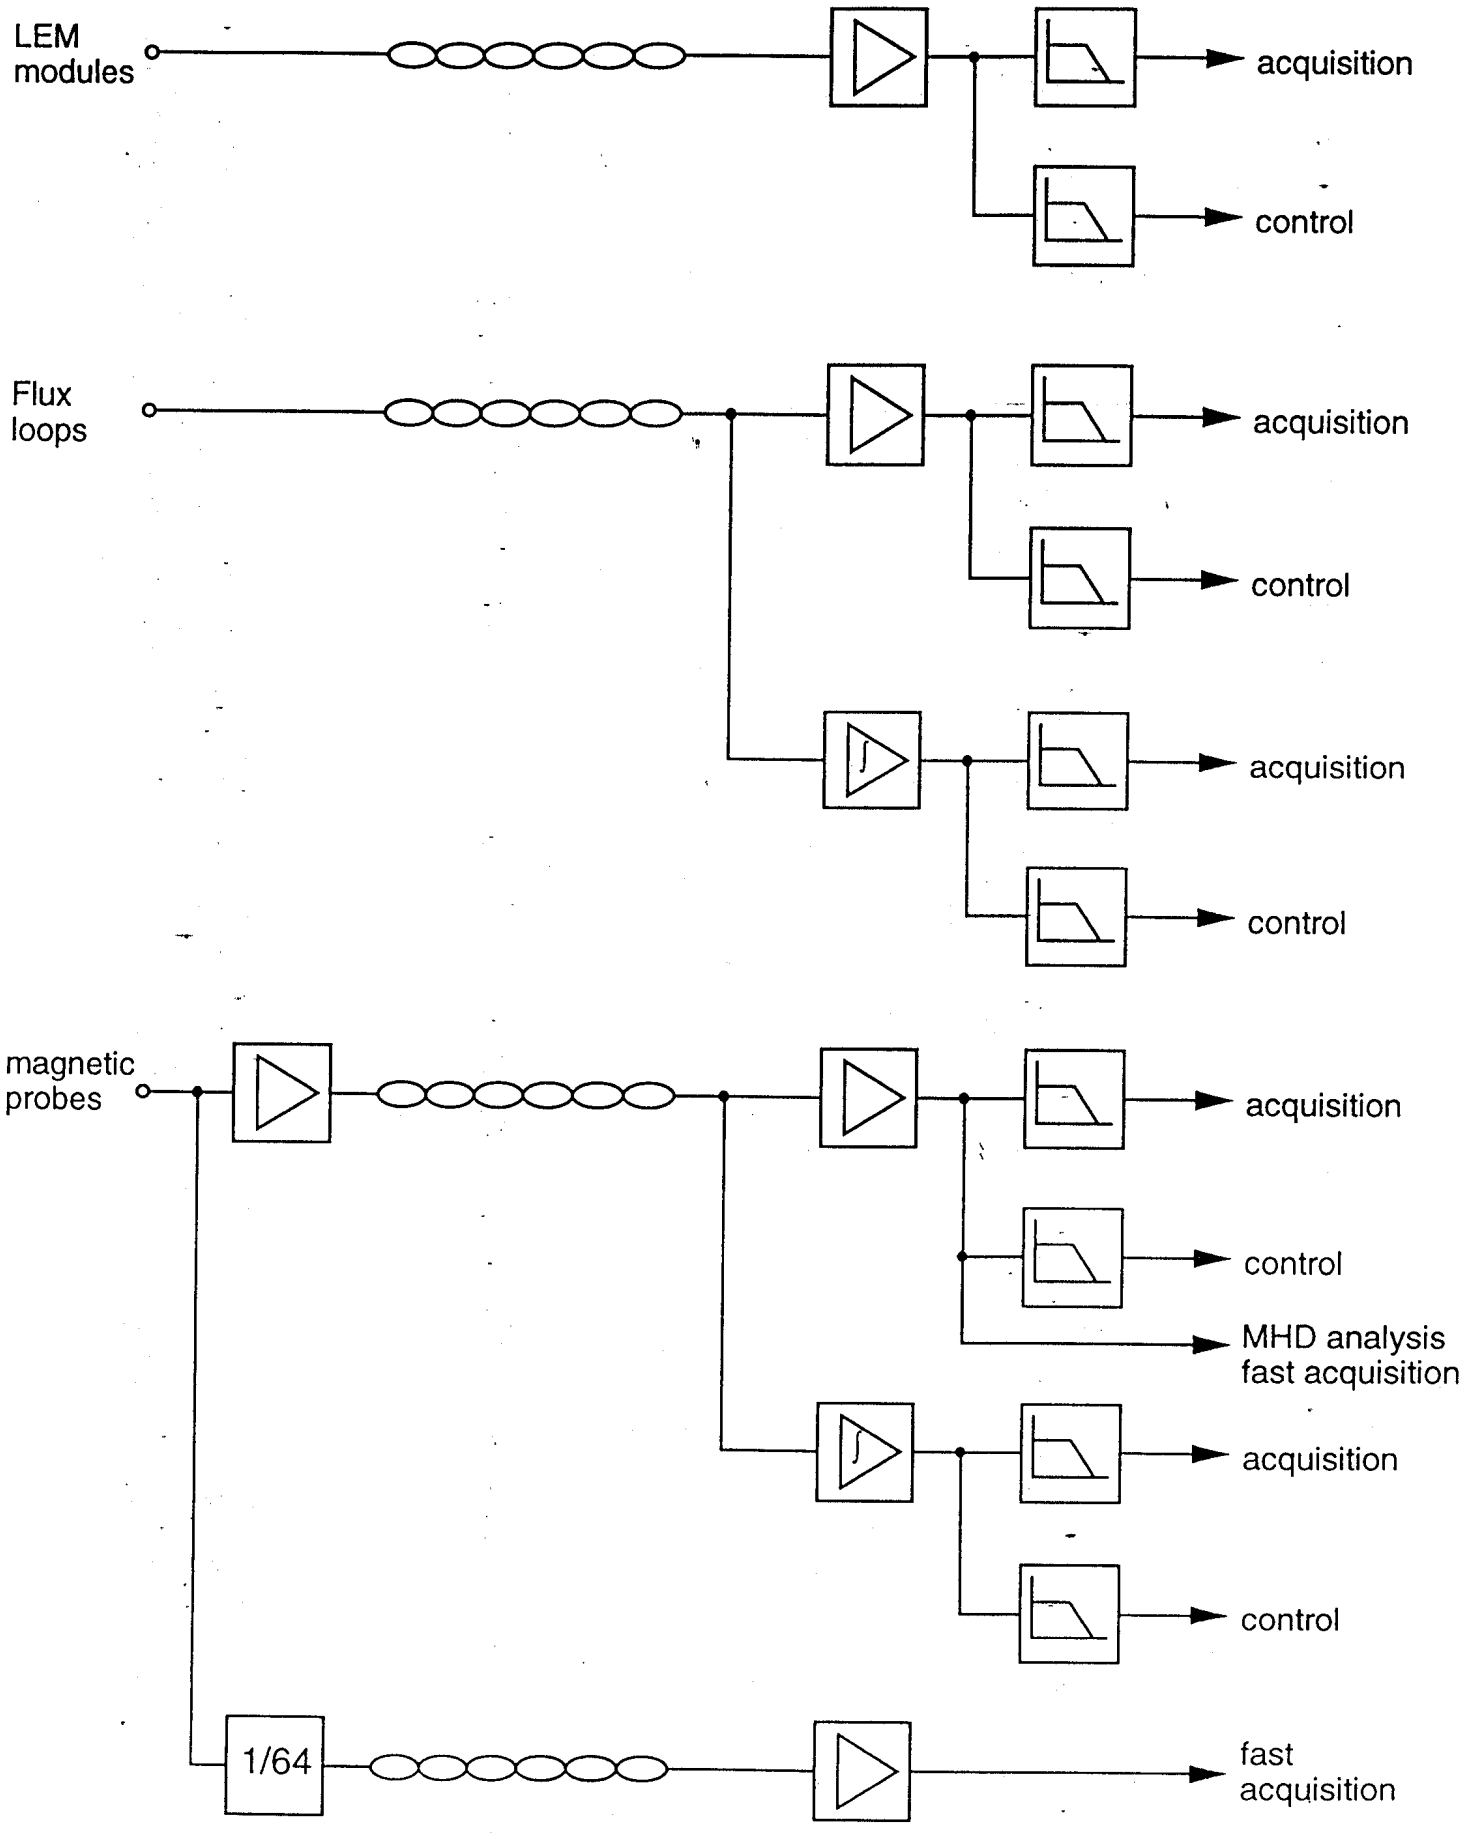

FIGURE 6 *Block diagram of the amplifying electronics for all magnetic signals.*

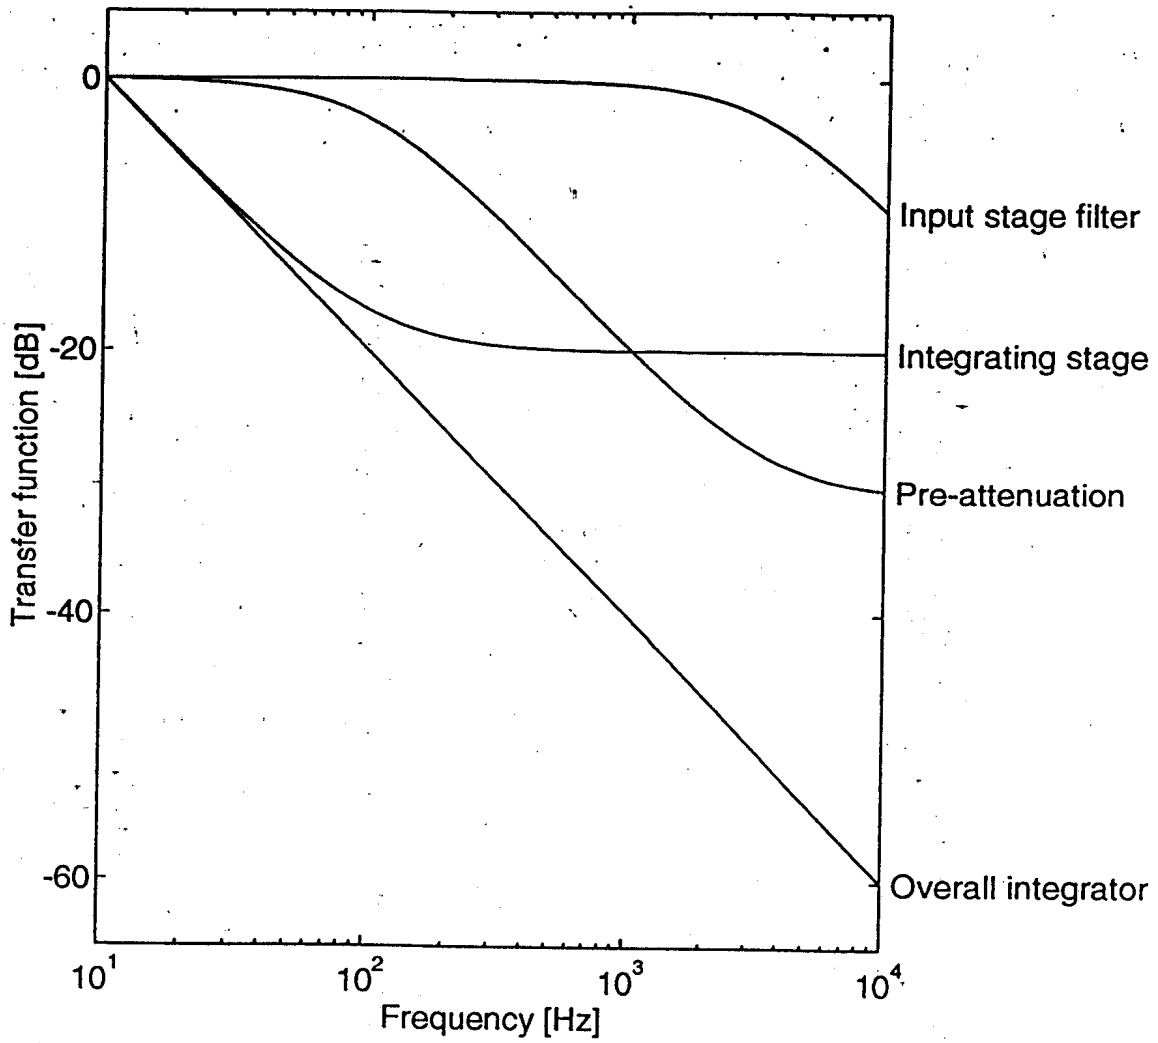

FIGURE 7 *Amplitude of the transfer function of the successive stages of the magnetic probe amplifying chain.*

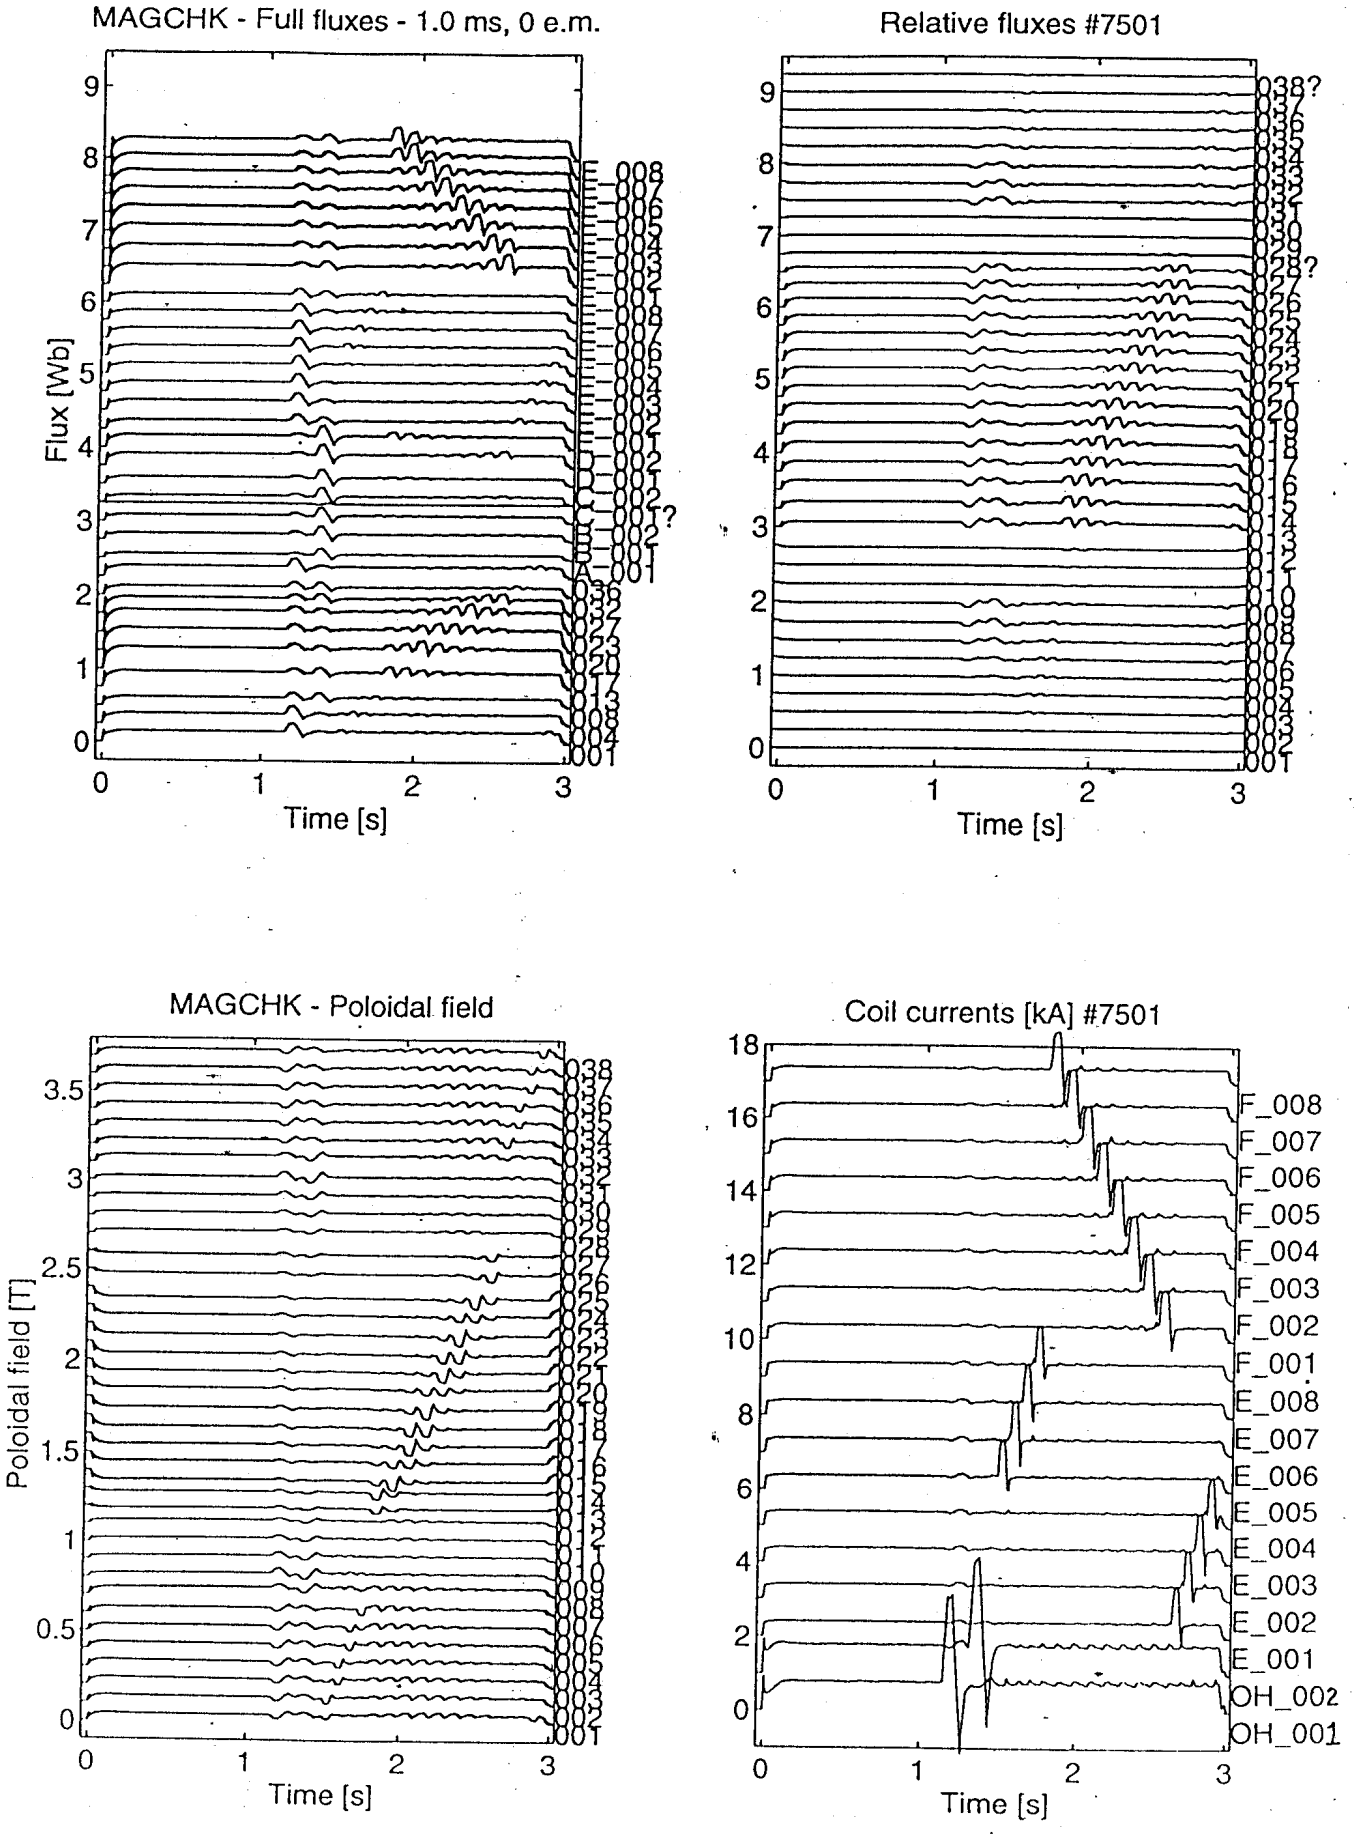

FIGURE 8 *Test signals used for daily check of the whole diagnostic. Poloidal coils are powered separately with a characteristic current spike. The magnetic signals are then compared with their theoretical value to detect any malfunction. Top Left: Fluxes measureed on the poloidal field coils (origins offset for display). Top Right: Flux with respect to inboard midplance loop 001 (Weber, origins are offset for display). Bottom Left: Poloidal field (Tesla). Bottom Right: Coil currents (kA)*

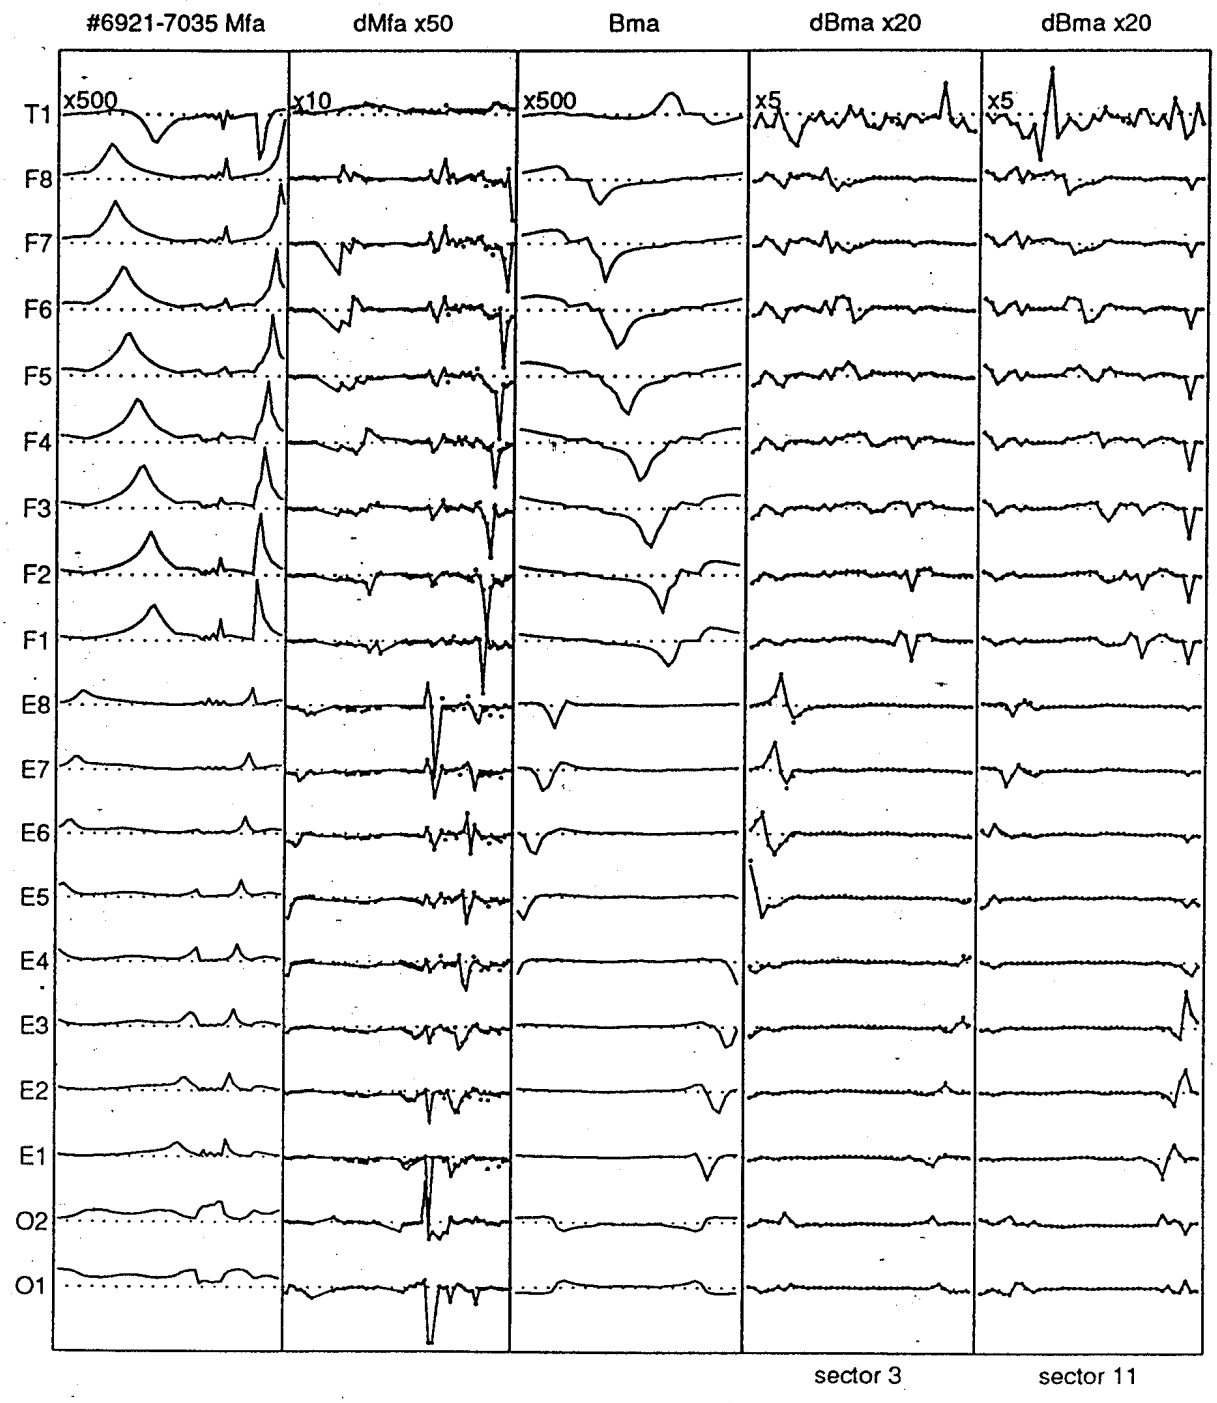

FIGURE 9 *Comparison of the expected flux ($M_{fa}I_a$) and field values ($B_{ma}I_a$) and their difference with respect to the corresponding measurements. Each curve consists of the values as a function of the loop or probe number for each of the 19 coils or coil sets. M_{fa} and B_{me} are the theoretical values, multiplied by 500 for the toroidal connections. Dots show the differences between these values and the measurements, on an expanded scale (50x for the fluxes and 20x for the magnetic fields). Superposed on the dots (solid line) are the differences implied by the derived corrections to the various calibrations and coil, loop and probe positions. The scales are constant from coil to coil (arb.).*

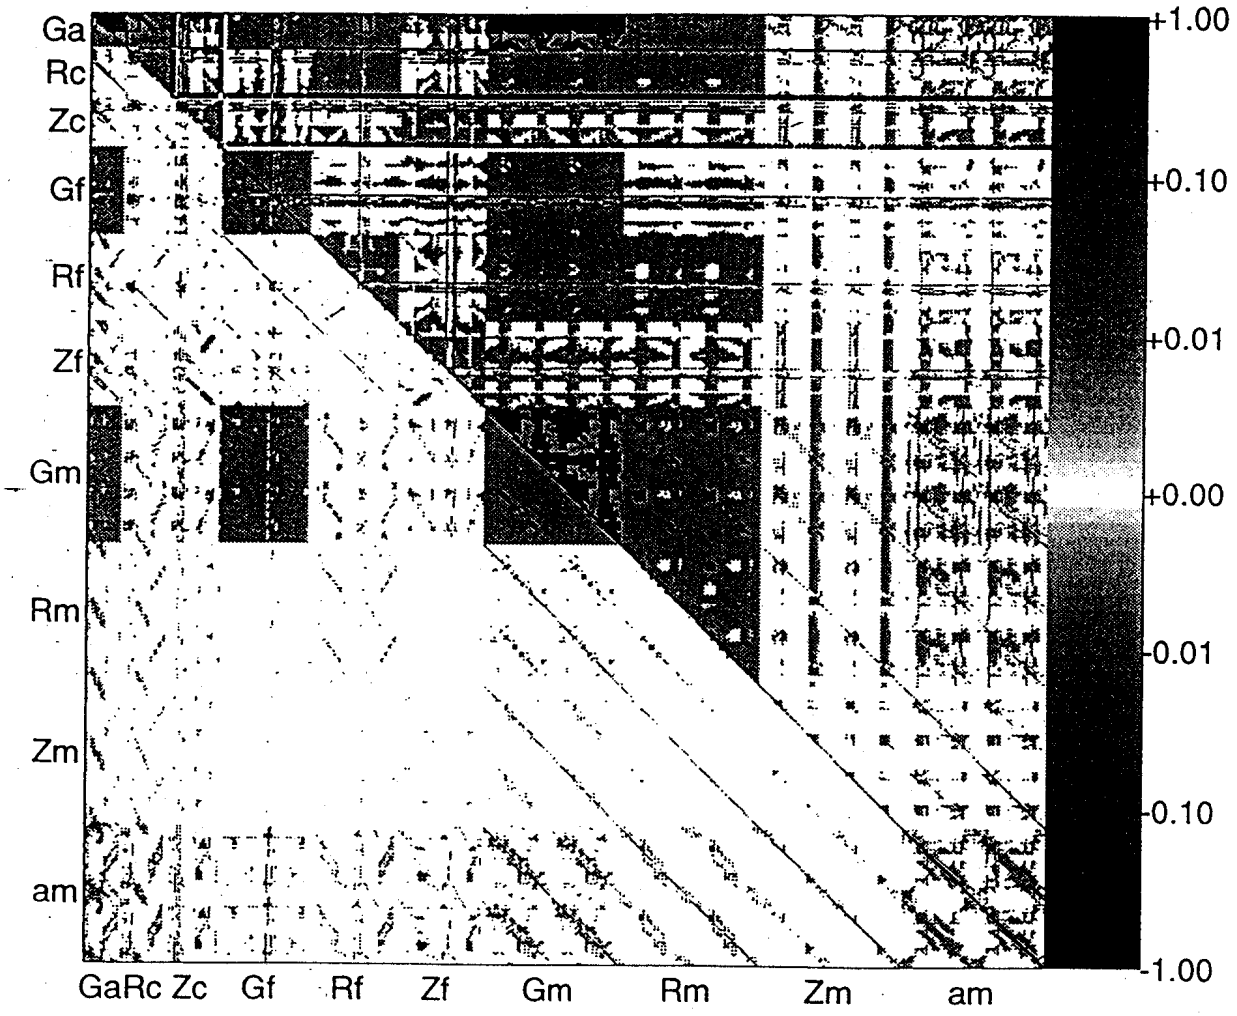

FIGURE 10 *Correction covariance matrix: upper and lower triangle with free and small parameter cost function respectively.*

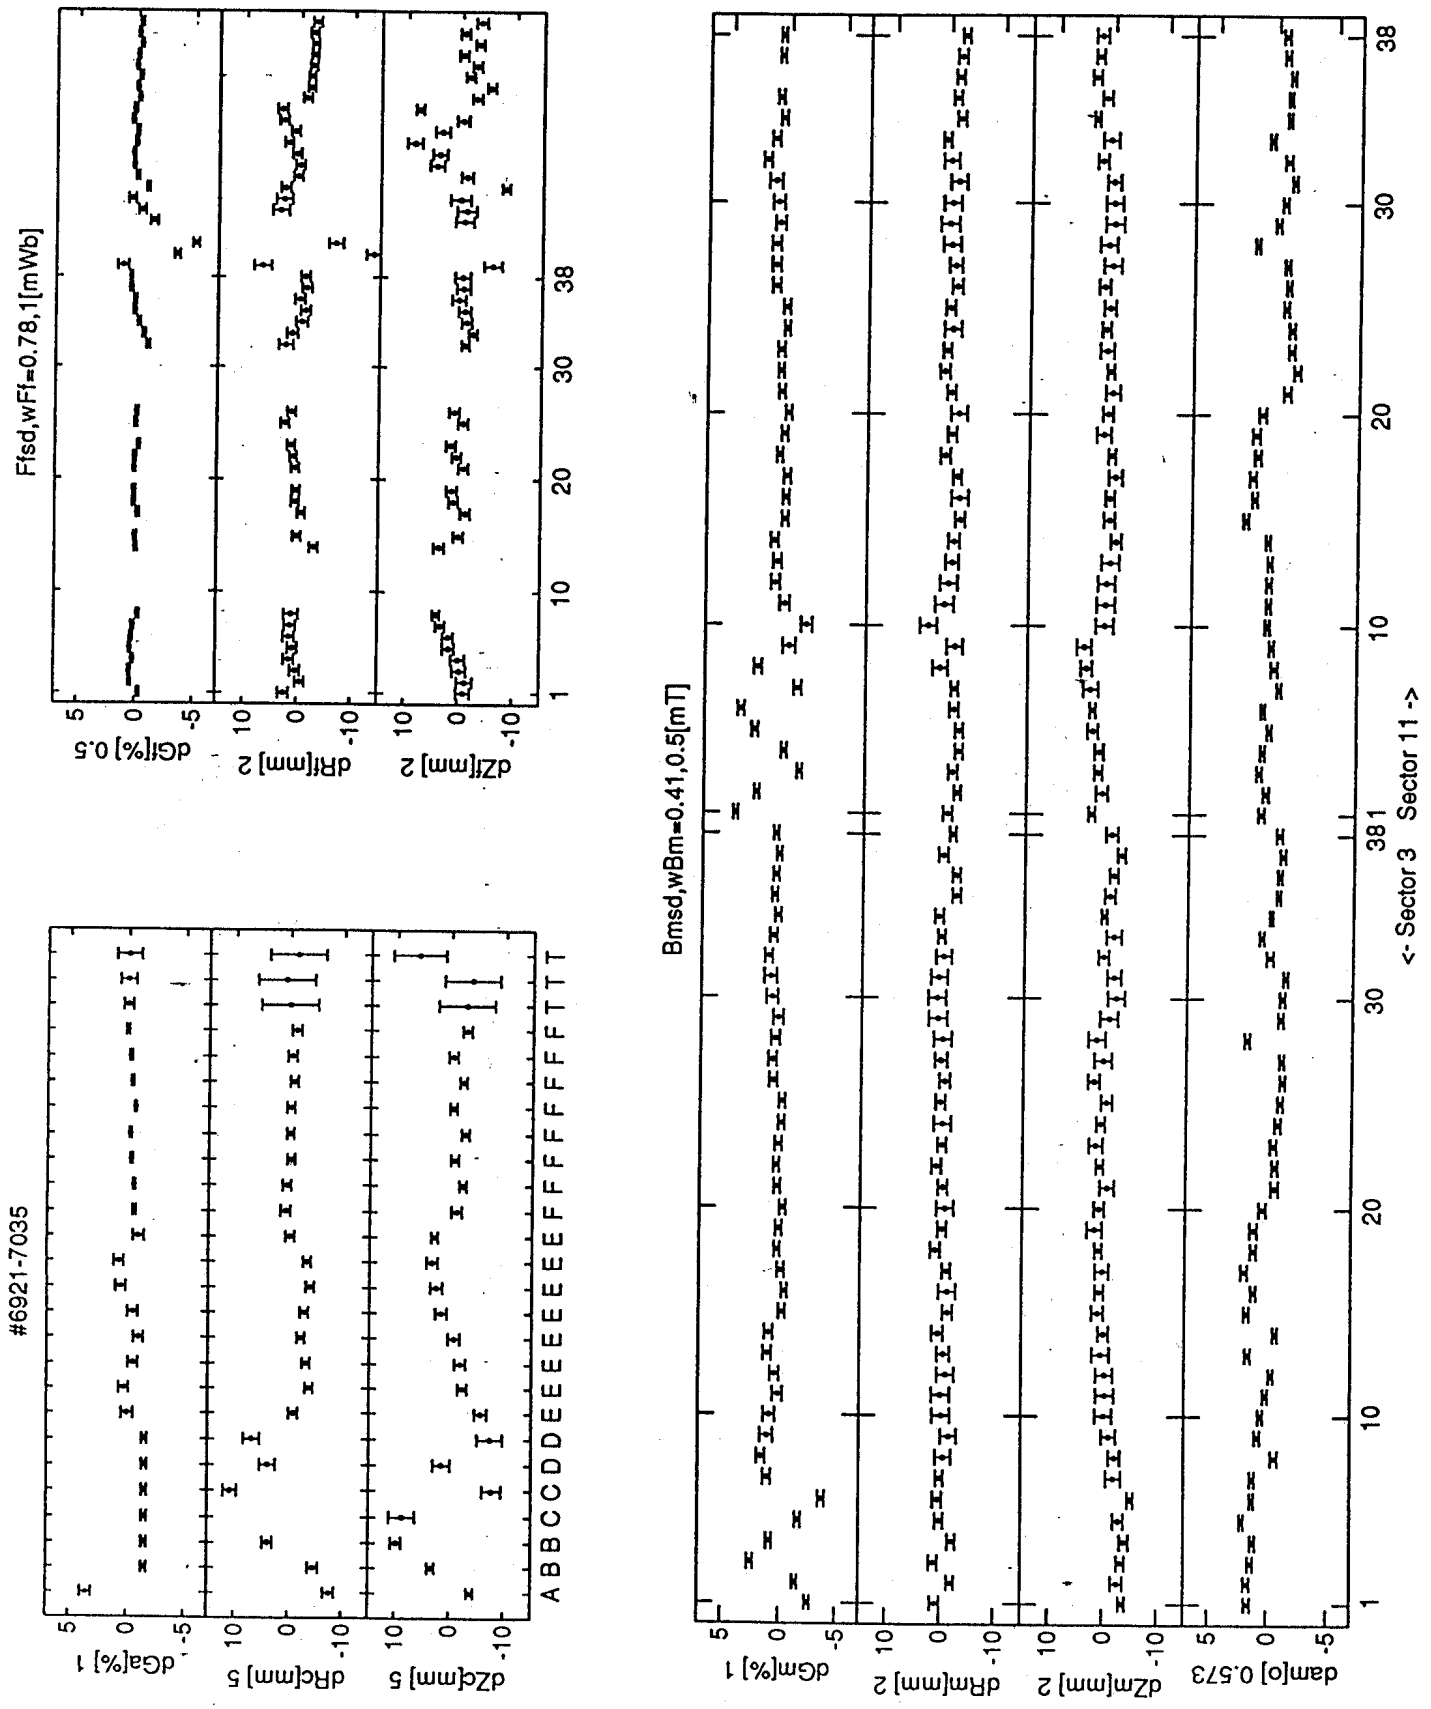

FIGURE 11 *Values of the corrections derived with a modified cost function that constrains the corrections to be small.*

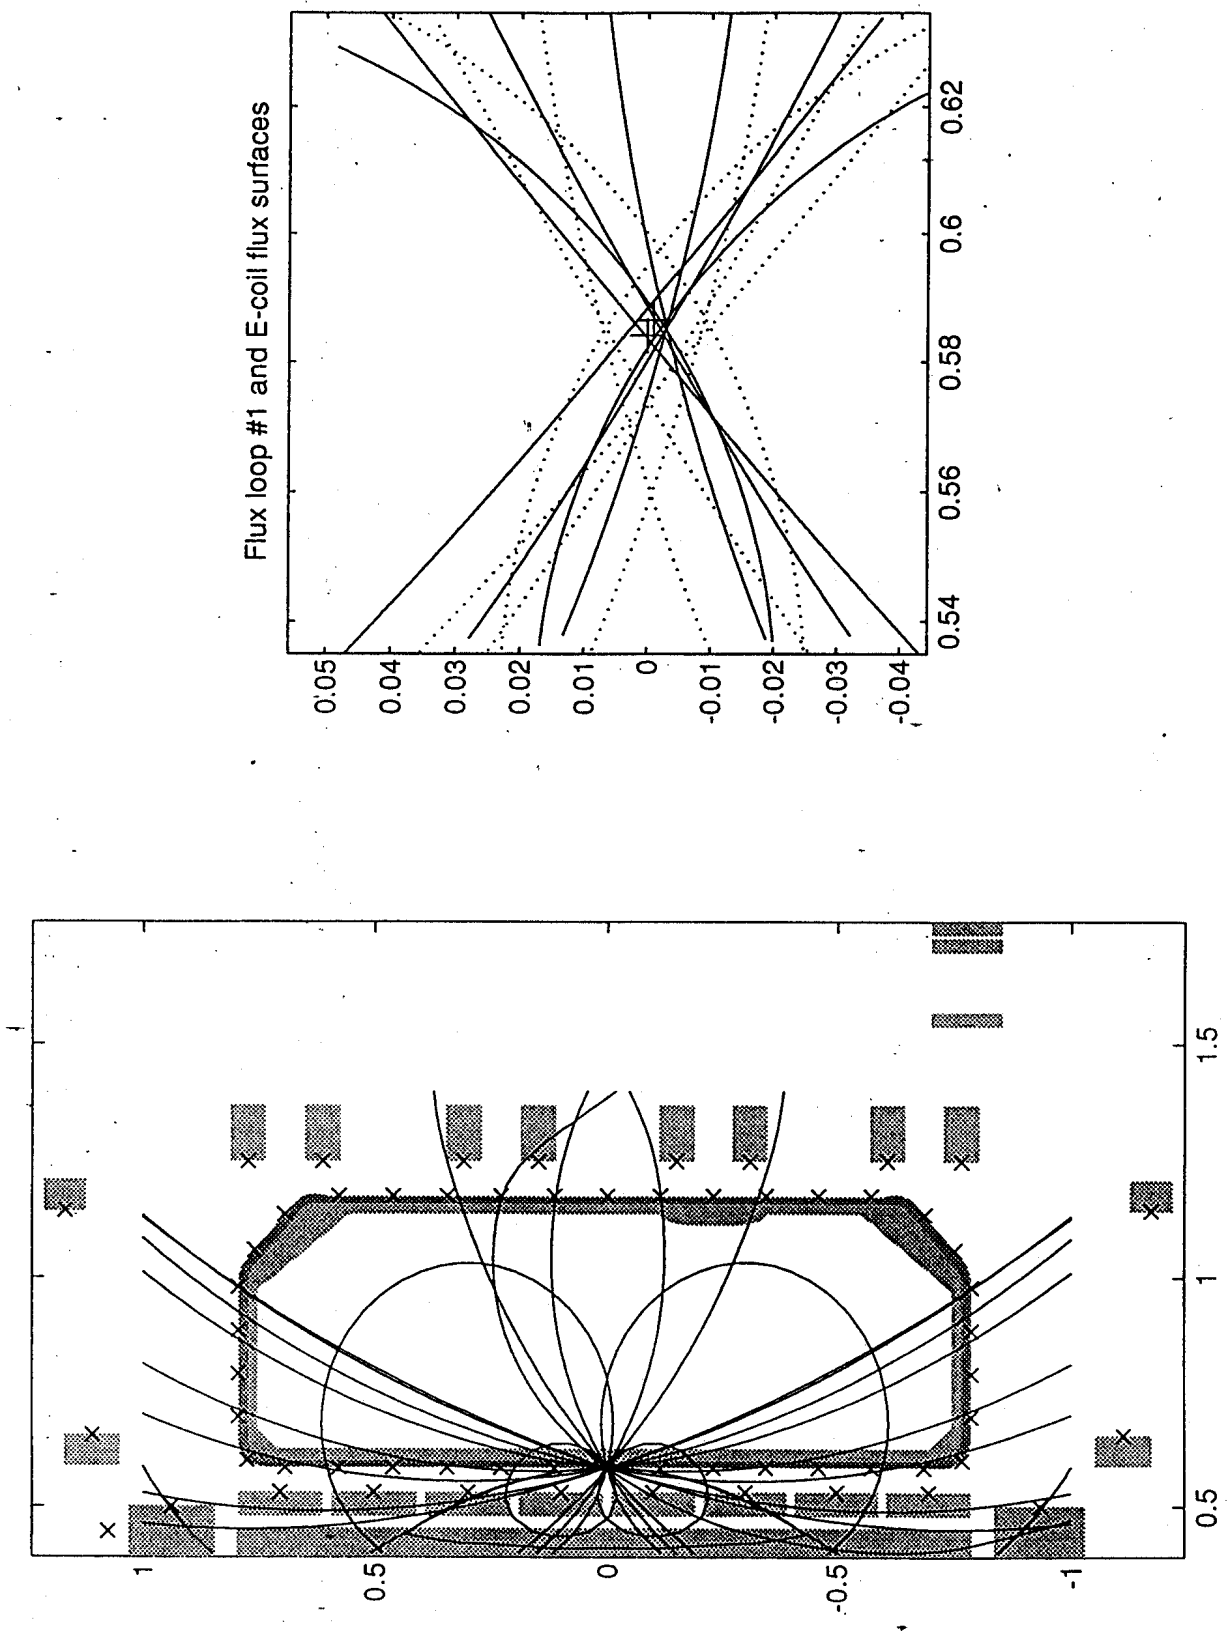

FIGURE 12 *Magnetic triangulation. The constant flux contours due to each E and F coil at the flux level measured by the inner equatorial loop cross at the loop location. On the enlarged plot these contours and the loop location are shown before applying the corrections (dotted lines) and with the corrections (solid lines).*

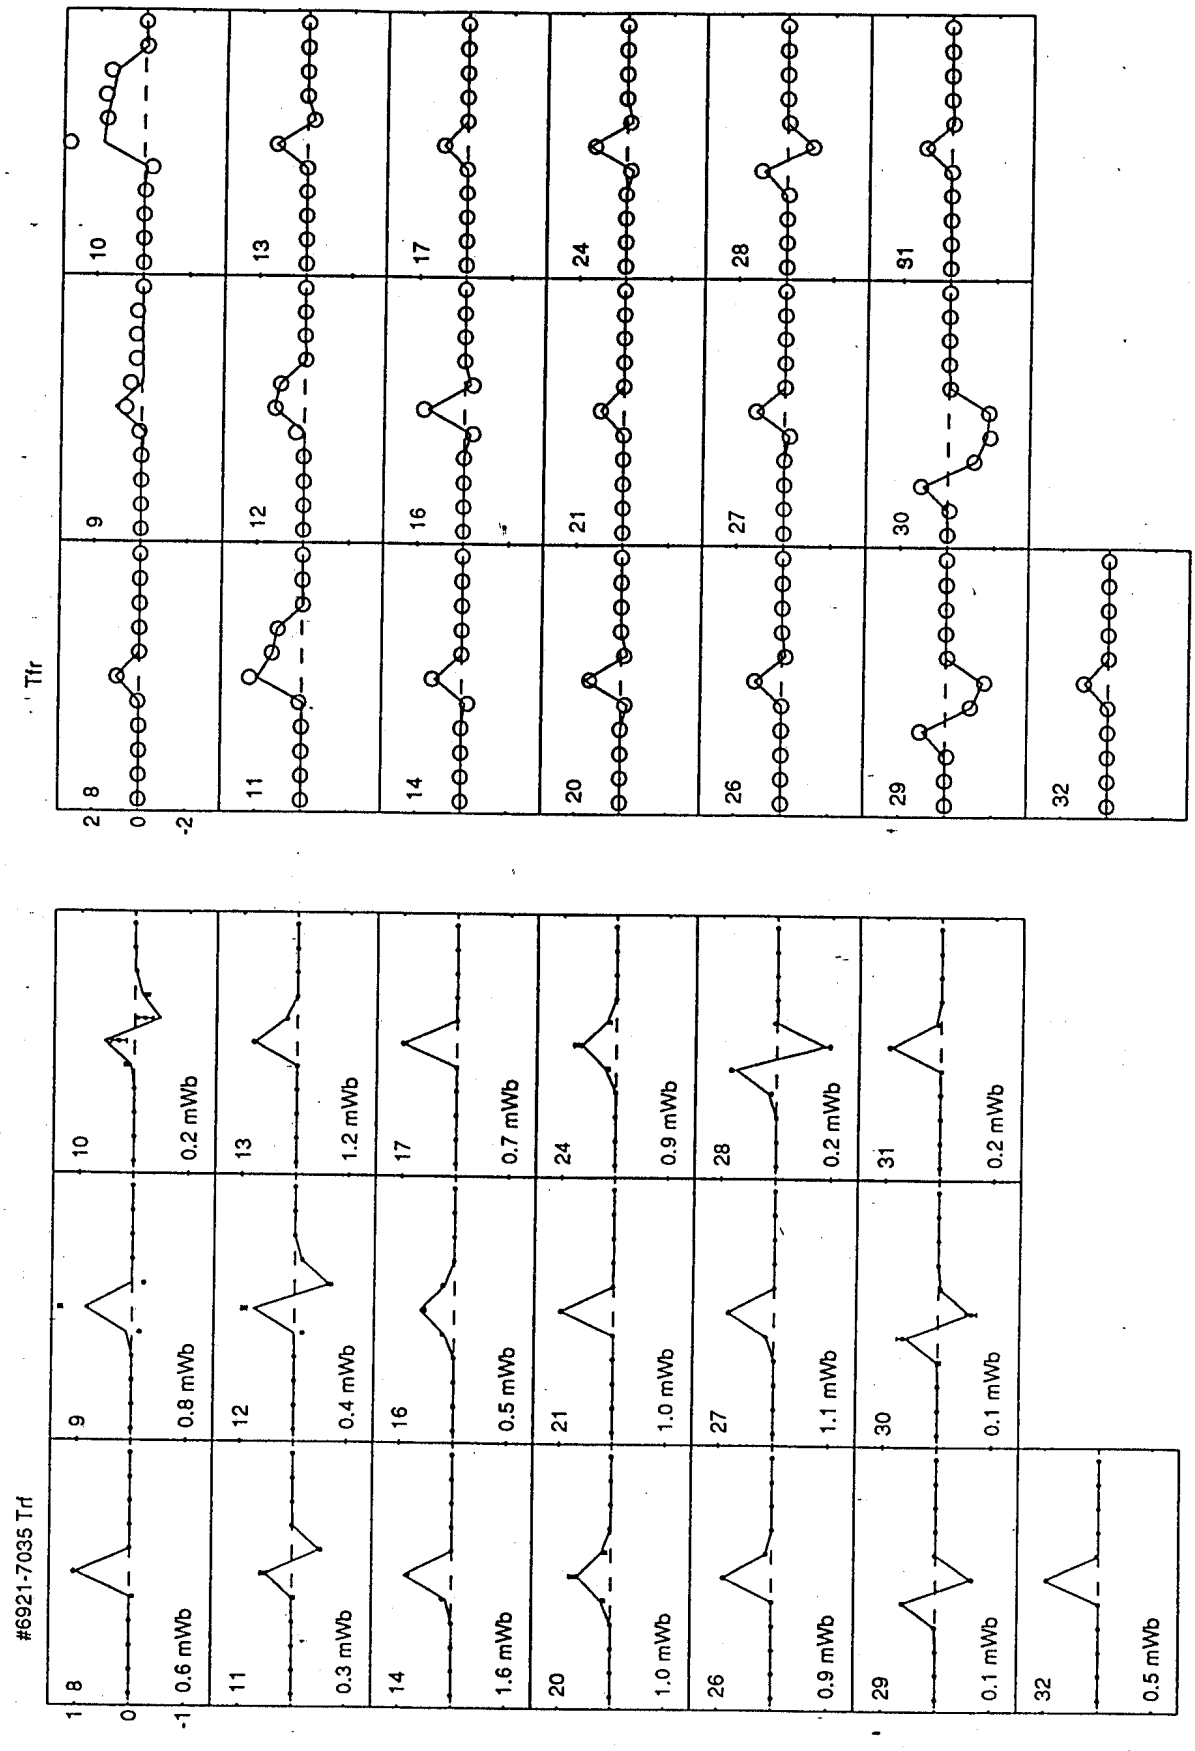

FIGURE 13 *Transformation from the relative flux loop measurements to the corresponding ideal fluxes. T_{fr} is a matrix containing the coefficients of this transformation. T_{rf} is that of the inverse transformation. For each partial flux loop, the coefficient for that loop is centrally plotted together with those of its neighbour. The coefficients derived from the constructional dimensions are plotted as solid lines while those derived from measurements are represented by dots. The axes are fixed and non-dimensional.*

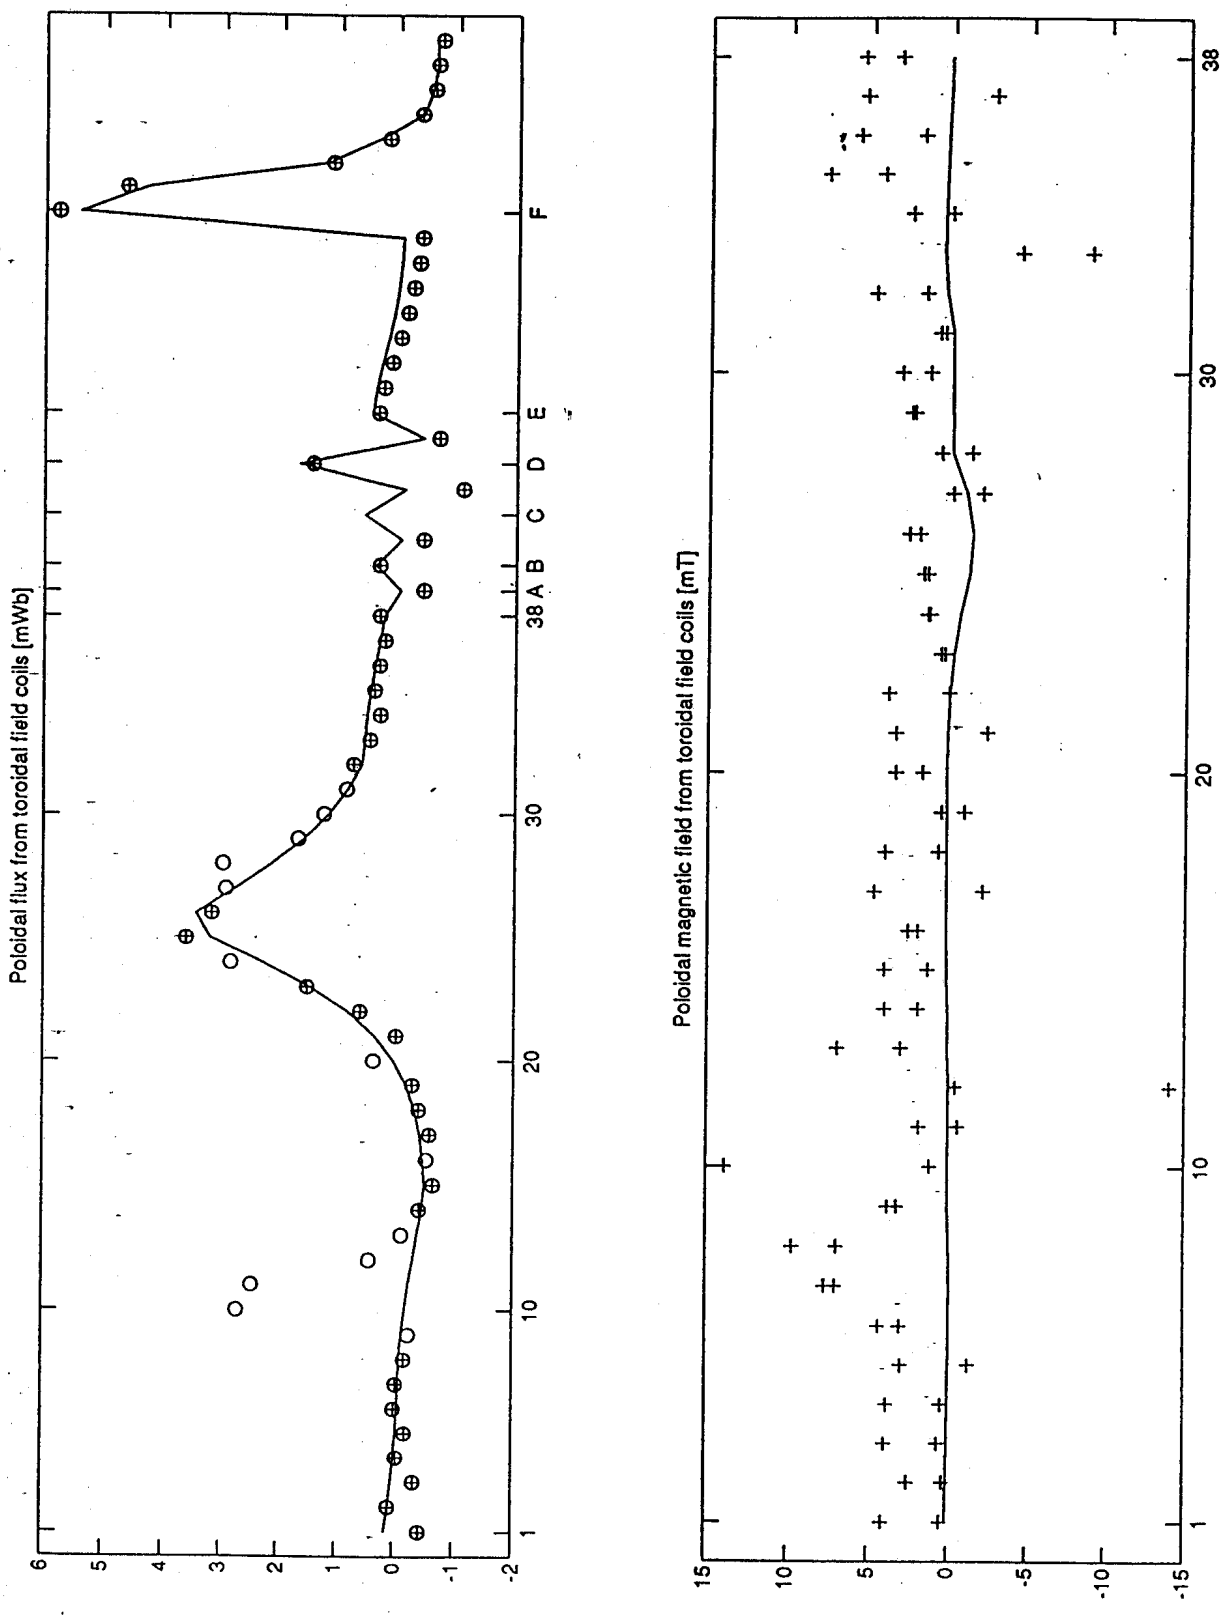

FIGURE 14 *Poloidal flux and poloidal magnetic field produced by the toroidal field coil connections at the location of the loops and probes. The solid line represents the values computed on the basis of the connection position while crosses (complete loops) and circles (détour loops) are the measured values. The main toroidal field was in this case 0.4 T.*

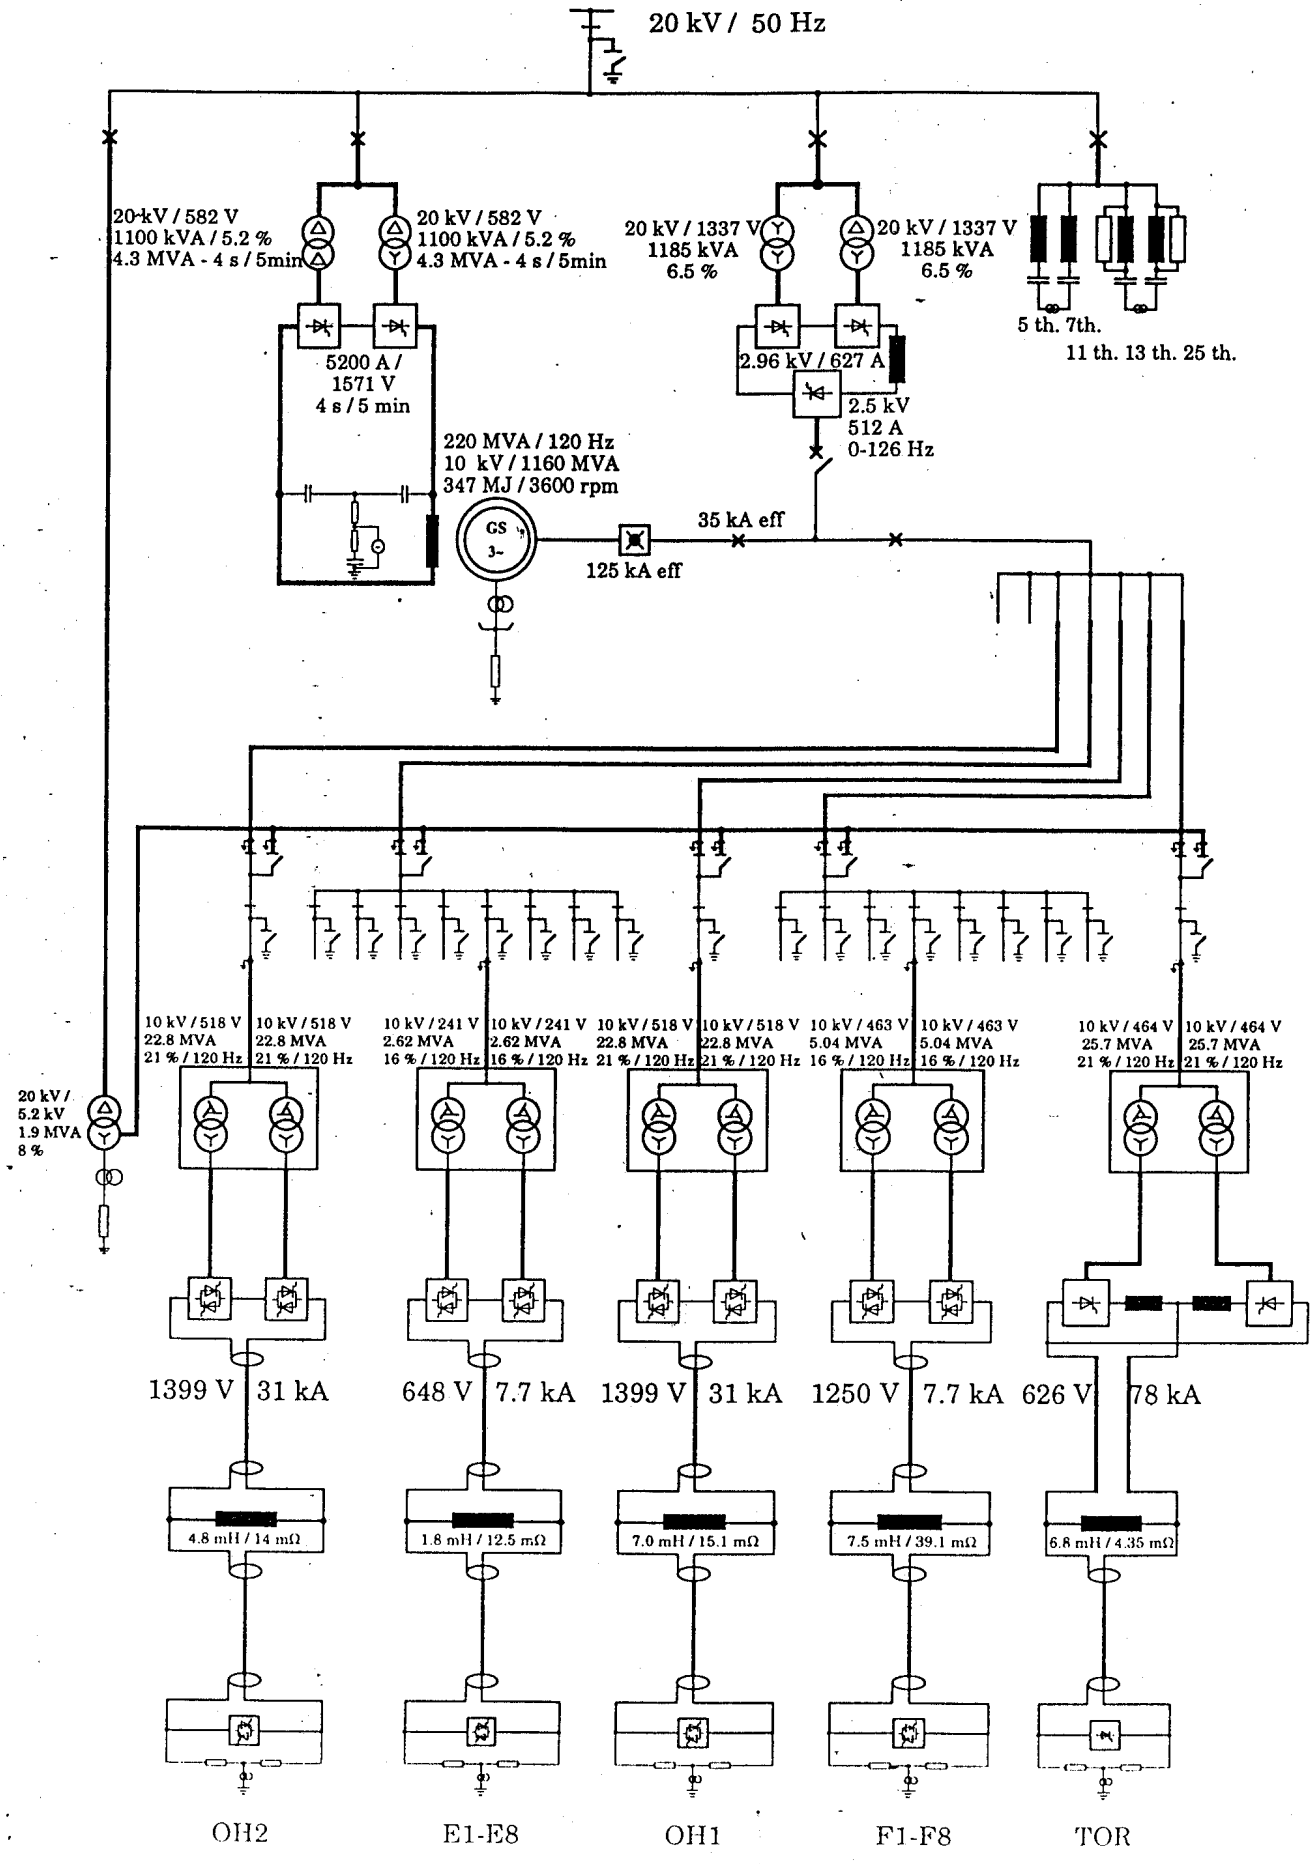

FIGURE 15 *Overall Outline of the TCV Power Systems. The input and output voltages, powers and short-circuit voltages are shown for each transformer.*

a) Current feedback

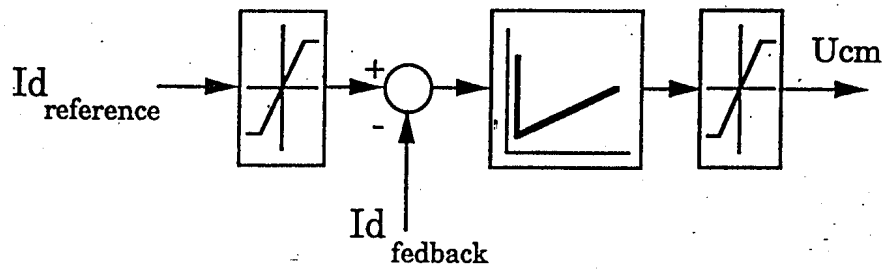

b) Open loop voltage driving (with current limitation)

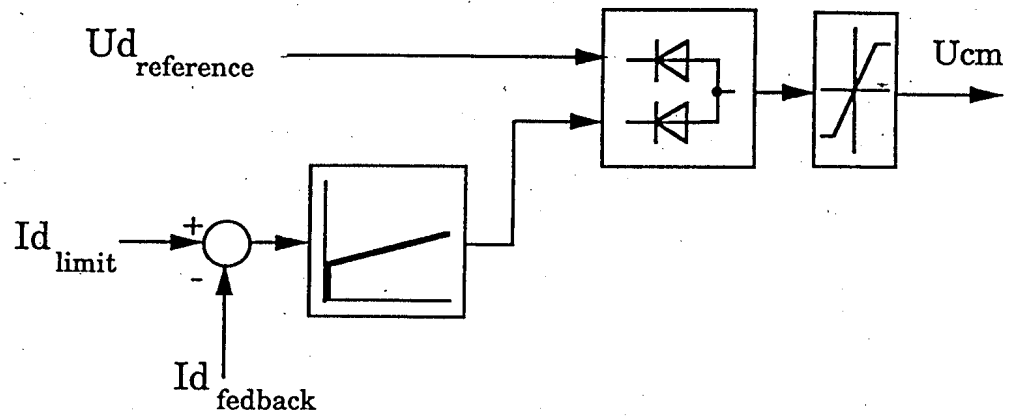

c) Hybrid driving (a and b)

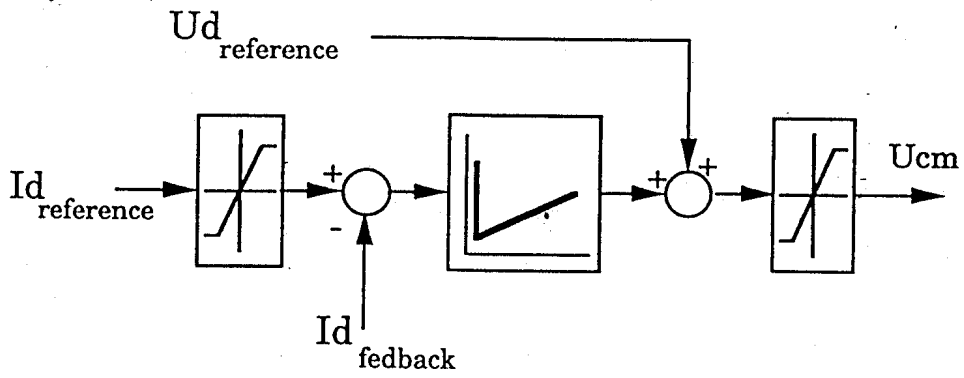

FIGURE 16 *Rectifier Control Modes.*

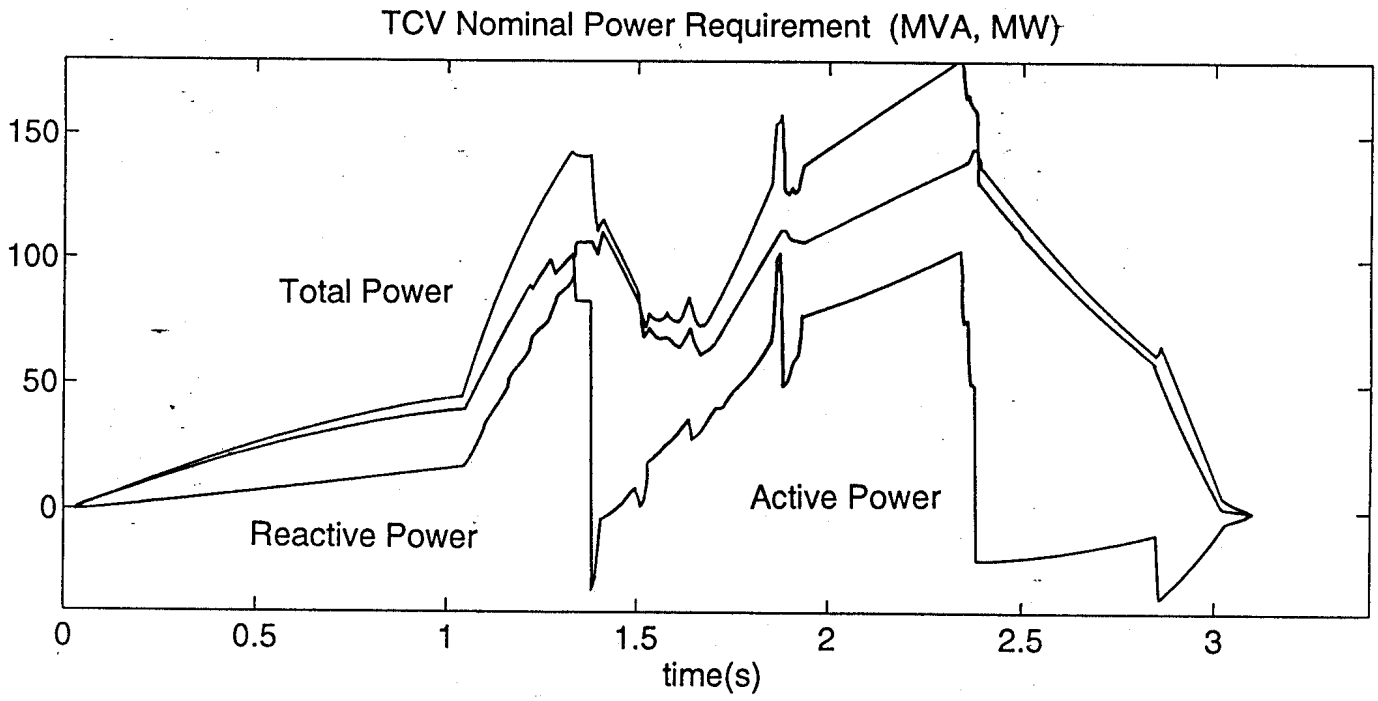

FIGURE 17 *Load characteristics of the nominal Tokamak Pulse. (Units are in MVA, MW aand MVAR),*

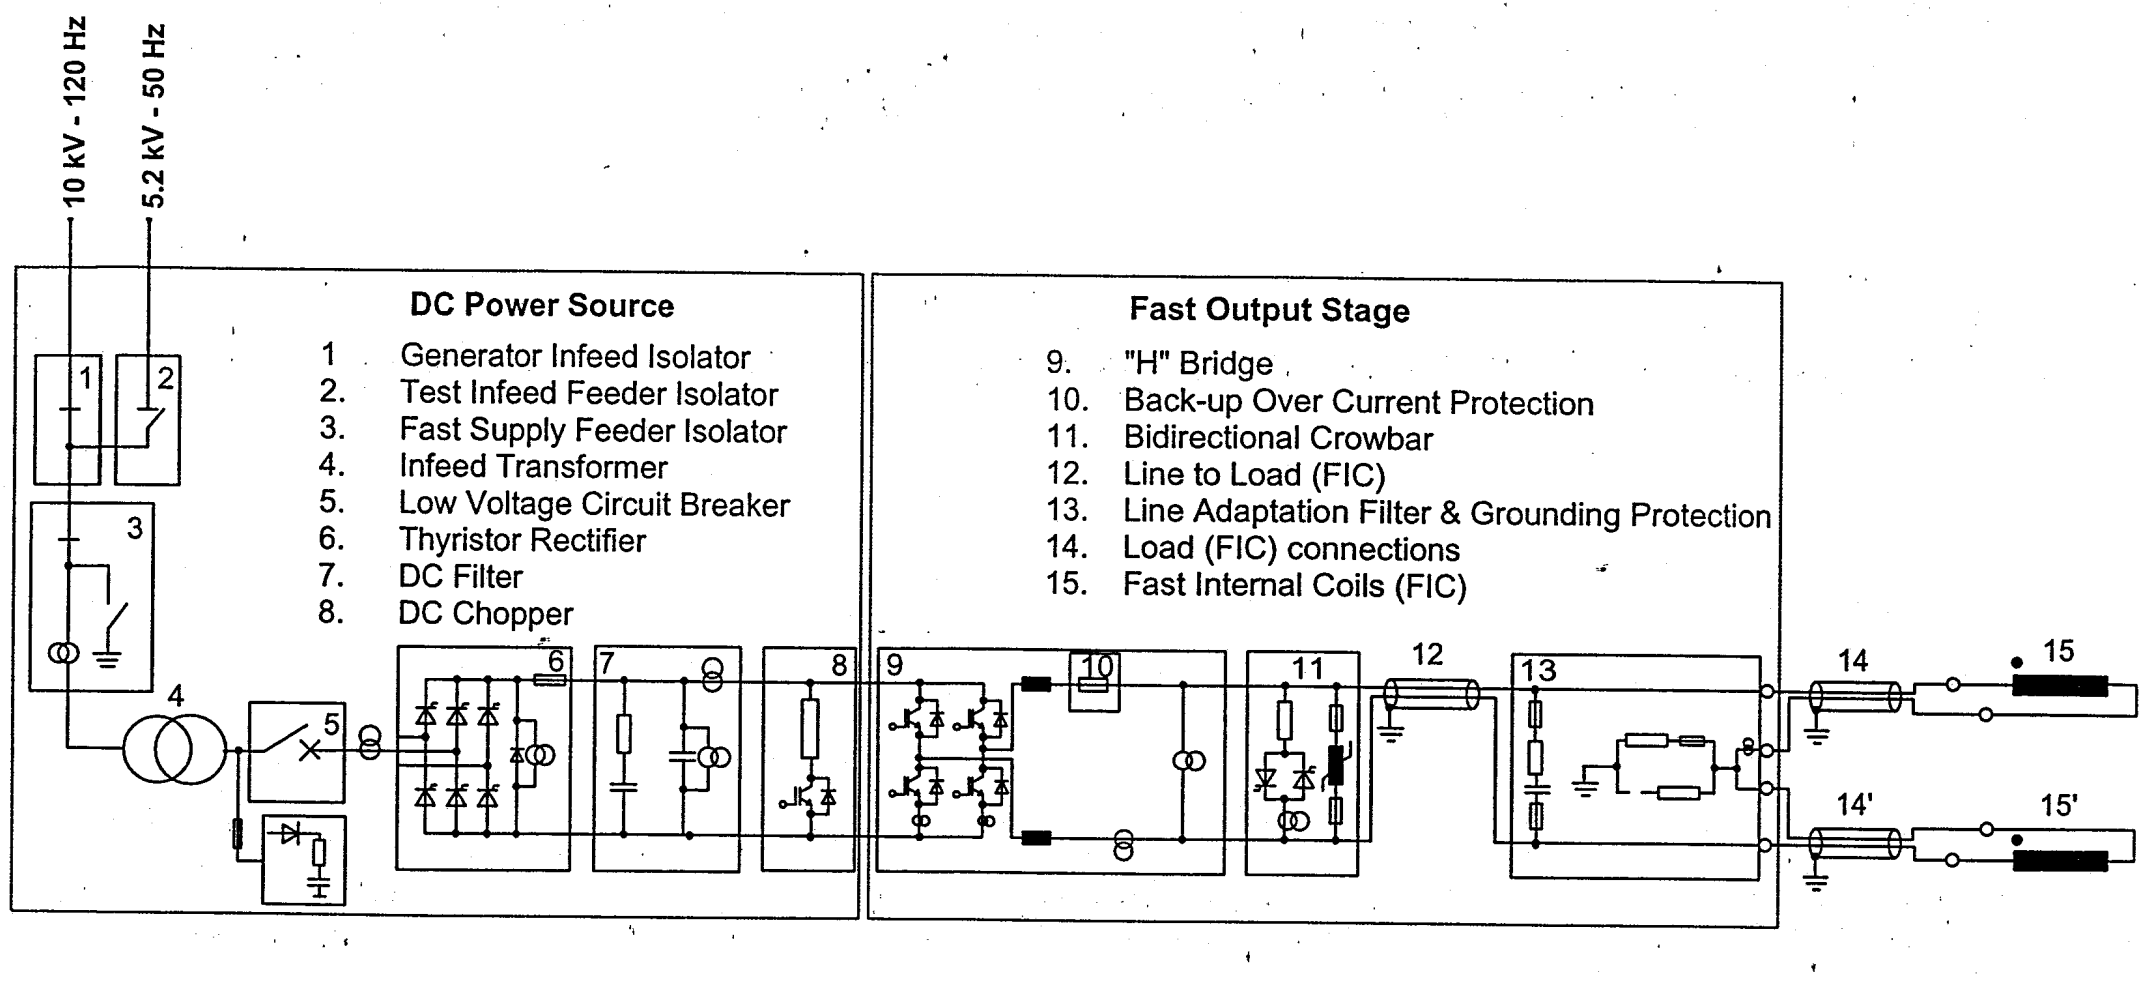

FIGURE 18 *Summary of the Fast Internal Coil Power Supply.*

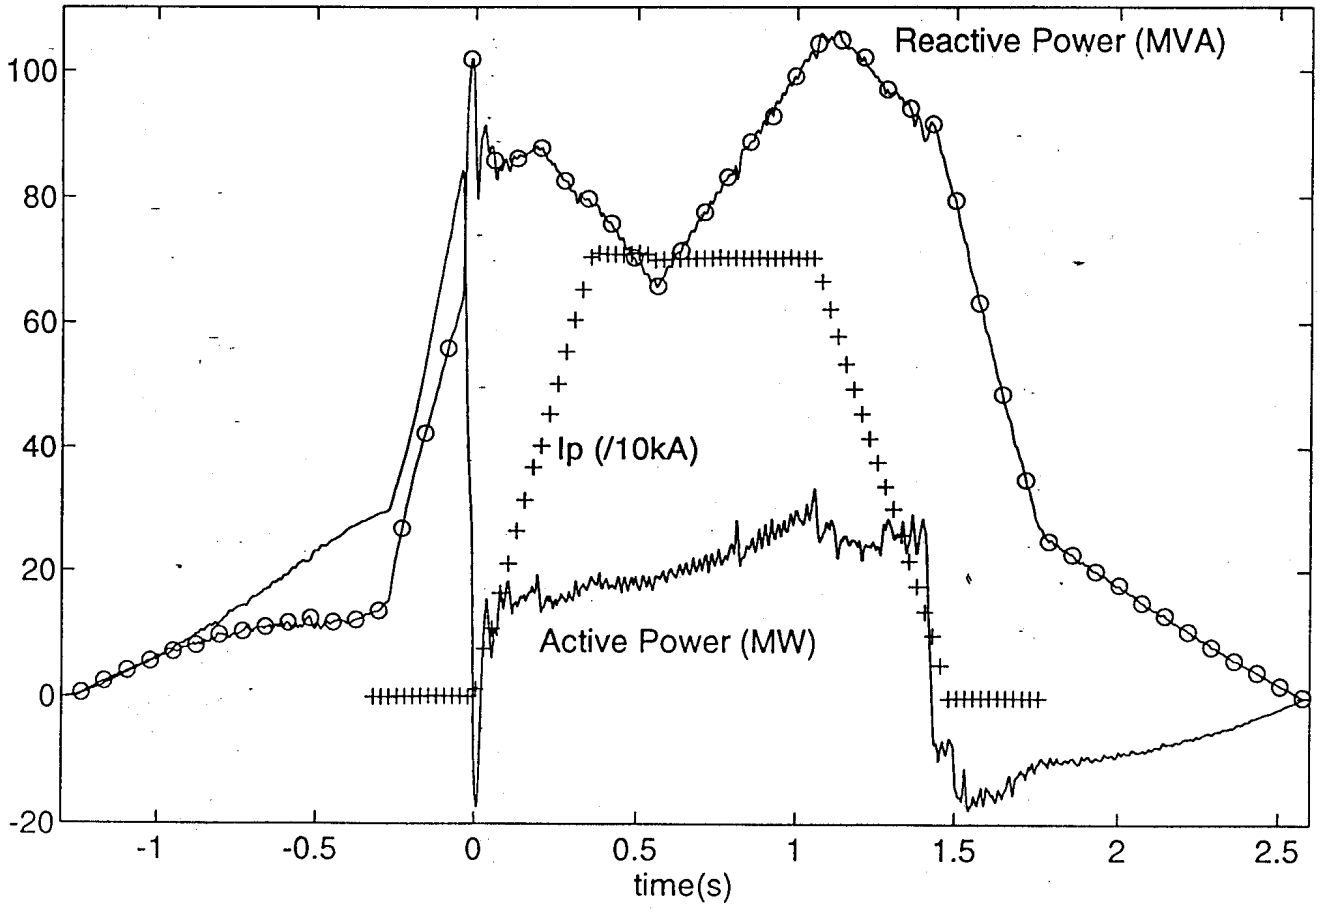

FIGURE 19 *Power requirements for a real TCV discharge. The 700 kA plasma current is shown as crosses, the reactive power as circles and the resistive power as a simple line. The scales are marked for each curve.*

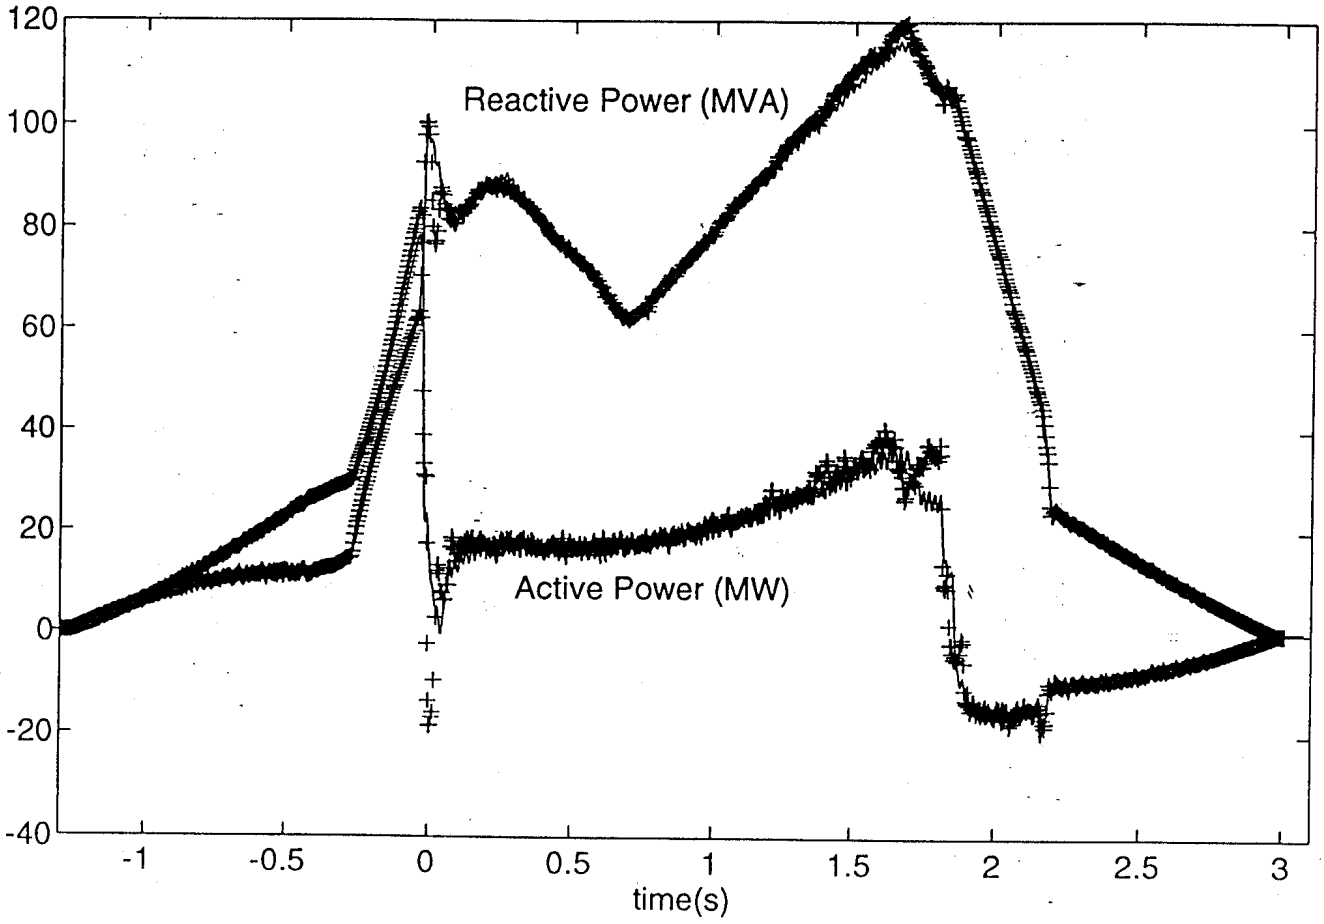

FIGURE 20 *Comparison between the required power for a pre-programmed discharge (solid line) and a real plasma discharge (crosses) with feedback requirements due to repeated equilibrium changes. The scales are marked for each curve.*

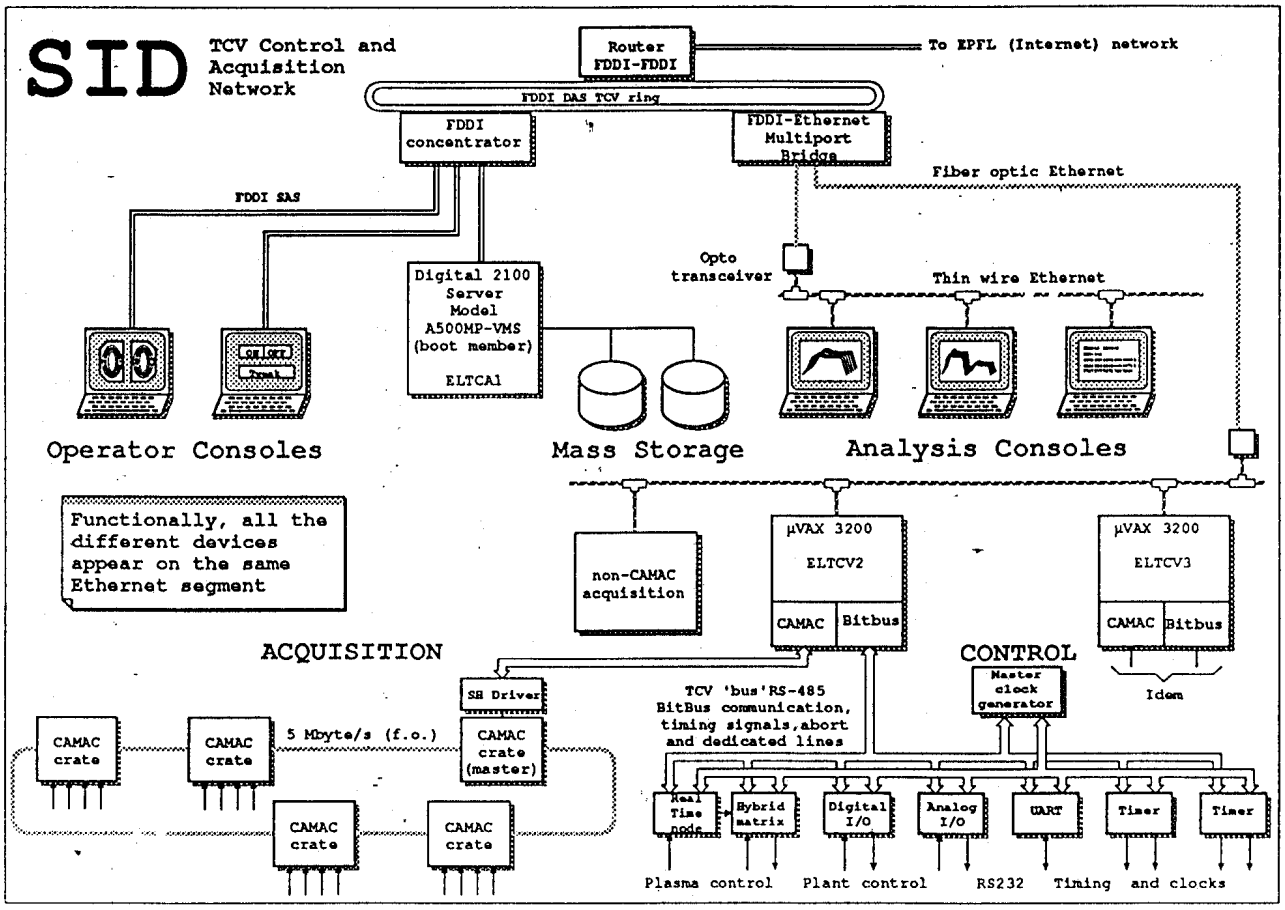

FIGURE 21 *TCV Système Intégré de Données. The integration of all of the computer hardware to control TCV,*

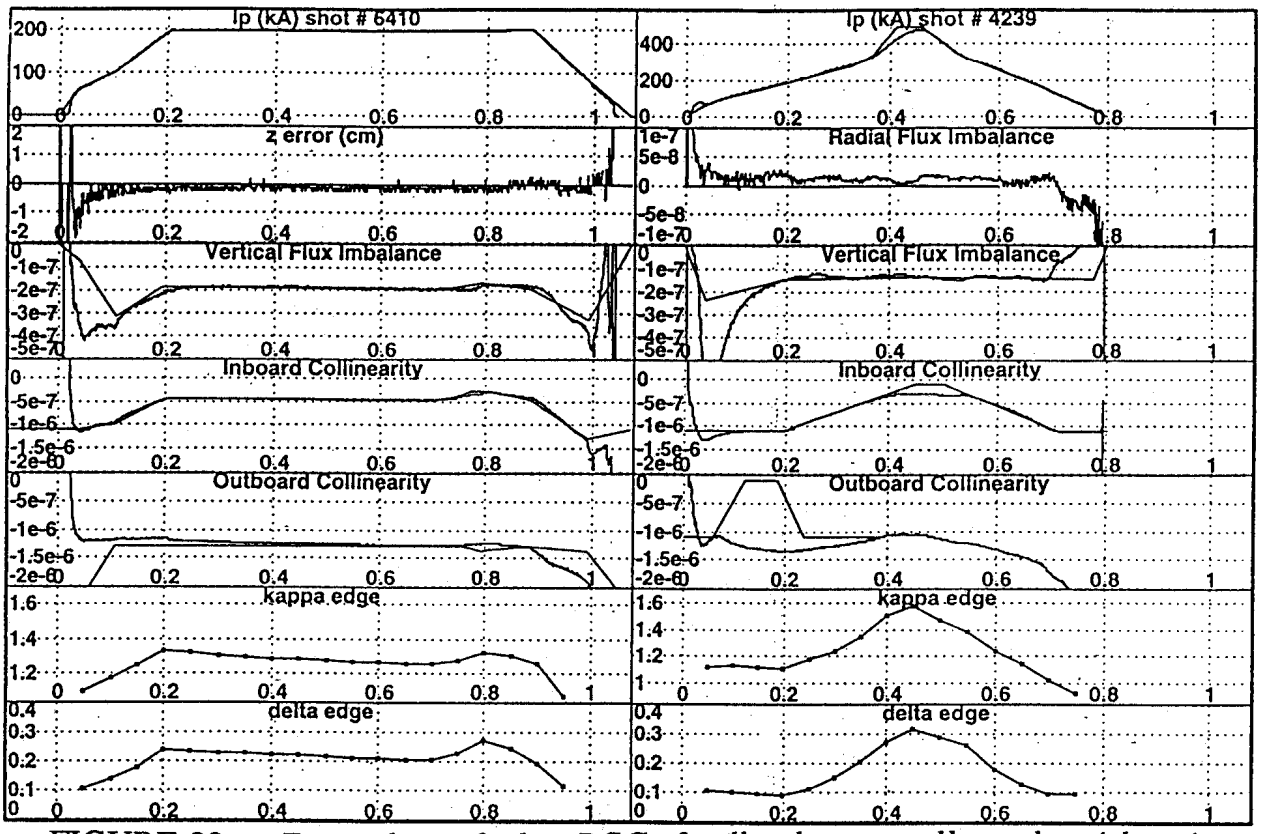

FIGURE 22 *Examples of the PCG feedback controller algorithm in annulling the feedback error. Left: Standard PCG feedback of a standard discharge. Right: No integral feedback gain in a discharge with small pre-programming errors.*

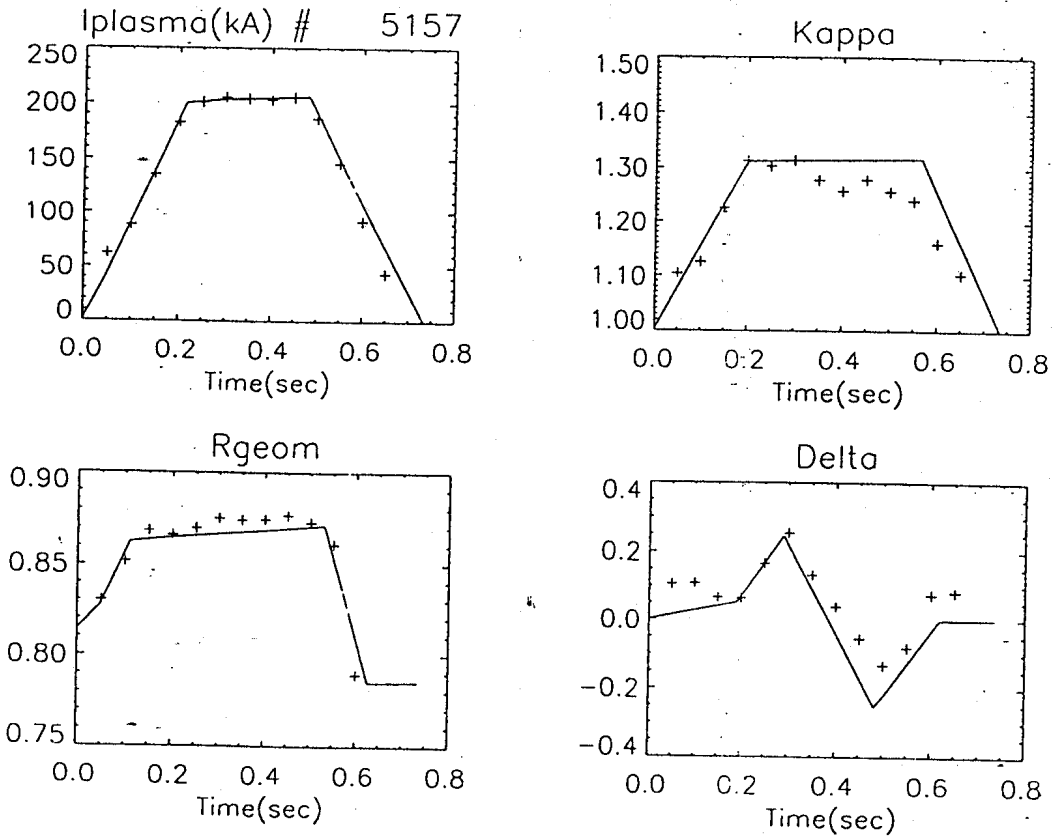

FIGURE 23 *Example of the pre-programmed physics waveforms and the reconstructed equilibrium parameters. The control variable reference waveforms are created by the Neural Network mapping. (Ip in kAmp, Rgeom in metres).*

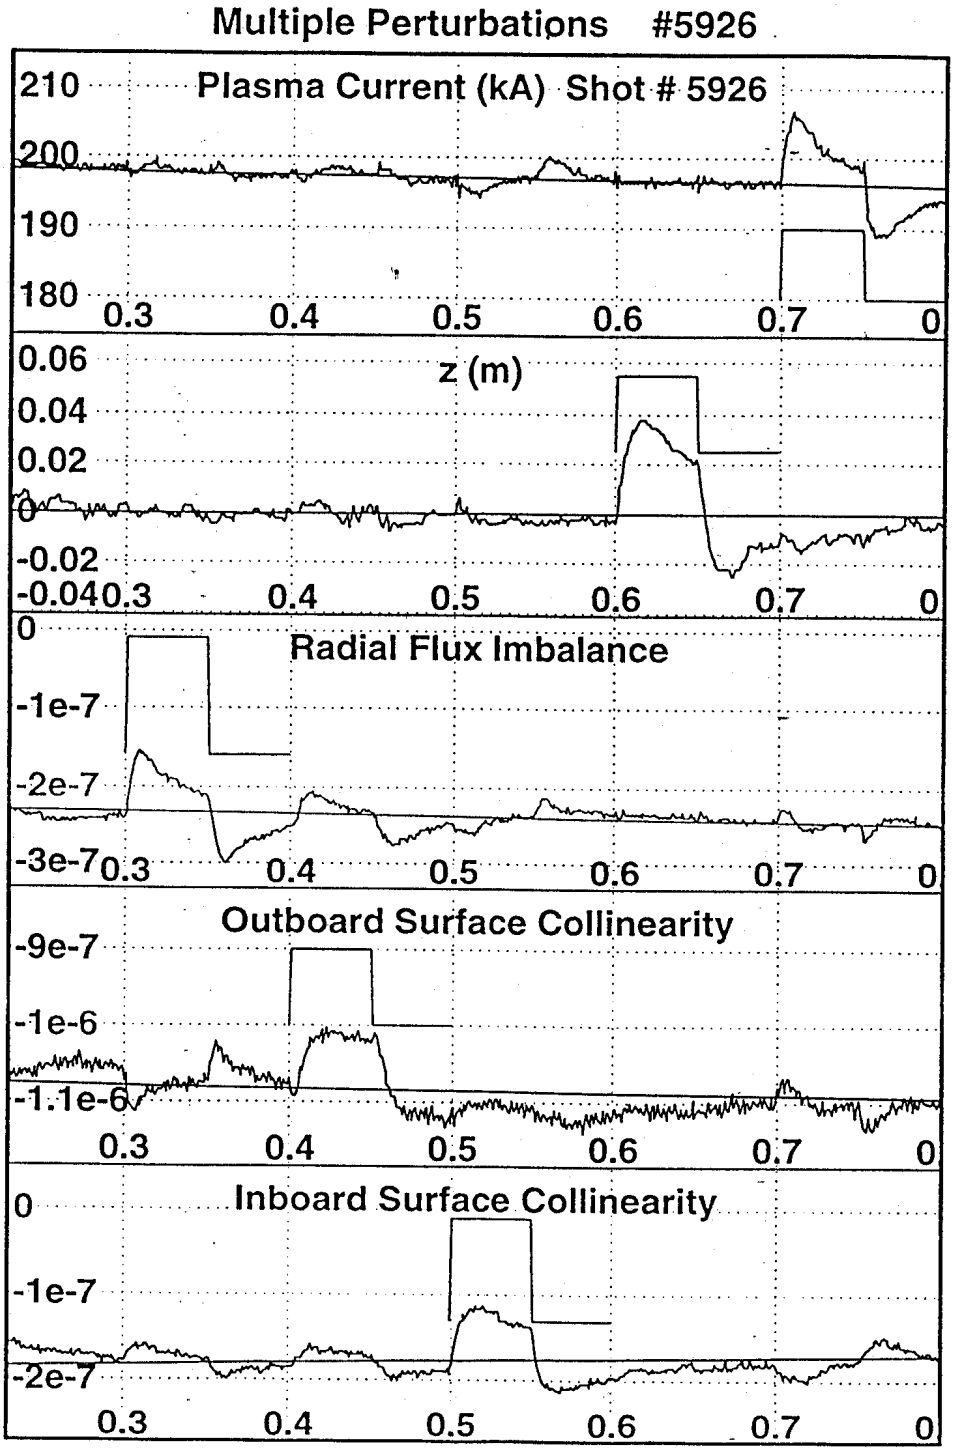

FIGURE 24 *Experimental measurement of the PCG controller decoupling.*

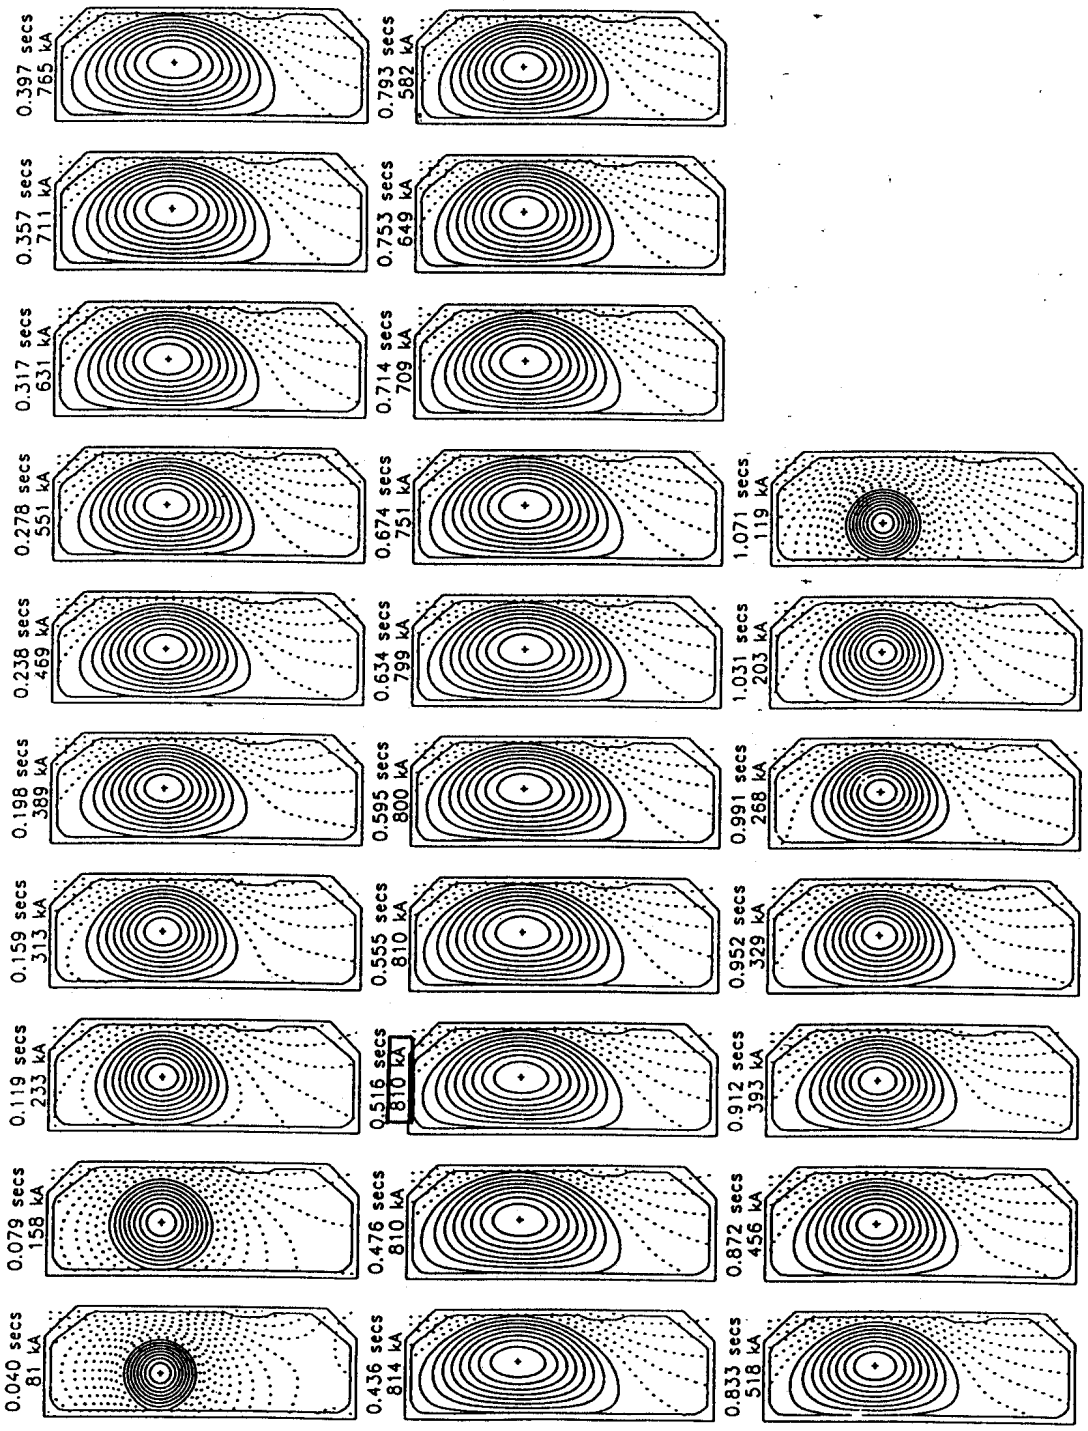

FIGURE 25 *The evolution of a TCV discharge, maximum plasma current 810 kA, maximum elongation 2.0.*

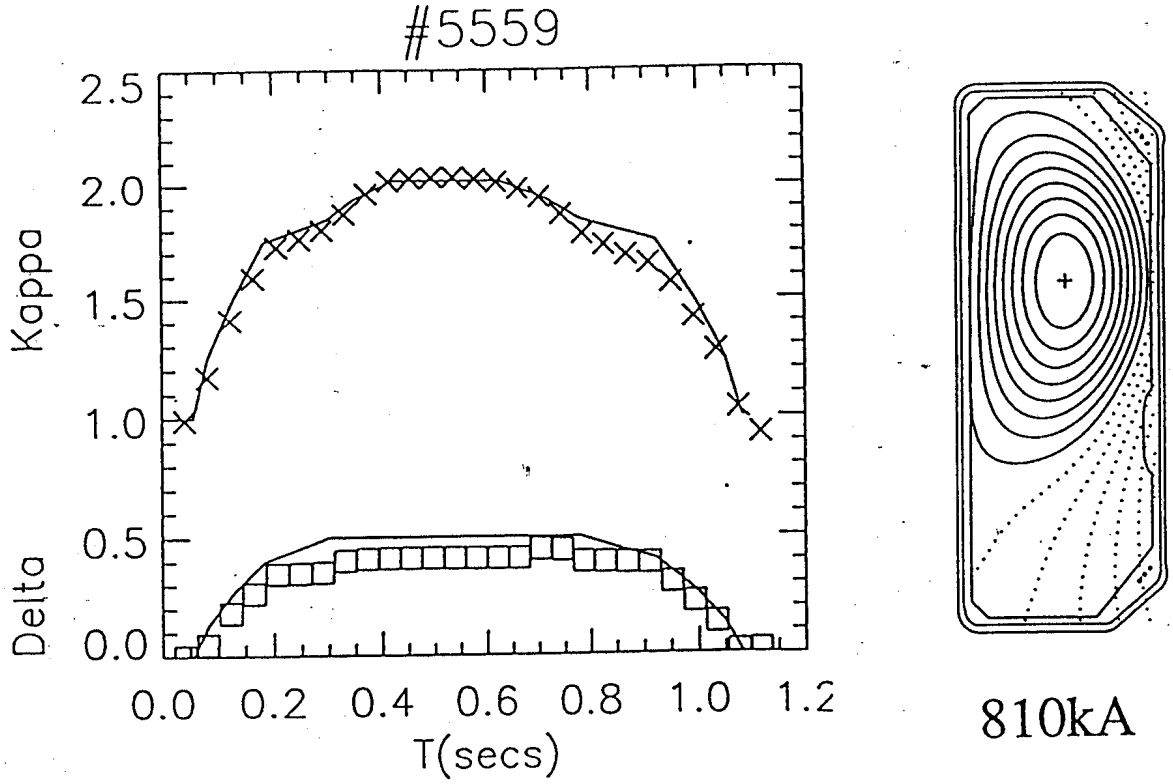

FIGURE 26 *Pre-programmed (lines) and measured (symbols) shape parameters for a limiter discharge with maximum plasma current 810 kA.*

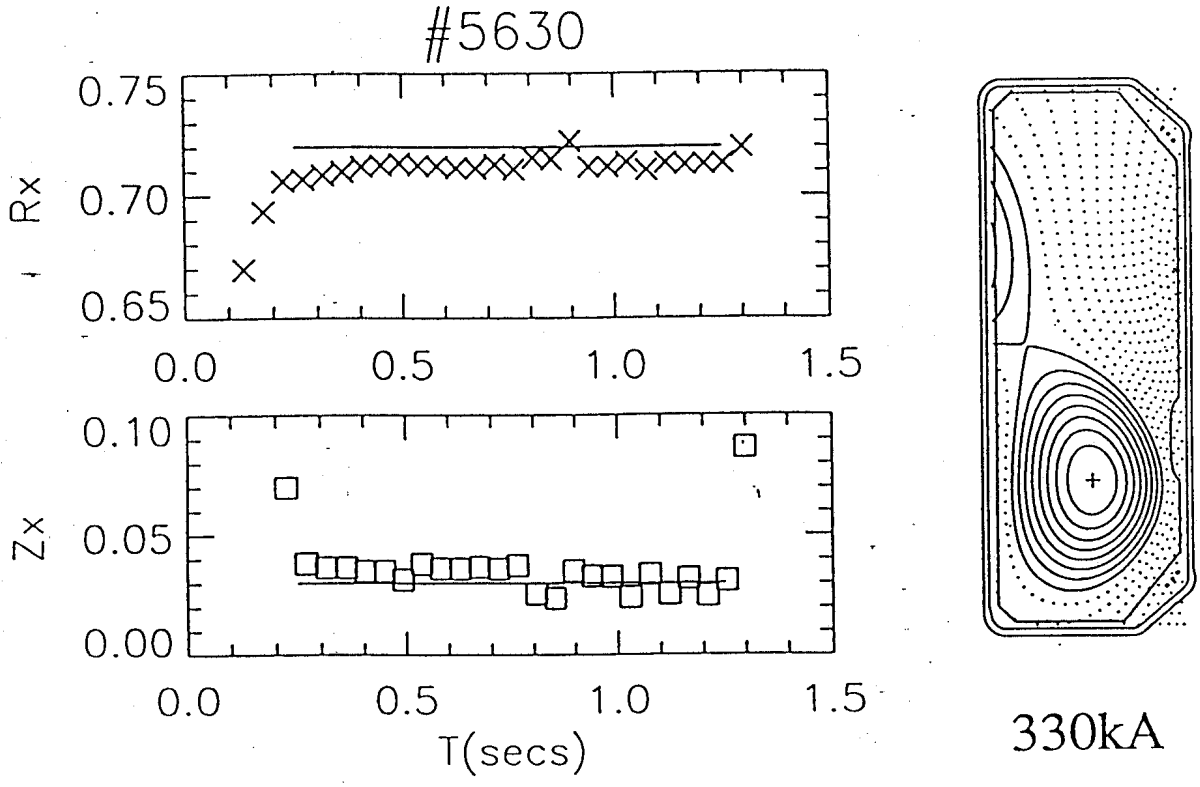

FIGURE 27 *Pre-programmed (lines) and measured (symbols) X-point coordinates for a single-null divertor discharge with a maximum plasma current of 330 kA.*

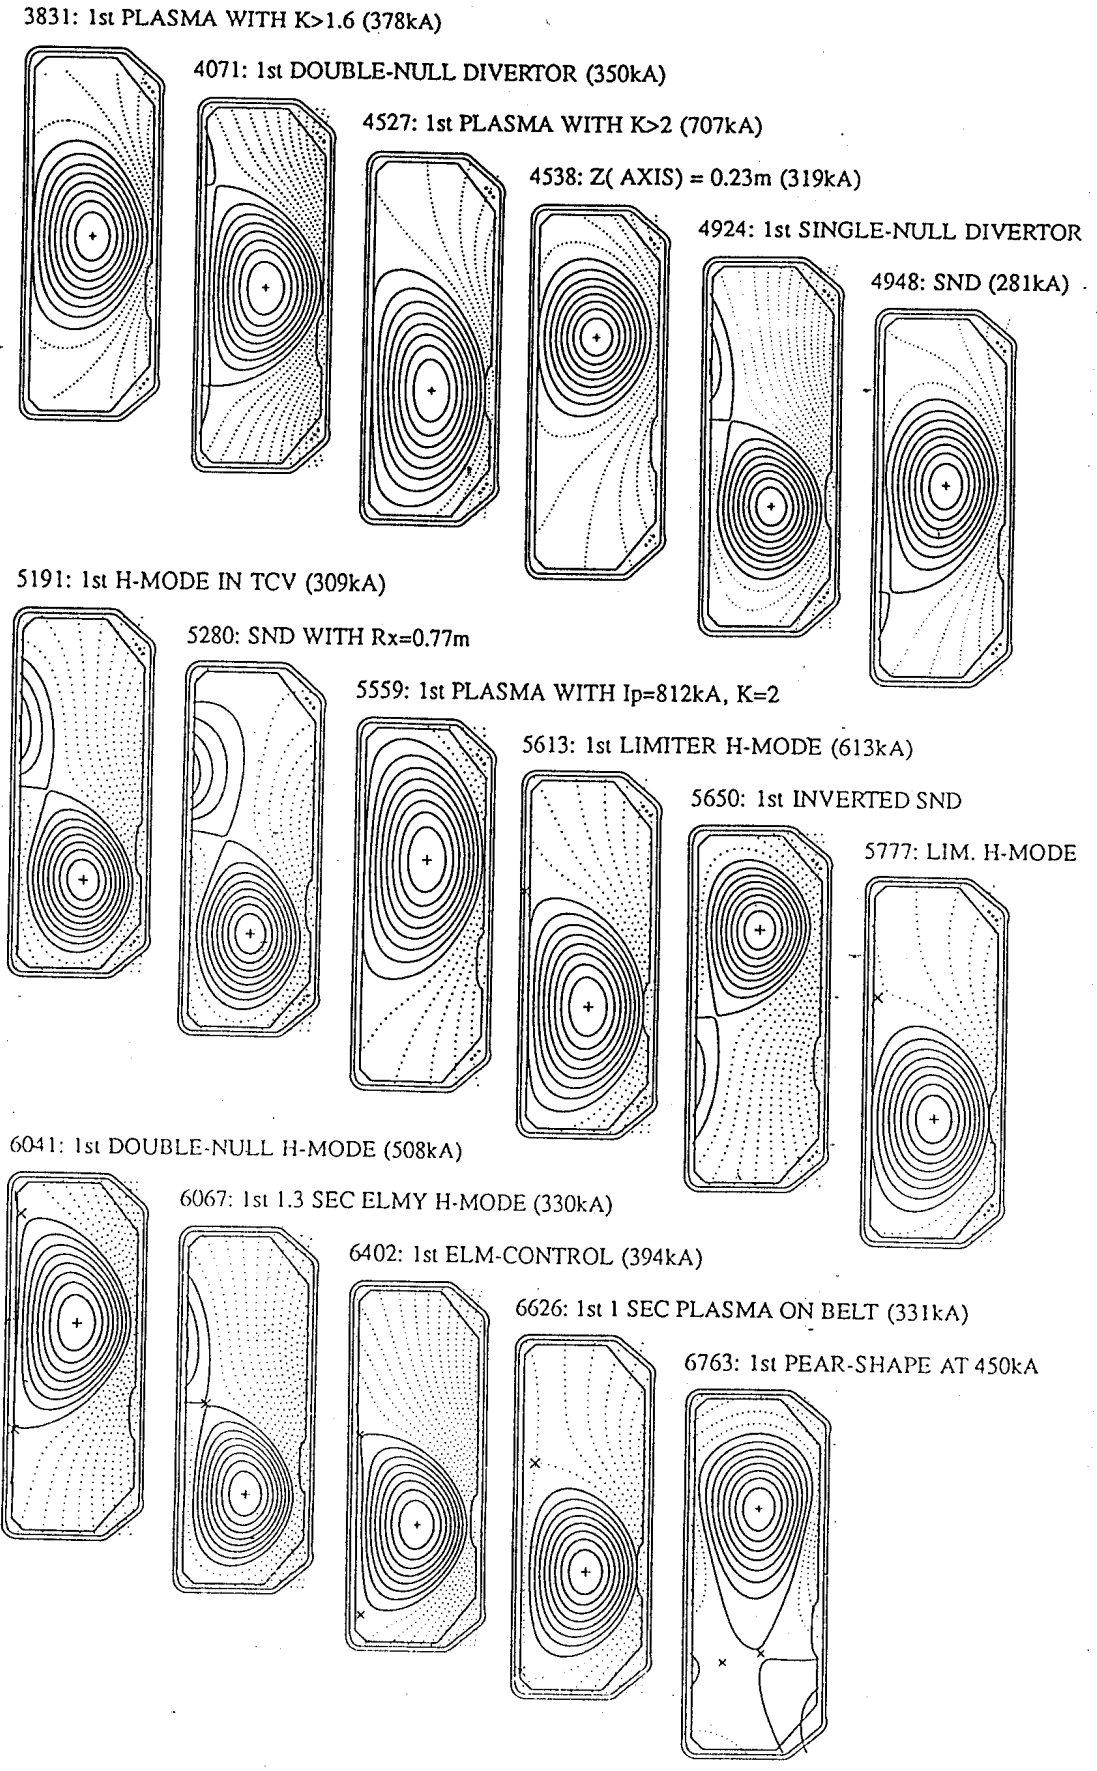

FIGURE 28 *Different equilibria generated by the MGAMS method.*

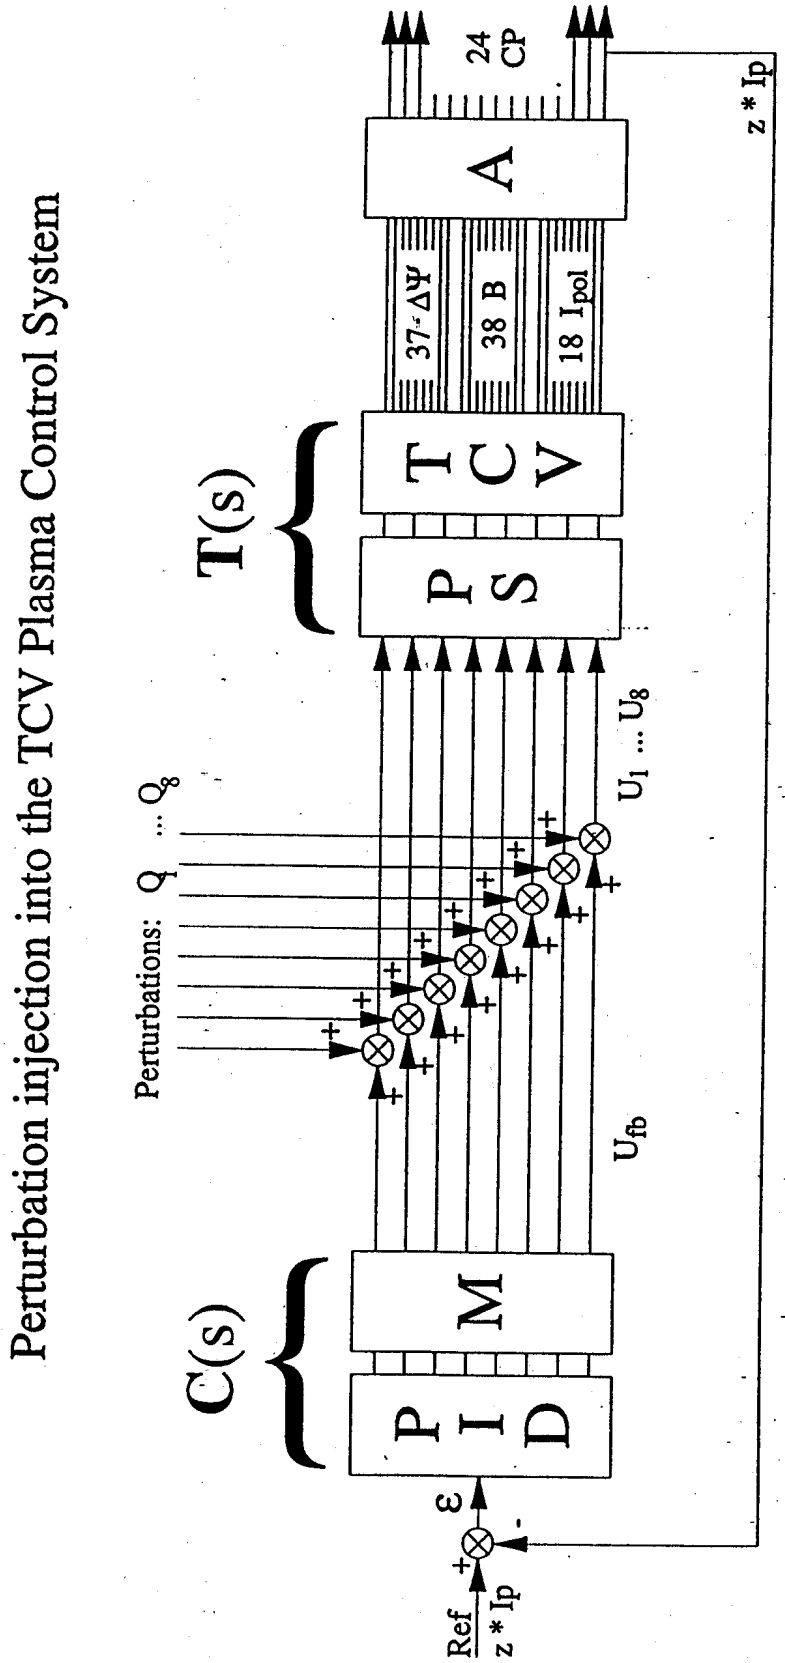

FIGURE 29 *Simple figure of the vertical position feedback control loop, showing the 8 anti-symmetric voltage demand signals (U) and perturbation inputs (Q). The transfer functions of the SIMO feedback controller, C(s), and that of the MISO comprising power supplies (PS) plus plasma-vessel-coil system, T(s), are also indicated.*

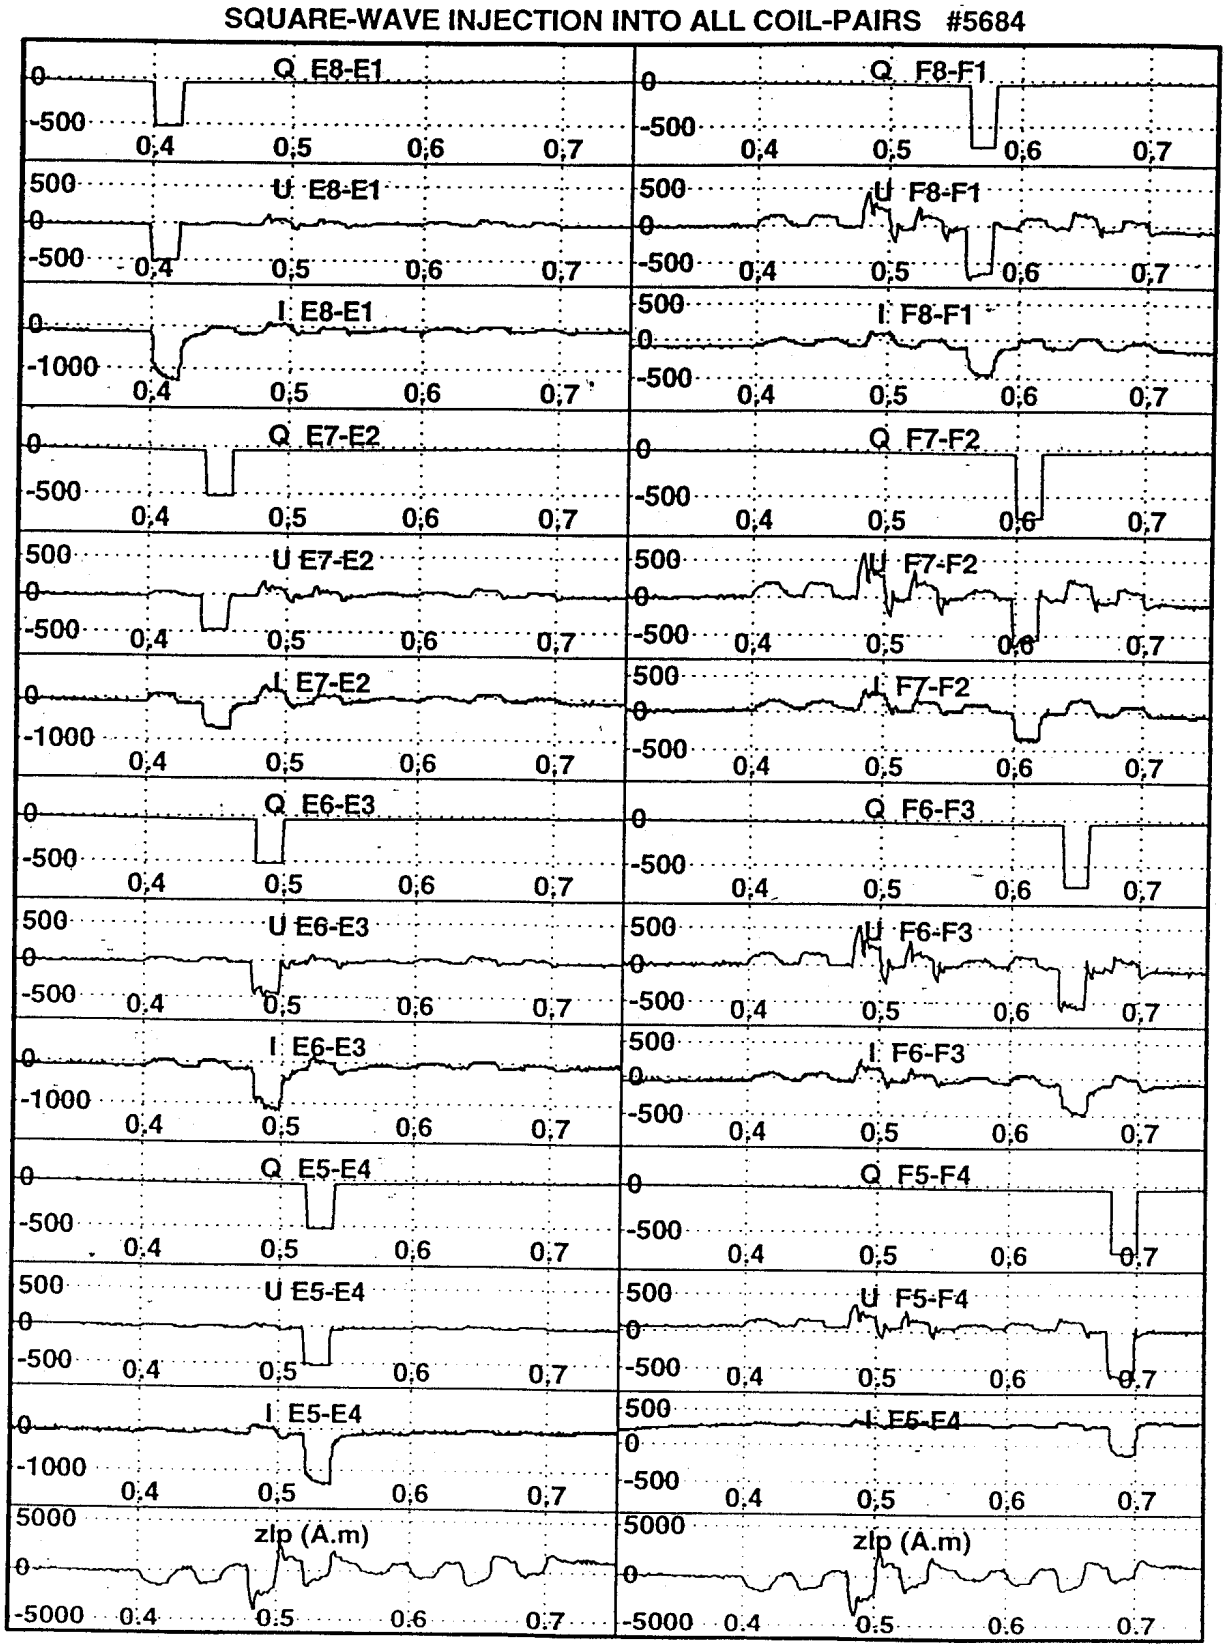

FIGURE 30. *Response of the system to square-wave pulse injection. For each anti-symmetric coil-pair we show the perturbation (Q), the voltage demand (U) and the coil-pair current (I). The vertical position zI_p is also shown.*

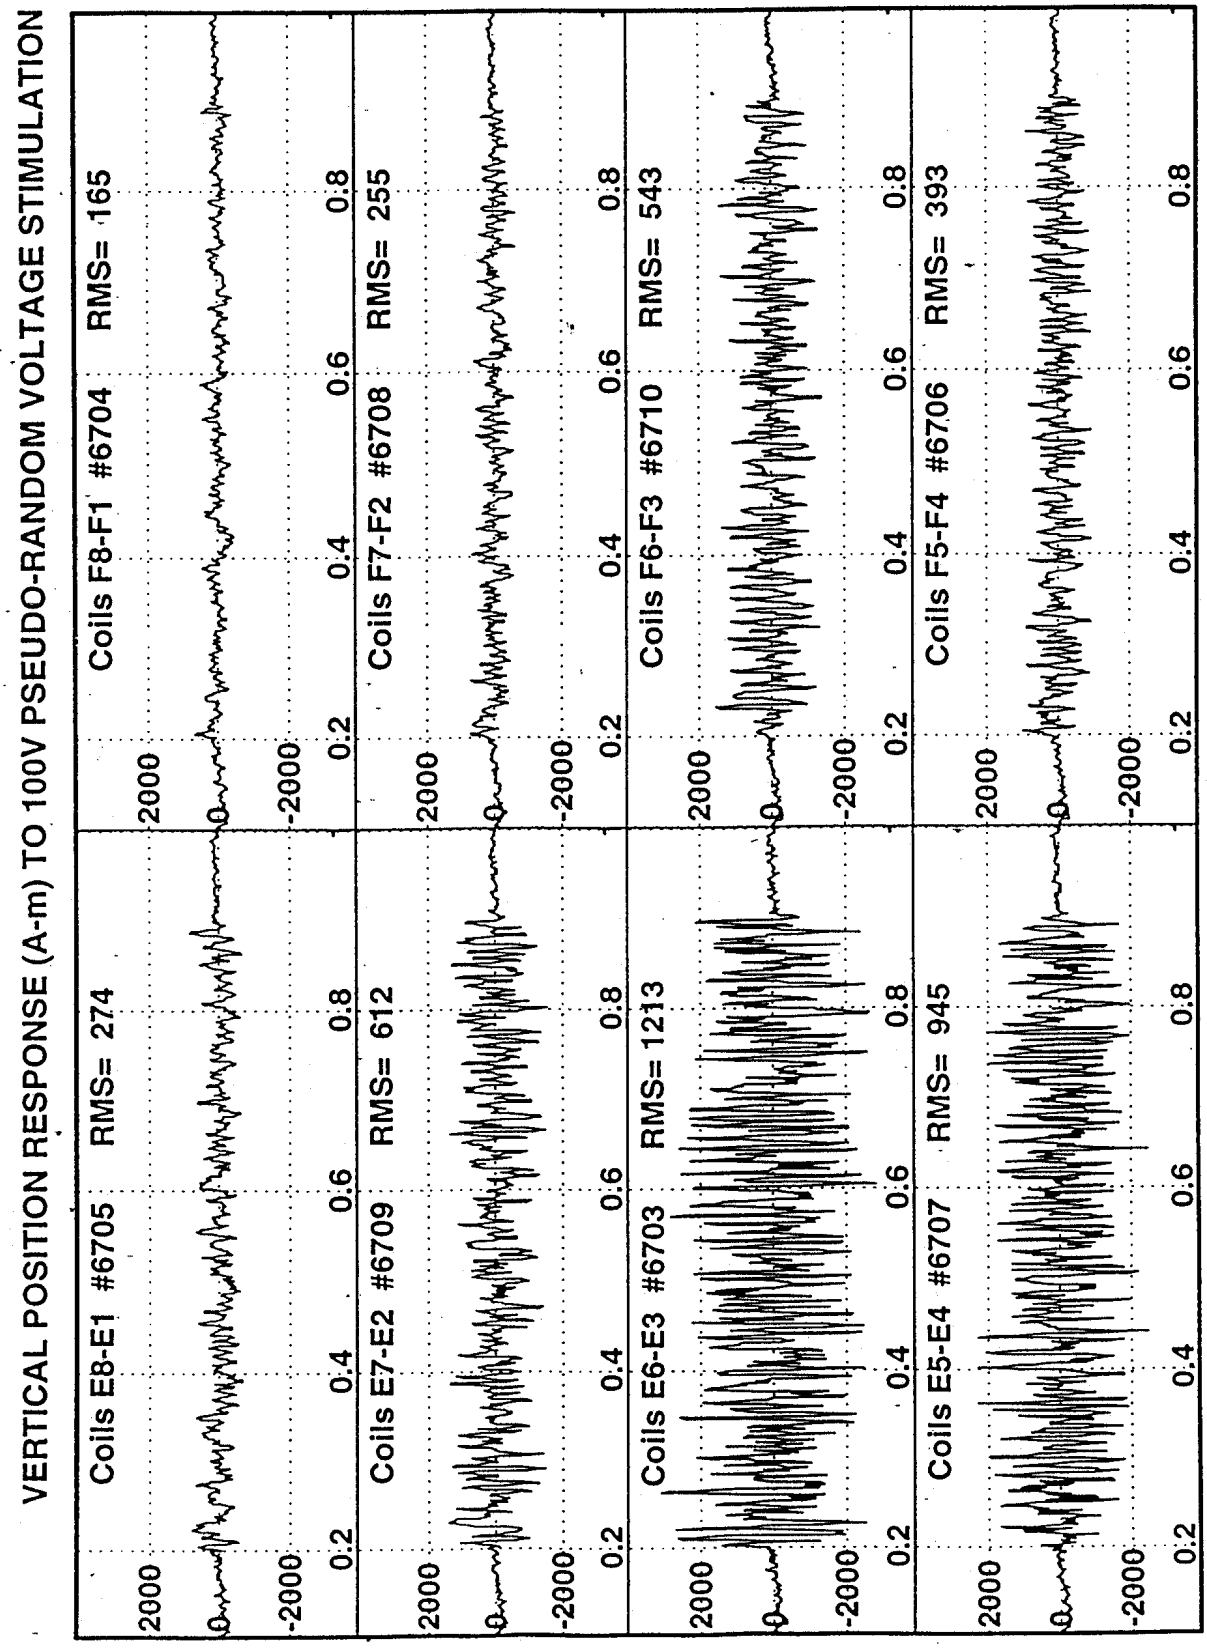

FIGURE 31 *Plasma position response to 700 msec. of continuous anti-symmetric excitation by a random binary sequence for each of the 8 TCV coil-pairs. The variation of the current-weighted vertical plasma position response is indicated for each coil-pair (RMS).*

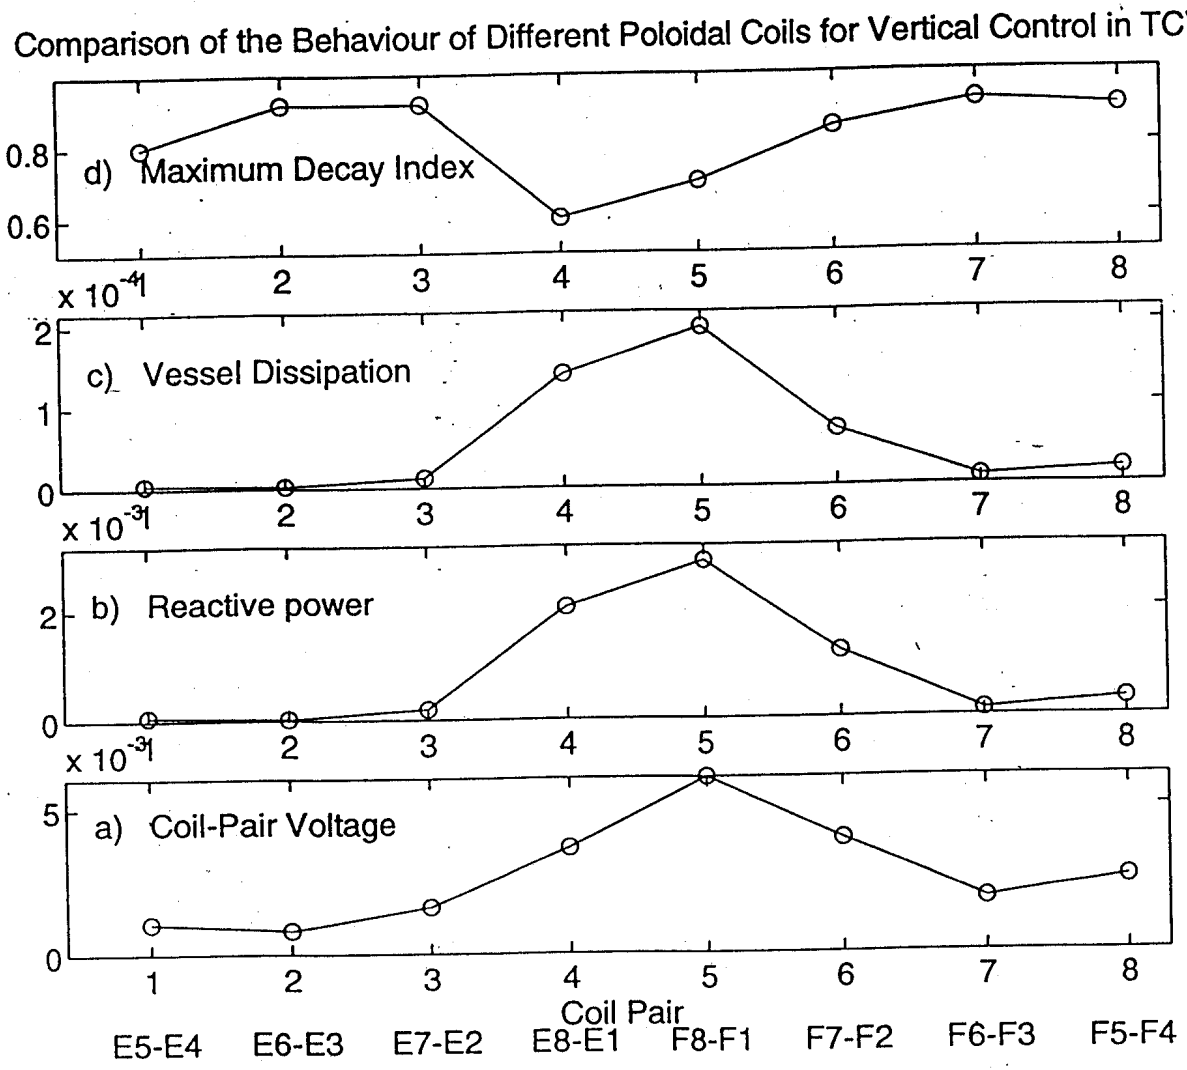

FIGURE 32 *The response of TCV to random continuous anti-symmetric perturbation of different coil-pairs. The coils are ordered according to their poloidal position (Fig. 2) a) voltage required to produce a given standard deviation of the vertical position response; b) Reactive power for unit modulation; c) Total ohmic dissipation in the vessel; d) Maximum field curvature for which the RCDM predicts vertical stabilisation using only that coil-pair, expressed as* α/α_{crit}.

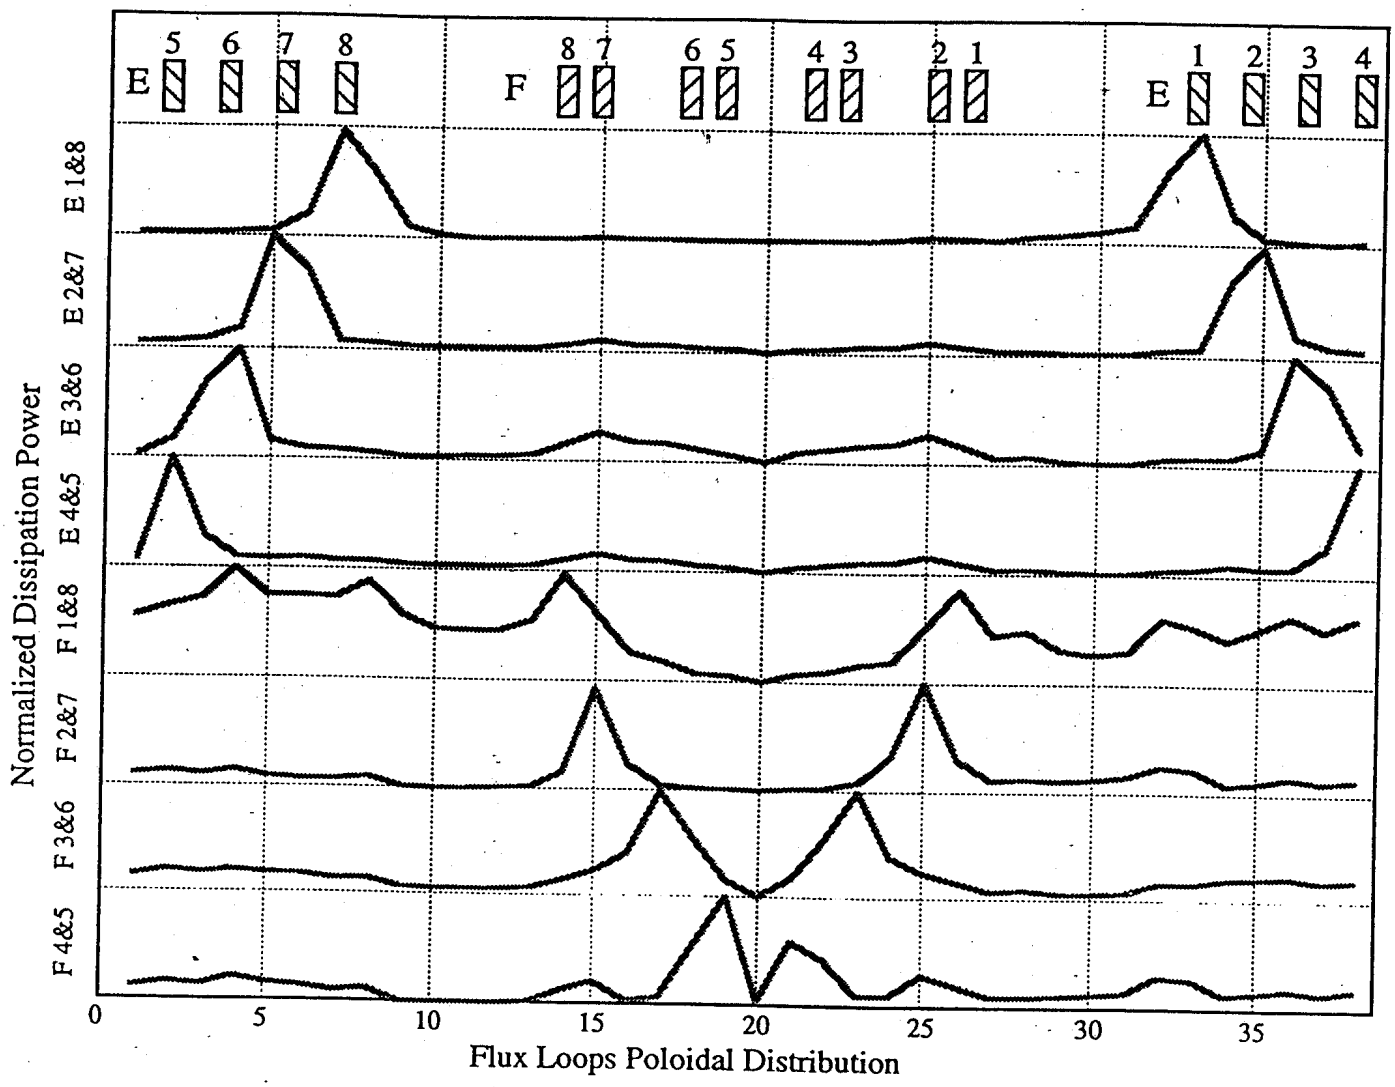

FIGURE 33 *Measured poloidal distribution of the induced vessel currents during the square-wave injection.*

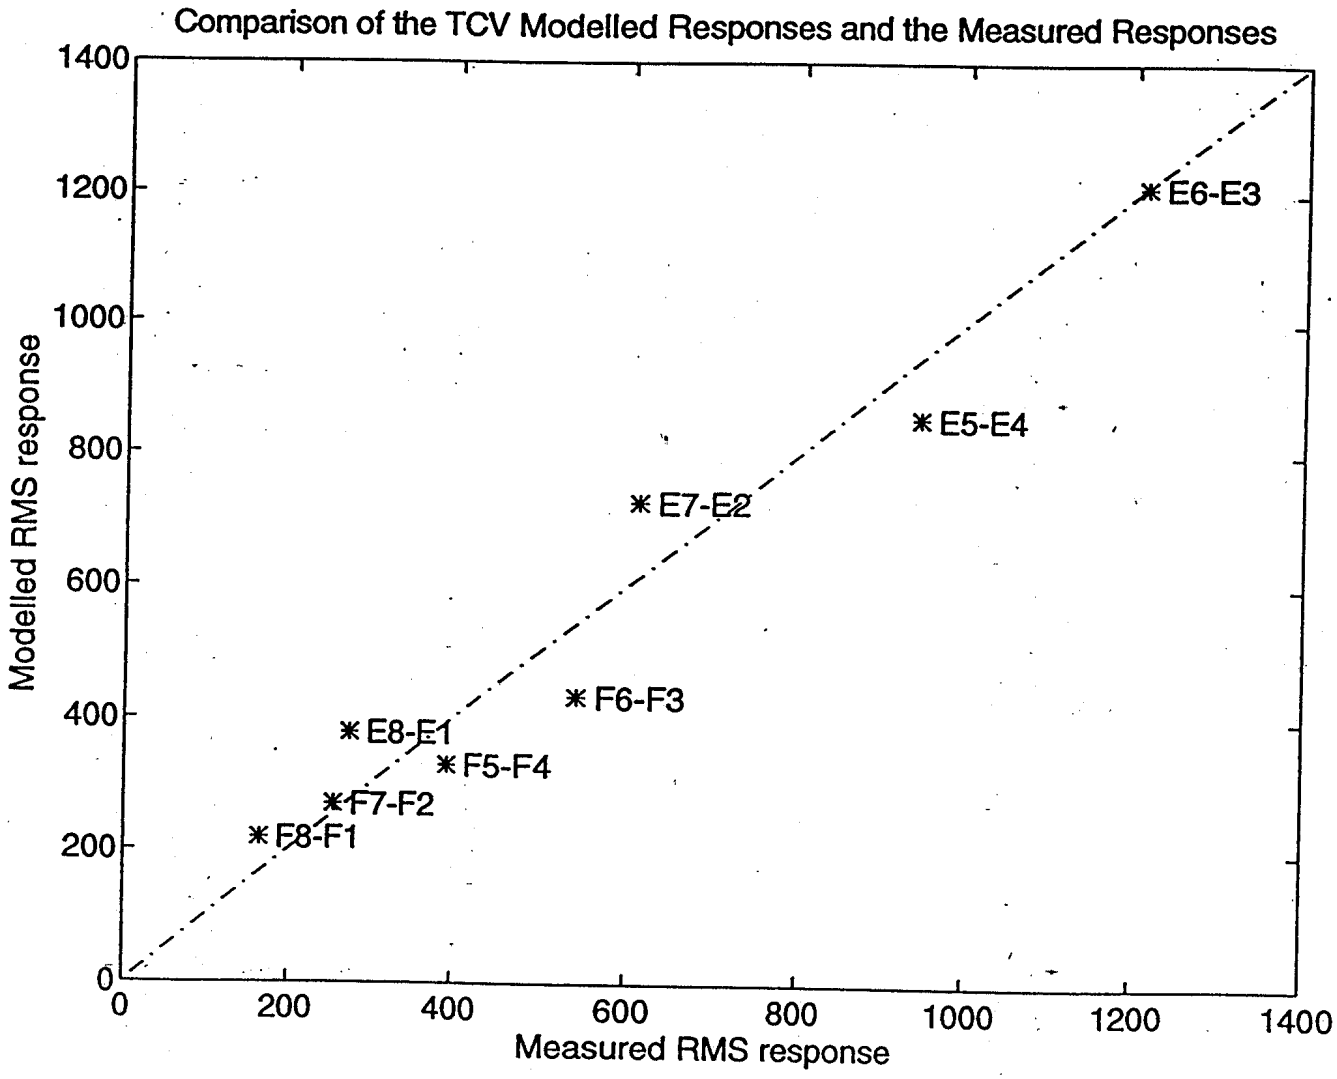

FIGURE 34 *Comparison between the zIp response found experimentally and that calculated using the RCDM.*

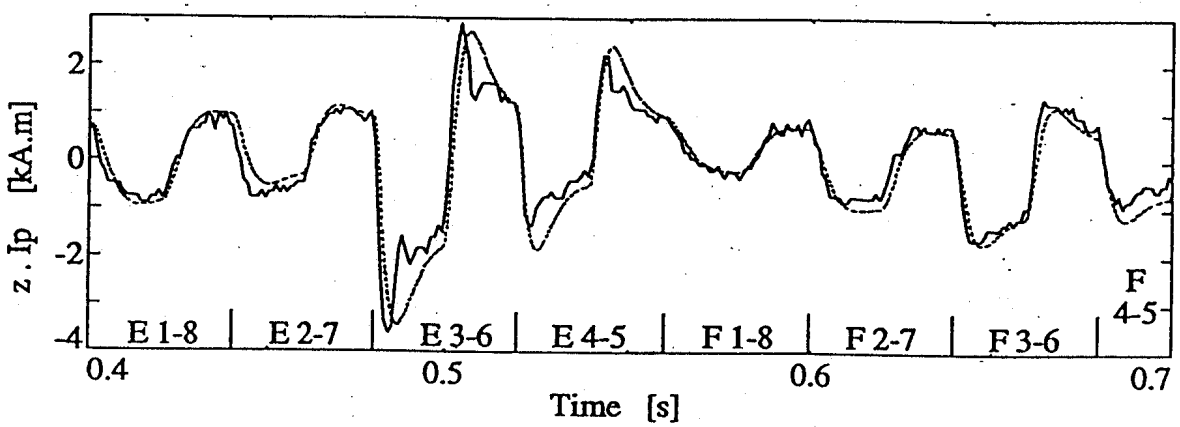

FIGURE 35 *Comparison between the $\sigma(zI_p)$ measured in the experiment and that predicted by the RCDM.*

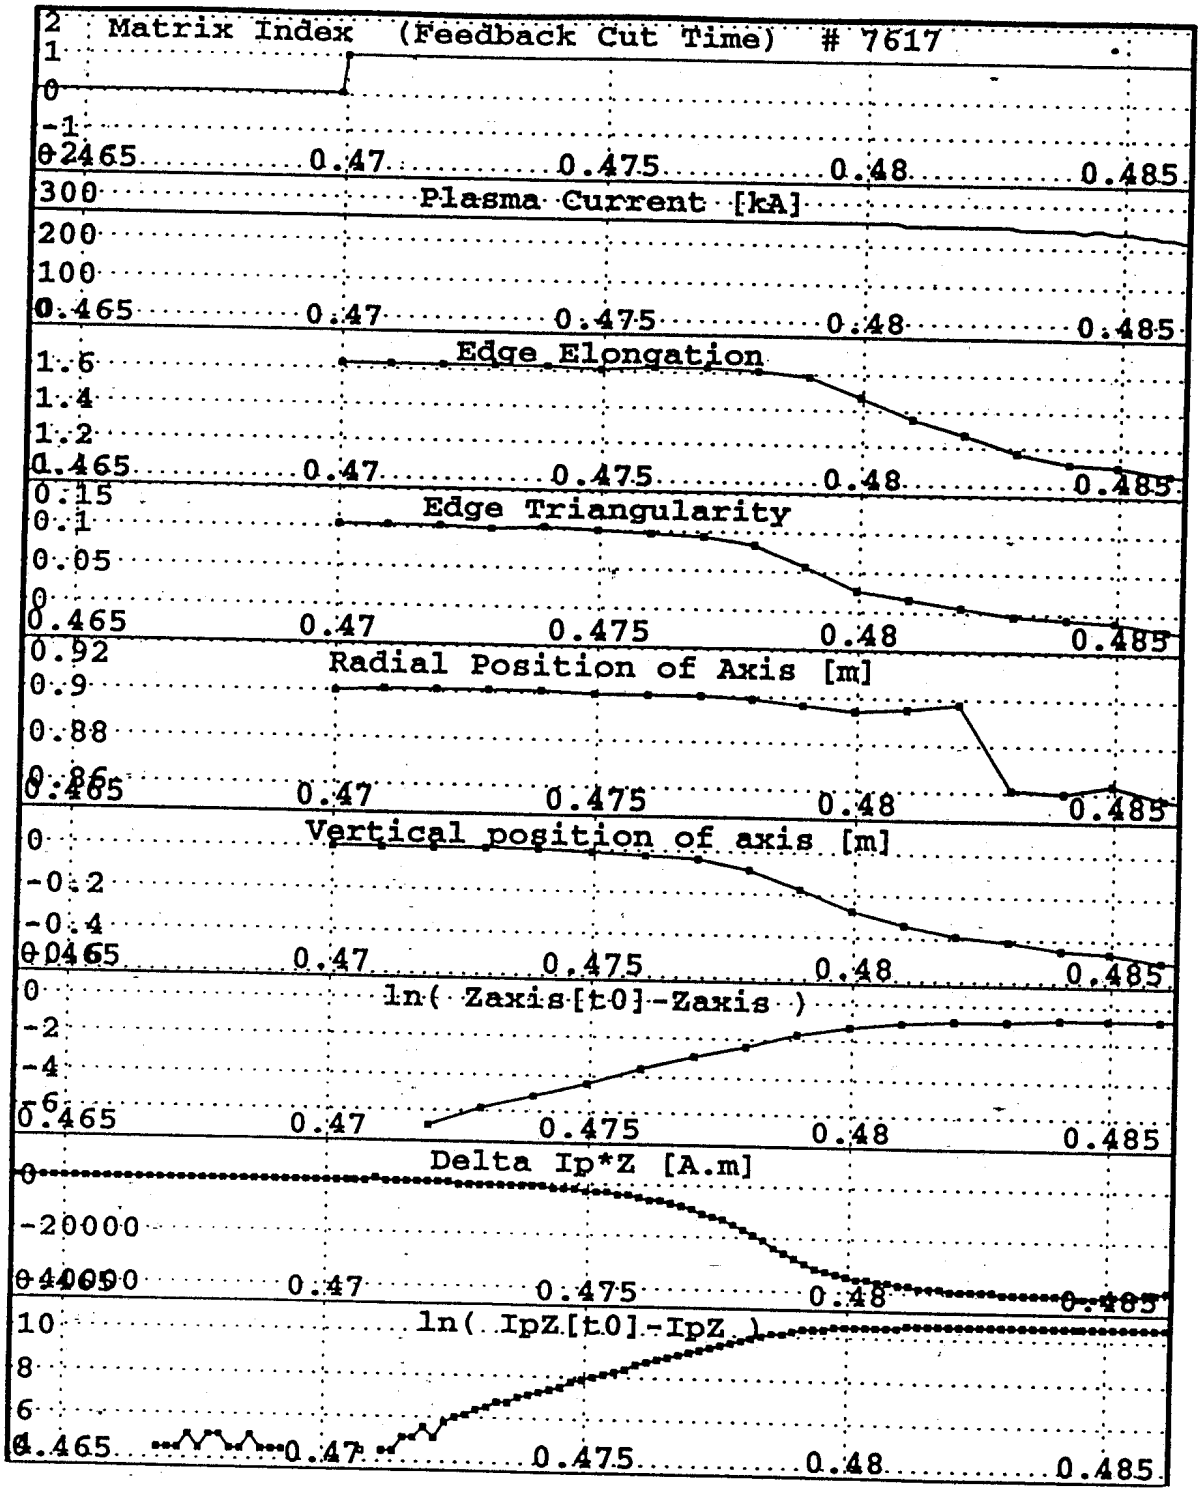

FIGURE 36 *Vertical position excursion following the inhibition of the vertical feedback control at t = 0.47 s.*

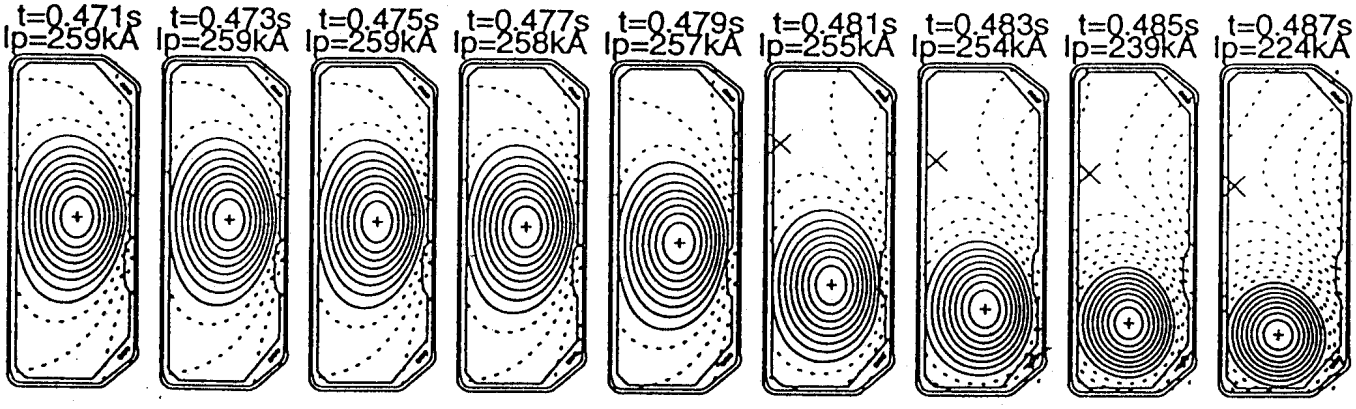

FIGURE 37 *Time evolution of the discharge shown in Fig. 36 after the vertical feedback cut.*

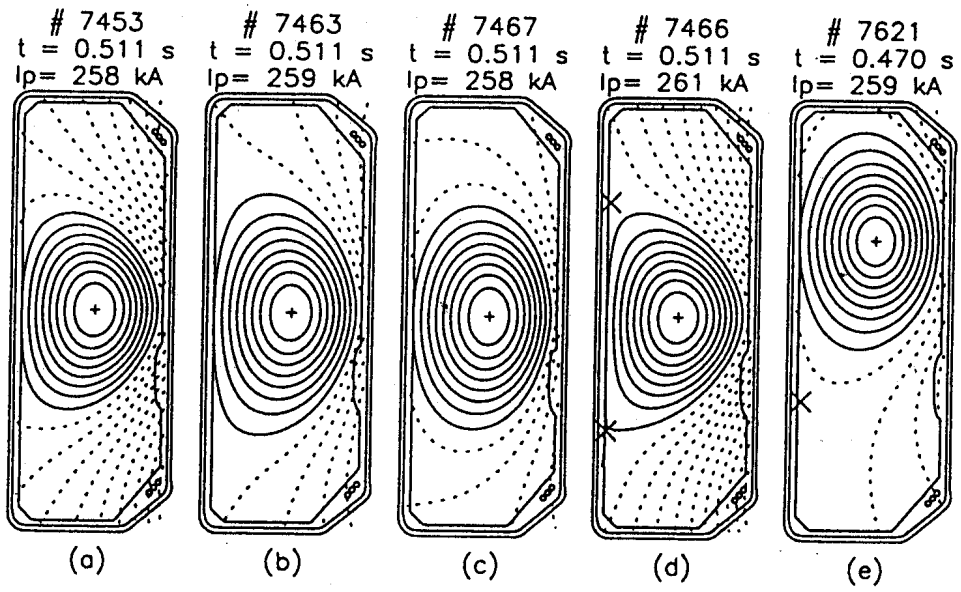

FIGURE 38 *Plasma configurations: extremes of elongation, triangularity and initial vertical position.*

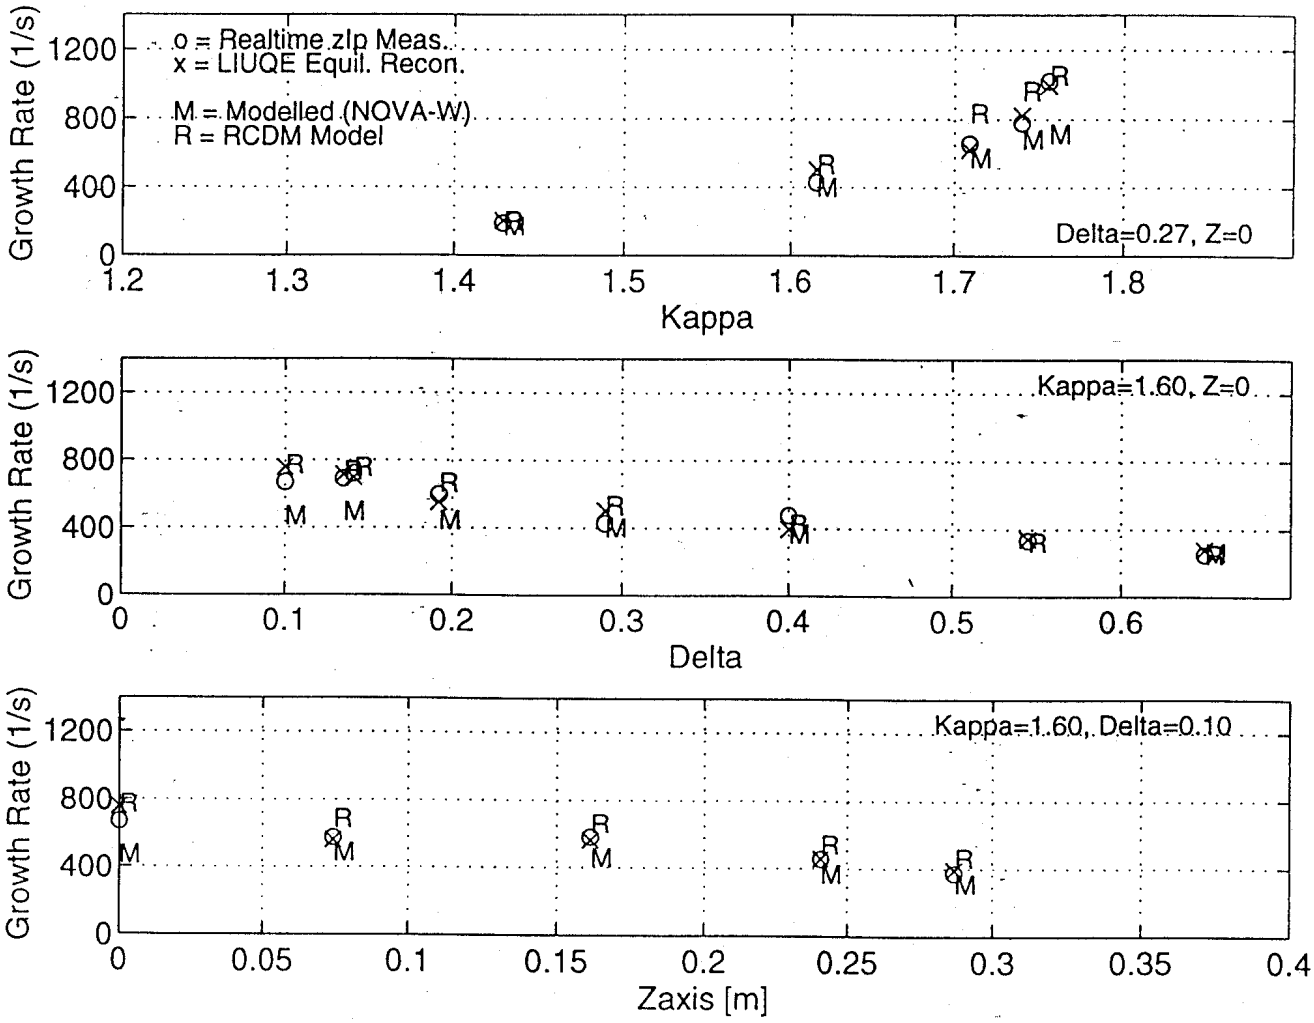

FIGURE 39 *Comparison between experimental and modelled vertical growth rates in TCV as a function of a) elongation, b) triangularity and c) initial vertical position. The results of the RCDM modelling are shown reduced by 10%.*

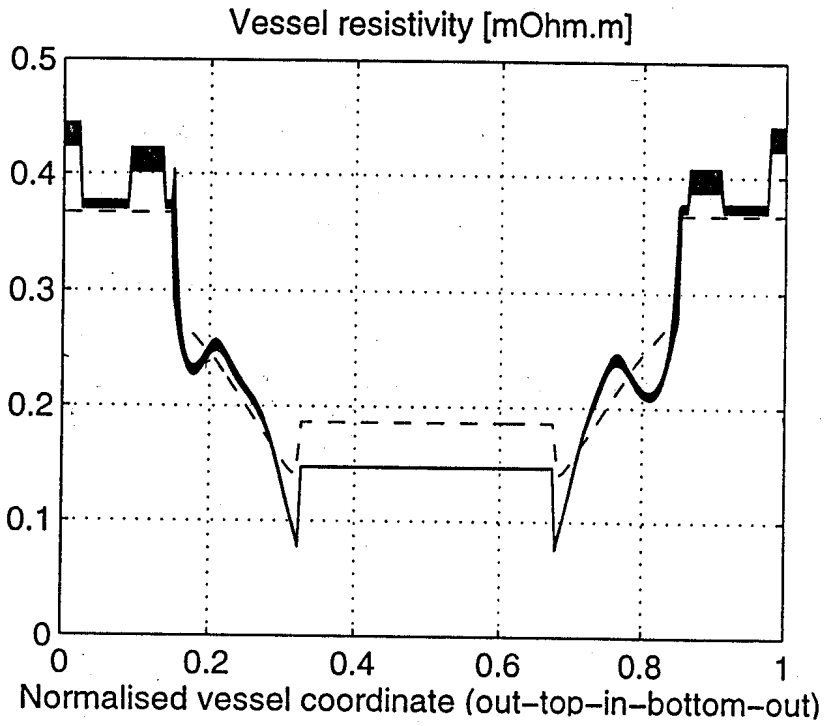

FIGURE 40 *The measured vessel resistivity, also indicating the rough estimation obtained from the vessel geometry and its specific resistivity.*

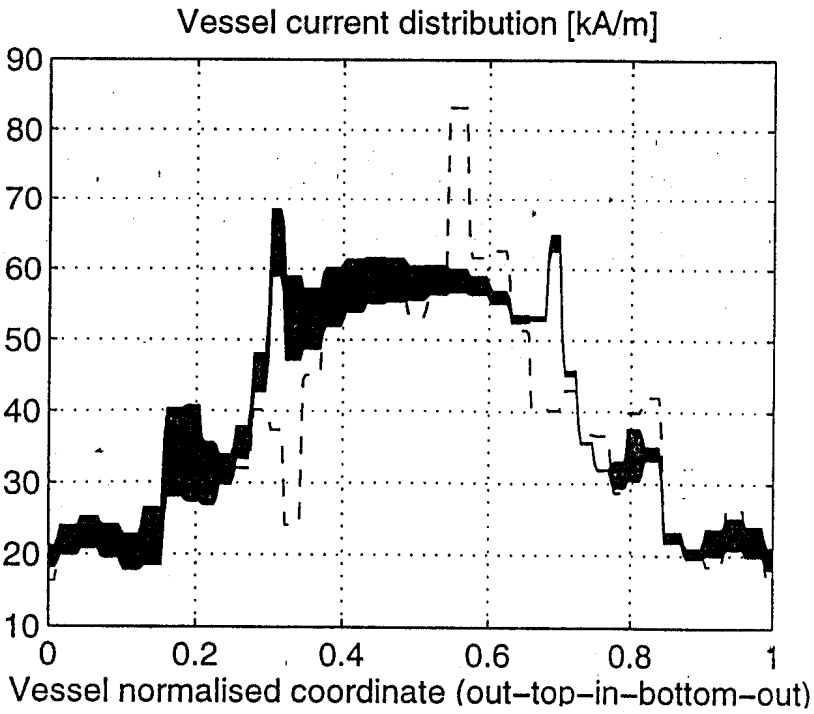

FIGURE 41 *The difference in the vessel current distribution estimated by the two methods discussed.*

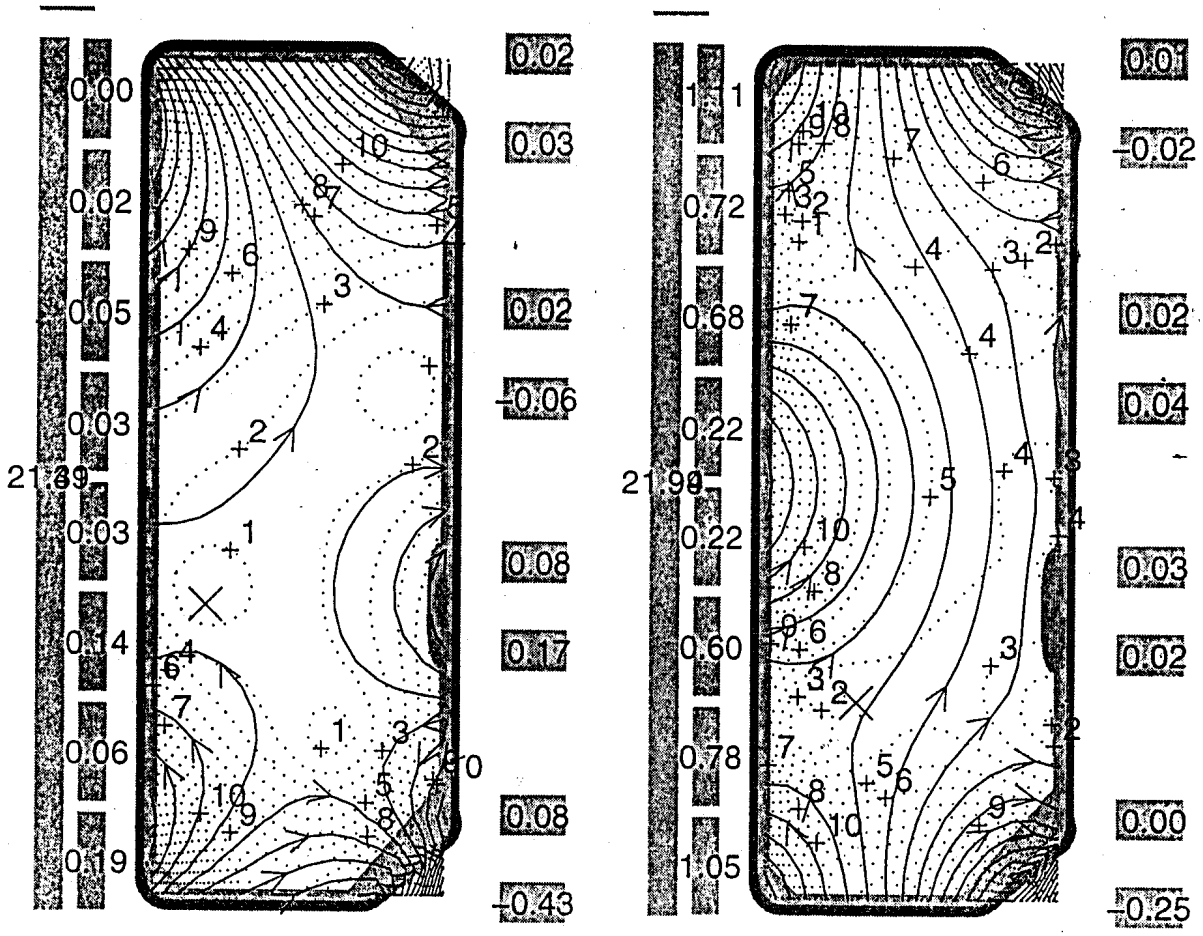

FIGURE 42 *An example of the magnetic field configuration prior to breakdown. Left: for z = -0.23 m the field amplitude is less than 1 mT in a region of more than 0.2 m in diameter. Right: the field configuration for a simultaneous breakdown at z = ±0.40 m used for the creation of a doublet.*

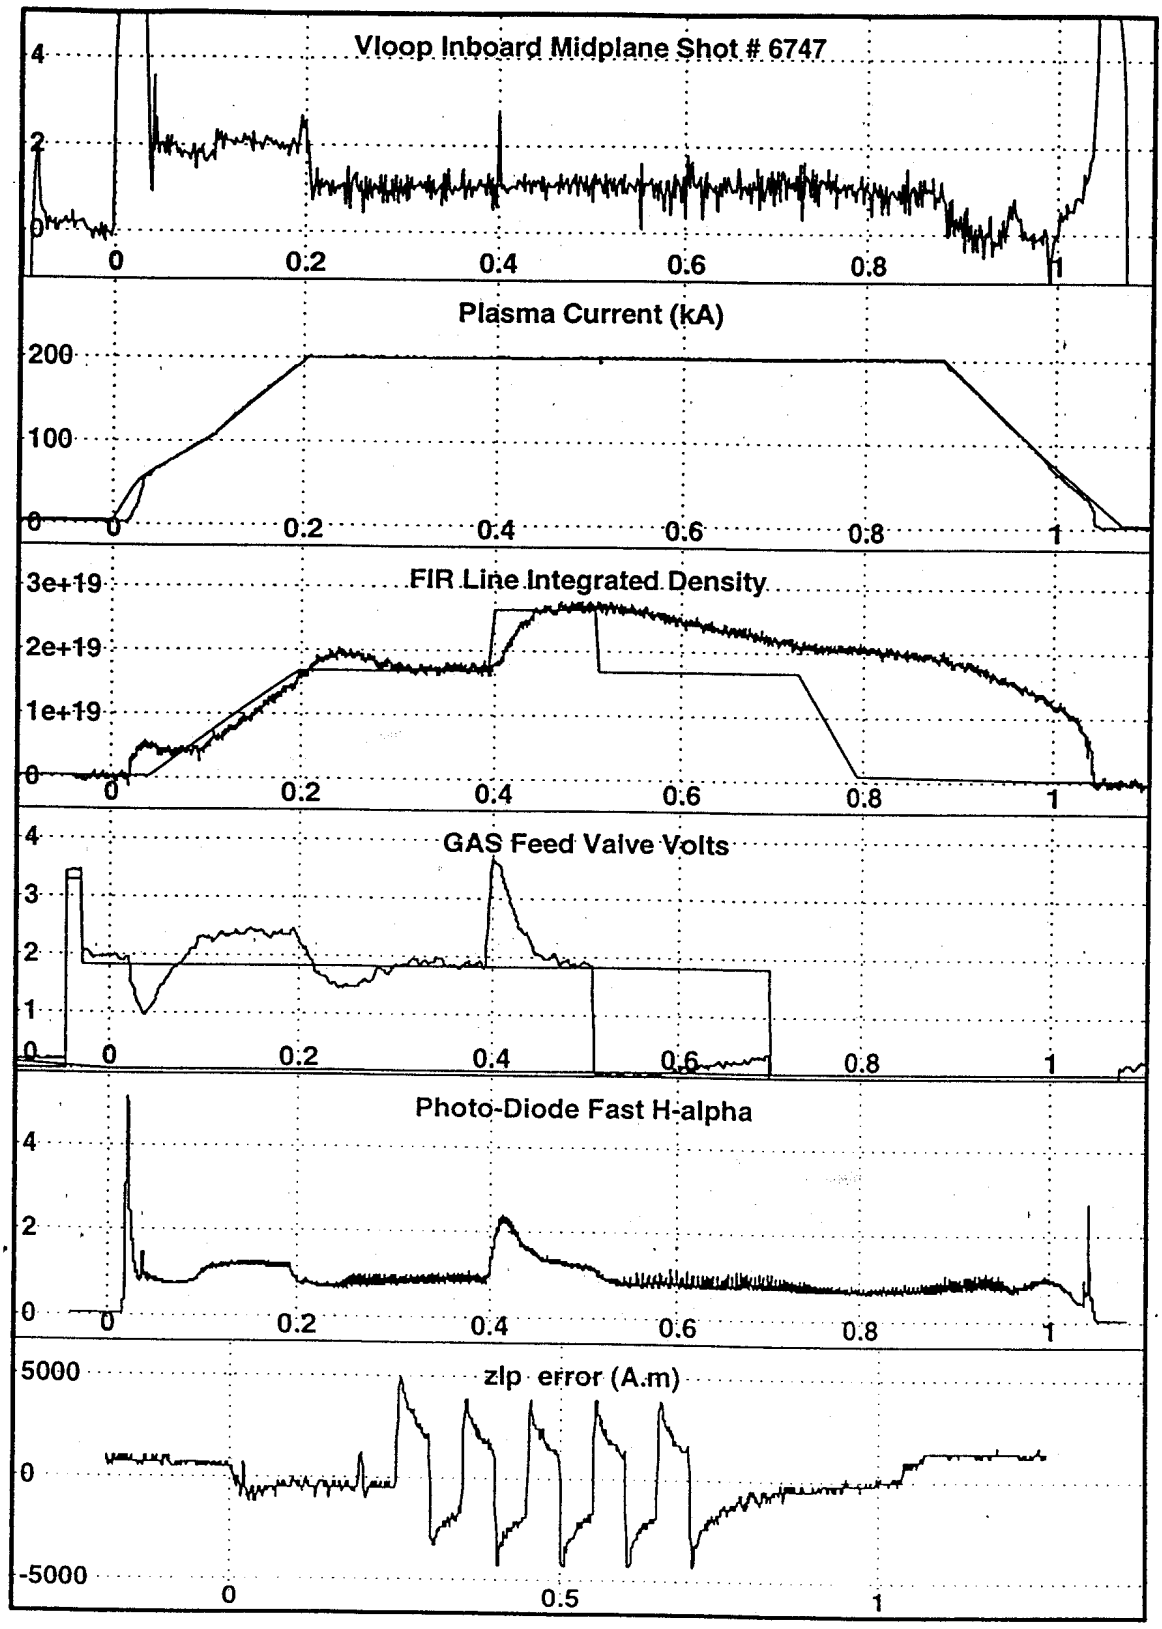

FIGURE 43 *Example of simple density feedback control.*

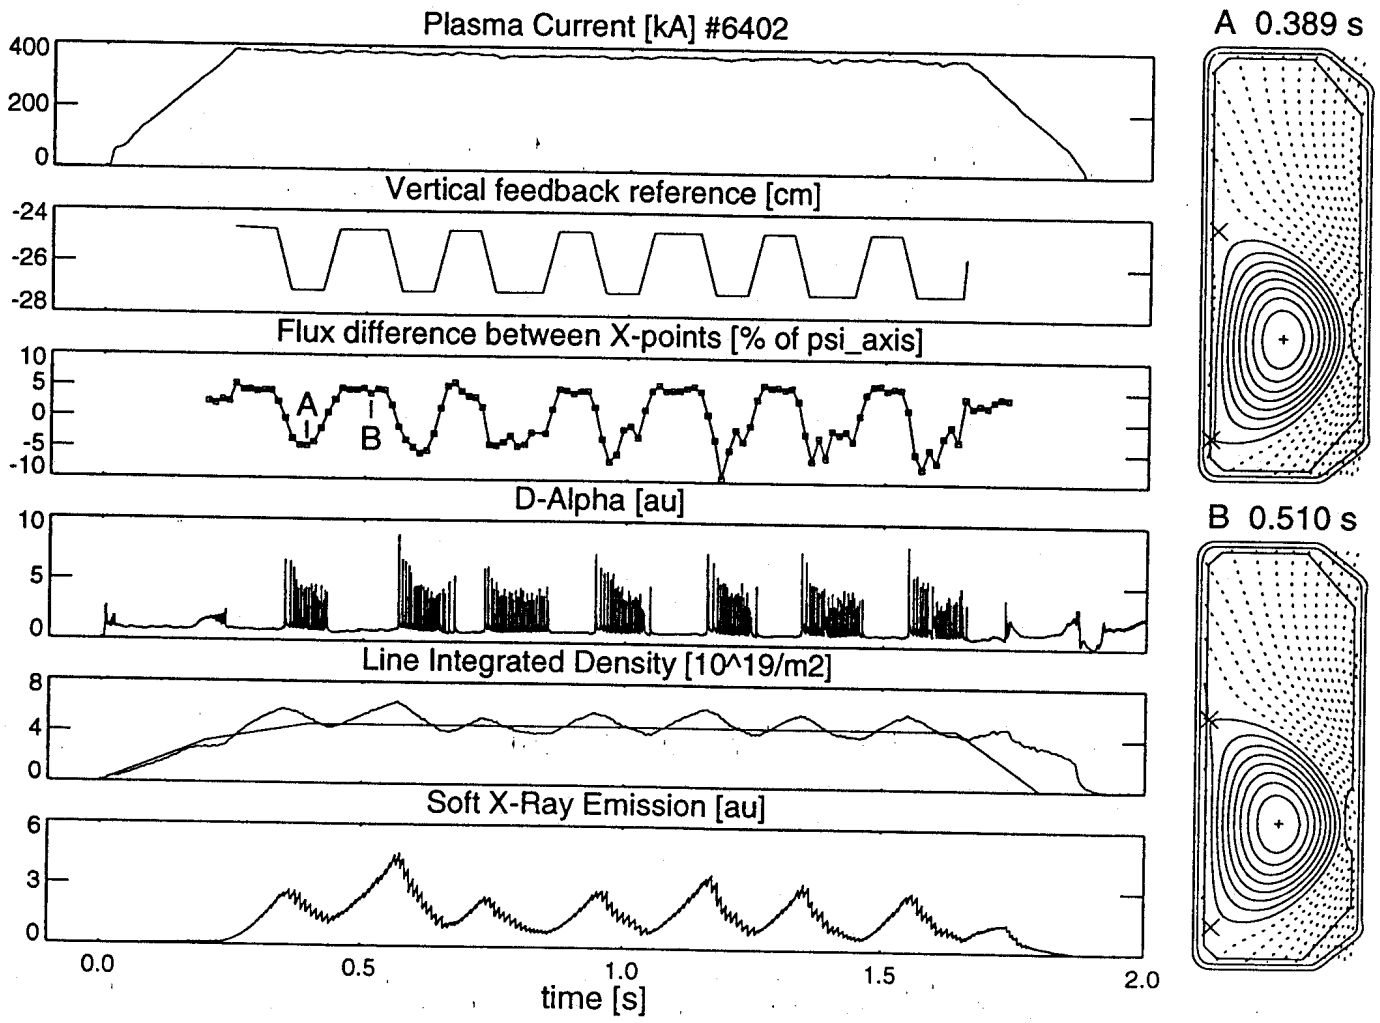

FIGURE 44 *Example of density control using Double-Null configurational modifications.*