

Plasma equilibrium response modelling and validation on JT-60U

J.B. Lister¹, A. Sharma², D.J.N. Limebeer², Y. Nakamura³,
J.P. Wainwright² and R. Yoshino³

¹ Centre de Recherches en Physique des Plasmas, Association EURATOM-Confédération Suisse, Ecole Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland

² Imperial College of Science, Technology and Medicine, Centre for Process Systems Engineering, Prince Consort Road, London SW7 2BY, UK

³ JAERI Naka Fusion Research Establishment, Naka-machi, Naka-gun, Ibaraki-ken 311-0193, Japan

E-mail: Jo.Lister@epfl.ch

Received 4 May 2001, accepted for publication 26 January 2002

Published 6 June 2002

Online at stacks.iop.org/NF/42/708

Abstract

A systematic procedure to identify the plasma equilibrium response to the poloidal field coil voltages has been applied to the JT-60U tokamak. The required response was predicted with a high accuracy by a state-space model derived from first principles. The *ab initio* derivation of linearized plasma equilibrium response models is re-examined using an approach standard in analytical mechanics. A symmetric formulation is naturally obtained, removing a previous weakness in such models. RZIP, a rigid current distribution model, is re-derived using this approach and is compared with the new experimental plasma equilibrium response data obtained from Ohmic and neutral beam injection discharges in the JT-60U tokamak. In order to remove any bias from the comparison between modelled and measured plasma responses, the electromagnetic response model without plasma was first carefully tuned against experimental data, using a parametric approach, for which different cost functions for quantifying model agreement were explored. This approach additionally provides new indications of the accuracy to which various plasma parameters are known, and to the ordering of physical effects. Having taken these precautions when tuning the plasmaless model, an empirical estimate of the plasma self-inductance, the plasma resistance and its radial derivative could be established and compared with initial assumptions. Off-line tuning of the JT-60U controller is presented as an example of the improvements which might be obtained by using such a model of the plasma equilibrium response.

PACS numbers: 52.55-s, 52.55.Fa

1. Introduction

The design and operation of next generation tokamaks will require accurate models of the dynamic plasma equilibrium response to poloidal field (PF) coil voltages. These models are needed to predict the behaviour during all types of operation and have particular relevance to the design of an effective feedback control system. The plasma equilibrium control system for a tokamak fusion reactor will have to be operated within stringent tolerances (e.g. the plasma boundary should be constrained to a few centimetres) and must be able to avoid current saturation of the superconducting PF coils and over-actuation of the PF coil power supplies. Typically, the plasma position and shape control will act to regulate a set of plasma equilibrium parameters at pre-set values whilst controlling the instability of the vertical position. Modern model-based control methodologies have been developed to explicitly address such issues.

There are many methods of constructing linearized plasma equilibrium models in a form suitable for feedback control design, typified by the CREATE-L deformable plasma response model [1], the DPM deformable plasma response model [2] and the RZIP rigid current displacement model [3]. Until recently there was no way of deciding whether these models are sufficiently accurate. The main restriction to the critical analysis of the suitability of a model was that the experiments always had to be performed in the presence of the stabilizing vertical position control loop. In practice this meant that the predicted behaviour of the model was significantly influenced by the control loop and information was partly concealed [4, 5].

In order to remedy the problem of only having access to closed-loop data for comparison with models, a series of identification experiments was performed earlier on the tokamak à configuration variable (TCV) [3, 6]. By taking data for the system identification inside the control loop, an

open-loop model of TCV (i.e. a model of the TCV dynamic response without the control loop) was created purely from the closed-loop experimental data.

The identified open-loop model was then compared with two tokamak models, CREATE-L and RZIP. The main conclusion of this earlier work on TCV was that both the models tested were reasonably accurate in predicting the dynamic response of TCV.

This previous work left several outstanding issues unresolved which fall into four broad categories, all addressed by this paper:

- Many linearized tokamak models do not consider both energy and flux conservation in a consistent manner.
- TCV is a relatively small tokamak with an electron temperature of about 1 keV when Ohmically heated; there may be higher-temperature phenomena not apparent in TCV which become important for larger machines, such as large edge currents.
- The method used to quantify model accuracy was *ad hoc* and as such was difficult to justify.
- Previous identification experiments did not discuss tuning the electromagnetic model in the case of disagreement, since the agreement was considered adequate.

This paper retraces some of the previous work in an improved way. It illustrates the full methodology which proceeds from the definition of a simple plasma equilibrium response model, through the execution and analysis of a set of experiments, concluding with a discussion of new techniques for adjusting the model to improve the agreement between the experimental results and the underlying circuit equations.

The first step presents a new formulation of the RZIP rigid current displacement model in section 2, created in order to address the issues of flux and energy conservation. This new method of derivation considers the Lagrangian of the system, a standard method in analytical mechanics, and incorporates resistive effects in a simple and natural way, yet makes a minimal set of assumptions. The resulting model supersedes the previously derived RZIP model [3]. All of the previously used linear tokamak models can, in principle, be derived using the Lagrangian method, and so can be expressed in the same framework. The existence of a general form for all models that satisfy the same few basic assumptions gives support to the grey-box modelling approach outlined in previous work [3]. By grey-box modelling we refer to an approach in which *a priori* knowledge of a system is used to define a certain structure, so that experimental information is used to refine the knowledge of parameters in the chosen model structure. In using this grey-box modelling technique it was assumed that substantial parts of the model description are accurate, especially those parts which are based on the electromagnetic description of the tokamak. Consequently, only those few values in the model which are dominated by the plasma behaviour were chosen for optimization.

The second step is a series of identification experiments carried out on the JT-60U tokamak [7]. The possibility that the positive result on TCV was limited to a relatively small and cool plasma has been eliminated since this tokamak has a much larger plasma cross-section (4 versus 0.5 m²), a much larger plasma current (2.7 versus 1 MA) and a higher plasma energy content than TCV.

Apart from its size, the JT-60U PF system differs from TCV in three significant respects. In normal plasma operation current demand signals are sent to the PF coils rather than voltage demand signals. In the work described, the voltages actually applied by the power supplies to the coils are measured, to obtain the plasma equilibrium responses to the voltages. The 43 coils present in the JT-60U PF coil system are connected as five independently powered, composite coil sets, each designed to control a different plasma equilibrium property. Finally, the time-scale of the vertical instability is much longer, of order 1 s as opposed to a few tens of milliseconds for the reported TCV experiments. The JT-60U tokamak and the generation of its plasmaless (i.e. purely electromagnetic) RZIP model are detailed in section 3.

The identification experiments were carried out in two phases, experiments without plasma and experiments with plasma. The third step, namely the design and analysis of the plasmaless experiment, is described in section 4.

The fourth step is an advance over previous work, namely an attempt to fine-tune the plasmaless model to agree better with the available experimental data, not necessary in the TCV work. The method of improving the basic model by using the plasmaless results is described in section 5. The agreement is quantified by defining a cost function which compares the difference between the model predictions and the experimental data. The final cost functions described in this paper were arrived at after an exhaustive iterative process, and are the result of a judgement of the best compromise between functionality and methodological purity. Quantifying the model agreement on the basis of a particular cost function suffers from not being generalizable. Using qualitative descriptions of 'good' or 'poorer' agreement is equally unsatisfying. Ultimately, the agreement between a model and a set of measurements depends on the purpose for which the agreement has been tested. If it is to answer a physics question, then the possibility of discriminating between competing models is the most important feature. If, as in the case of this paper, the underlying motivation is the development of models which are adequate for designing a feedback controller, then other considerations are important and the behaviour of the closed loop becomes the main issue, addressed illustratively in section 8.

The fifth step is the set of plasma experiments detailed in section 6. These experiments were carried out with Ohmically heated discharges as well as discharges heated by additional neutral beam injection (NBI), allowing generalization of the results from TCV to a more reactor-relevant plasma. The sixth step in section 7 presents the grey-box adjustment of the plasma model, using the same method as section 5.

After this sixth step, we have completed the route from a nominal model to an experimentally adjusted model in agreement with the available data. The methodology of this somewhat long and complex chain can be summarized as follows:

- Establish a nominal model on the basis of well-known concepts and the Lagrangian techniques, on the assumption of displacements of a fixed normalized current distribution.
- Perform a series of experiments in the absence of plasma to establish the response of all measurements to PF coil voltage variations at different frequencies.
- Provide a means of quantifying the agreement between the modelled and experimental measurements.

- Modify the coefficients in the circuit equations governing the electromagnetic couplings and the diagnostic sensitivities in order to minimize the differences between the modelled and experimental responses.
- Perform a series of experiments with plasma to establish the responses of all the measurable quantities.
- Adjust the coefficients in the circuit equations which depend on the presence of the plasma, to minimize the differences between the modelled and experimental responses.

In section 8 we demonstrate that the grey-box RZIP model accurately reproduces the closed-loop behaviour of JT-60U by comparison with experiment. An illustration of the benefits of this methodology is the demonstration that a controller tuned with this model can counteract undesirable closed-loop behaviour, specifically some observed coil-set cross-couplings. Section 9 discusses the implications of our results for assessing the model uncertainty, especially useful for evaluating the performance of controllers in the presence of such uncertainties. We conclude in section 10.

2. Derivation of the RZIP linearized model

Low-order accurate models of the dynamic response of tokamak equilibria are needed for the design of multi-variable model-based tokamak controllers. By rigorously defining a simple lumped-parameter linear tokamak equilibrium model from a clear set of assumptions, we facilitate comparisons with data from open-loop system identification experiments. Present low-order tokamak modelling methods often rely on equations separately derived using assumptions that are not necessarily consistent. As a result, these models do not always explicitly conserve energy or flux [1, 3]. To address this we derive a tokamak model from a minimal set of assumptions within a Lagrangian formalism, obtaining a model that consistently considers flux, momentum and energy. The resulting equations are found to have the same structure as those derived in the previous work, given the same choice of the total plasma current, the tokamak vessel and coil currents, and the plasma vertical and radial position as system states.

In this model we assume that the tokamak system is fully described by the toroidal current density in the plasma and surrounding structure. A cylindrical coordinate system (R, z, ϕ) is used. We describe the plasma by a small number of degrees of freedom: the current profile, the plasma current, the plasma thermal energy, and the plasma position.

We make a number of initial assumptions:

- all quantities are independent of the toroidal angle ϕ (axisymmetry),
- the plasma has negligible mass, therefore reaches an equilibrium instantaneously,
- the system may be perturbed about that equilibrium,
- poloidal currents in the plasma and structure may be ignored,
- plasma transport effects can be ignored,
- the equilibrium PF coil currents are constant.

The tokamak's physical structure can be classified into two parts, the active structure and the passive structure. The active structure is the coil system, to which we can externally

apply voltages, whilst no voltages are applied to the plasma or the passive structure. The passive structure therefore carries induced eddy currents. This model considers the different parts of the tokamak separately. The continuous conducting parts of all the structure are discretized into a number of toroidal elements each with an individual toroidal current. The poloidal elements of the currents in the coils and passive structure are taken to be zero.

These assumptions yield a set of four non-linear differential equations, which can be linearized about a given tokamak equilibrium.

2.1. Derivation of the model equations by a Lagrangian method

Here we use the following notation:

- $\delta\alpha$ is a small variation in α from a static equilibrium α^0 such that $\alpha = \alpha^0 + \delta\alpha$;
- $\dot{\alpha}$ is the continuous time derivative d/dt of α ;
- $\partial\alpha/\partial\varepsilon|_0$ is the derivative $\partial\alpha/\partial\varepsilon$ evaluated at equilibrium.

The total plasma current will be represented by I_p , and the currents in the structure elements will be represented by the vector I_s . In the case of JT-60U $I_s = [I_D I_F I_T I_V I_H I_{\text{passive}}^T]^T$, where the quantity I_x is the current in coil x , and the I_{passive} is the vector of currents flowing through the elements of the passive structure. The voltages applied to the structure elements will be represented by the vector V_s . The equilibrium plasma current density distribution j_ϕ is estimated by an inverse equilibrium reconstruction code.

The plasma radial position R is defined by a current-weighted average of plasma element radial positions [10],

$$R = \frac{\int_{\text{plasma}} jr \, dS}{\int_{\text{plasma}} j \, dS}, \quad (2.1)$$

where S is the plasma cross-section, and r is the major radial coordinate. The vertical position z is defined similarly.

We define the effective plasma self-inductance L_p via the equivalent energy of the total current distribution:

$$\frac{1}{2} L_p I_p^2 = \frac{1}{2} \iint_{\text{plasma}} j_i M_{ik} j_k \, dS_i \, dS_k, \quad (2.2)$$

where M_{ik} is the mutual inductance between two elements i and k , for i different to k , and the self-inductance of element i for $i = k$. The effective mutual inductance matrix between the plasma and structure M_{ps} is

$$I_p M_{ps} I_s = \int_{\text{plasma}} j_i M_{is} I_s \, dS_i, \quad (2.3)$$

where M_{is} is the mutual inductance between an element i and the vector of structure element currents.

We define W_T as the thermal energy of the plasma.

2.2. Choice of generalized coordinates

In the derivation of the tokamak governing equations, we will use (Q_s, Q_p, R, z) as our generalized coordinates. The quantities R and z are, respectively, the plasma radial and vertical position. The quantities Q_s and Q_p represent the charge which has flowed since time t_0 through the structure and the plasma, respectively. These constitute generalized coordinates in the classical theory [11]. It is easy to see that the corresponding currents are $I_s = Q_s$ and $I_p = Q_p$.

2.3. Derivation of the model equations

The Lagrangian, L , of a system is defined as

$$L = T - V,$$

where T is the generalized kinetic energy of the system and V is the generalized potential energy of the system.

For our system we can write

$$\begin{aligned} T &= \frac{1}{2} I_s^T L_s I_s + \frac{1}{2} L_p I_p^2 + I_p M_{ps} I_s, \\ V &= -Q_s^T V_s - W_T. \end{aligned} \quad (2.4)$$

This gives the expression for the Lagrangian

$$L = \frac{1}{2} I_s^T L_s I_s + \frac{1}{2} L_p I_p^2 + I_p M_{ps} I_s + Q_s^T V_s + W_T \quad (2.5)$$

that we will now associate with the Euler–Lagrange equation:

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{q}_i} \right) - \frac{\partial L}{\partial q_i} = \frac{\partial P}{\partial \dot{q}_i}, \quad i = 1, 2, \dots, n$$

in which the q_i 's are the generalized coordinates. The right-hand side is associated with dissipation effects due to the plasma and structure resistances. In our case the power dissipation function can be expressed as

$$P = -\frac{1}{2} (I_p^T \Omega_p I_p + I_s^T \Omega_s I_s). \quad (2.6)$$

The equations for each generalized coordinate are then:

- (Q_s) The structure flux conservation equation

$$\frac{d(L_s I_s)}{dt} + \frac{d(M_{ps} I_s)}{dt} + \Omega_s I_s = V_s. \quad (2.7)$$

- (Q_p) The plasma flux conservation equation

$$\frac{d(L_p I_p)}{dt} + \frac{d(M_{ps} I_s)}{dt} + \Omega_p I_p + \frac{d}{dt} \left(\frac{\partial W_T}{\partial I_p} \right) = 0. \quad (2.8)$$

- (z) The vertical force balance equation

$$\frac{1}{2} \frac{\partial (L_p I_p^2)}{\partial z} + \frac{\partial (I_p M_{ps} I_s)}{\partial z} + \int \frac{\partial \Omega_p I_p}{\partial z} dQ_p + \frac{\partial W_T}{\partial z} = 0. \quad (2.9)$$

- (R) The radial force balance equation

$$\frac{1}{2} \frac{\partial (L_p I_p^2)}{\partial R} + \frac{\partial (I_p M_{ps} I_s)}{\partial R} + \int \frac{\partial \Omega_p I_p}{\partial R} dQ_p + \frac{\partial W_T}{\partial R} = 0. \quad (2.10)$$

The methodology used to derive equations (2.7)–(2.10) does not assume the rigid current displacement of the plasma. These equations have a structure comparable to the equations used in previous work [3]. A plasmaless model will simply have all the plasma-related terms equal to zero, resulting in a purely electromagnetic model.

2.4. Linearization of equations

The equations derived have the variables (R, z, I_p, I_s) . However, for consistency with earlier work [3] we use the products $(R I_p^0, z I_p^0, I_p)$. The quantities (R, z, I_p) are simply related to $(R I_p^0, z I_p^0, I_p)$. In consideration of this, we can define a state vector x using these quantities;

$$x = \begin{bmatrix} I_s - I_s^0 \\ (z - z^0) I_p^0 \\ (R - R^0) I_p^0 \\ I_p - I_p^0 \end{bmatrix} = \begin{bmatrix} \delta I_s \\ (\delta z) I_p^0 \\ (\delta R) I_p^0 \\ \delta I_p \end{bmatrix}. \quad (2.11)$$

The four physics equations are therefore linearized in x about the fixed point $x^0 = 0$, to give four linear equations in x and \dot{x} . These linear matrices can then be cast in the standard state-space model form

$$\begin{aligned} \dot{x} &= Ax + Bu, \\ y &= Cx + Du. \end{aligned} \quad (2.12)$$

The use of the state-space representation of a linear dynamical system is purely for convenience, due to the large number of techniques developed for analysing multi-dimensional transfer functions as a function of frequency or of time. It would be difficult to imagine a convenient method avoiding this standard usage.

2.5. The state-space model of a Tokamak

The linearized structure circuit equation, plasma circuit equation, and plasma force balance equations can be represented as follows:

$$\begin{bmatrix} L_s & \frac{\partial M_{sp}}{\partial z} \Big|_0 & \frac{\partial M_{sp}}{\partial R} \Big|_0 & M_{sp}^0 \\ \frac{\partial M_{ps}}{\partial z} \Big|_0 & \left(\frac{\partial^2 M_{ps}}{\partial z^2} \Big|_0 \frac{I_s^0}{I_p^0} \right) & \left(\frac{\partial^2 M_{ps}}{\partial z \partial R} \Big|_0 \frac{I_s^0}{I_p^0} \right) & 0 \\ \frac{\partial M_{ps}}{\partial R} \Big|_0 & \left(\frac{\partial^2 M_{ps}}{\partial z \partial R} \Big|_0 \frac{I_s^0}{I_p^0} \right) & \left(\frac{1}{2} \frac{\partial^2 L_p}{\partial R^2} \Big|_0 + \frac{\partial^2 M_{ps}}{\partial^2 R} \Big|_0 \frac{I_s^0}{I_p^0} \right) & \left(\frac{1}{2} \frac{\partial L_p}{\partial R} \Big|_0 + \frac{\partial M_{ps}}{\partial R} \Big|_0 \frac{I_s^0}{I_p^0} + \frac{\partial^2 W_T}{\partial I_p \partial R} \frac{1}{I_p^0} \right) \\ M_{ps}^0 & 0 & \left(\frac{1}{2} \frac{\partial L_p}{\partial R} \Big|_0 + \frac{\partial M_{ps}}{\partial R} \Big|_0 \frac{I_s^0}{I_p^0} + \frac{\partial^2 W_T}{\partial I_p \partial R} \frac{1}{I_p^0} \right) & L_p^0 + \frac{\partial^2 W_T}{\partial I_p^2} \end{bmatrix} \dot{x} + \begin{bmatrix} \Omega_s^0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{\partial \Omega_p}{\partial R} \Big|_0 \\ 0 & 0 & \frac{\partial \Omega_p}{\partial R} \Big|_0 & \Omega_p^0 \end{bmatrix} x = \begin{bmatrix} \delta V_s \\ 0 \\ \frac{\partial^2 W_T}{\partial R \partial \alpha} \dot{\alpha} \\ \frac{\partial^2 W_T}{\partial I_p \partial \alpha} \dot{\alpha} \end{bmatrix}. \quad (2.13)$$

Here the plasma internal energy W_T is allowed to vary with an introduced variable α that is independent of the states. For clarity we have taken W_T and the plasma current distribution j to be invariant under translation in z (consequently the plasma resistance and self-inductance are also invariant under translation), though it is not strictly necessary at this stage. This is of the required form $M\dot{x} + Px = u$. Considering the magnetic field probe and flux probe outputs as functions of the currents in the tokamak provides the output matrix C and feed-forward matrix D . Comparing (2.13) with (2.12) gives the matrix of coefficients A , and the control matrix B :

$$\begin{aligned} A &= -M^{-1}P, \\ B &= M^{-1}. \end{aligned} \quad (2.14)$$

Note the symmetry of M and P , which implies that hysteresis is not a feature of the plasma model. This is a necessary consequence of the formalism used. A brief inspection shows that P and the bottom right-hand four terms of M contain the only purely plasma response terms.

From a minimal set of assumptions we have derived a linear, time invariant model in state-space form. All linearized tokamak models with the same choice of states can be expressed in this structural form. Within this formalism we can derive all models that assume toroidal symmetry and perturb about an MHD equilibrium.

2.6. Approximations

Further to the assumptions detailed in section 2.1, the RZIP model also assumes a rigid current displacement, namely that the normalized current profile is independent of movements in the R and z directions and of changes in plasma current. This allows us to calculate the plasma mutual and self-inductance derivatives simply and directly.

For example, the radial derivative of the mutual inductance between two plasma elements f and g ,

$$\frac{\partial M_{fg}}{\partial R} = \frac{\partial M_{fg}}{\partial R_f} \frac{\partial R_f}{\partial R} + \frac{\partial M_{fg}}{\partial R_g} \frac{\partial R_g}{\partial R}.$$

A pure rigid current displacement assumption fixes the two radius relations as

$$\frac{\partial R_f}{\partial R} = \frac{\partial R_g}{\partial R} = 1.$$

Given this assumption we can approximate $\partial L_p / \partial z = 0$. Similarly we can take $\partial W_T / \partial z = 0$.

The poloidal beta (the ratio between averaged kinetic and magnetic pressures) is defined as

$$\beta_p = \frac{\int_S p \, ds / \int_S ds}{B_a^2 / 2\mu_0}, \quad (2.15)$$

where p is the pressure in the plasma, S is the plasma cross-section, and

$$B_a = \frac{\mu_0 I_p}{l},$$

l being the poloidal circumference of the plasma.

If we take $W_T = \int_{\text{plasma}} p \, dV$ (where V is volume), then by writing $dV = 2\pi R \, dS$ we can substitute for β_p , immediately giving

$$W_T = \pi \mu_0 \frac{S}{l^2} \beta_p R I_p^2.$$

From this expression we may take $\alpha = \beta_p$ in (2.13) and the values for the derivatives of W_T . Note that this results in perturbative terms in $\dot{\beta}_p$ on the right-hand side of equation (2.13). We have neglected the B_ϕ term as we have chosen to ignore all poloidal currents.

There are several approaches for estimating the total inductance of the plasma current distribution, which are more or less computationally intensive. In order to keep the determining equations simple, we use the following approximation to the plasma self-inductance defined in equation (2.2):

$$L_p = \mu_0 R \left(4\pi R \frac{\bar{B}_v}{\mu_0 I_p} - \beta_p - \frac{1}{2} \right),$$

where $\bar{B}_v S = \overline{M_{ps} I_s}$ and where the overscript implies averaging over the plasma current distribution as in equation (2.1). This approximation gives almost identical values compared with longer calculations. The expression has the advantage that the positional derivatives are simply calculable. If we were to use a formula for L_p that included the internal inductance l_i as an independent variable, such as $L_p = \mu_0 R (\ln(8R/a) + l_i/2 - 2)$, where a is the minor radius, this would also result in terms in \dot{l}_i on the right-hand side of equation (2.13). Strictly speaking however, allowing l_i to vary would change the current profile, which is a violation of the rigid current distribution assumption.

These formulae are all based on the assumption that the current distribution is displaced vertically and horizontally when the vacuum poloidal flux pattern is changed. This assumption is crude and neglects changes in the plasma cross-sectional area, changes in the plasma current profile, changes in the force due to the toroidal field, the induction and dissipation of all currents in the poloidal plane and the question of helicity conservation. These effects are considered in detail in other models, but are excluded from the present simple model. The last formulae for the inductance are approximate, but are practical and agree well with other estimates.

3. The nominal mode of JT-60U

The JT-60U tokamak is the second largest operational tokamak with major radius 3.2 m, minor radius 0.95 m, toroidal magnetic field 4.5 T, plasma current 2.7 MA. Figure 1 shows a poloidal cross-section of the tokamak and illustrates the position of the PF coils, as well as the outline of the vacuum vessel and limiter. There are 43 separate PF coils connected in series as five independently powered PF coil sets, labelled differently on the figure. They provide poloidal flux (F coil set), vertical field (V coil set), horizontal field (H coil set), divertor field (D coil set) and a field to control the plasma cross-sectional triangularity (T coil set). Most of the five coil sets are only weakly coupled to each other but some produce a net poloidal flux and are therefore magnetically coupled by construction. The plasma elongation is controlled indirectly in

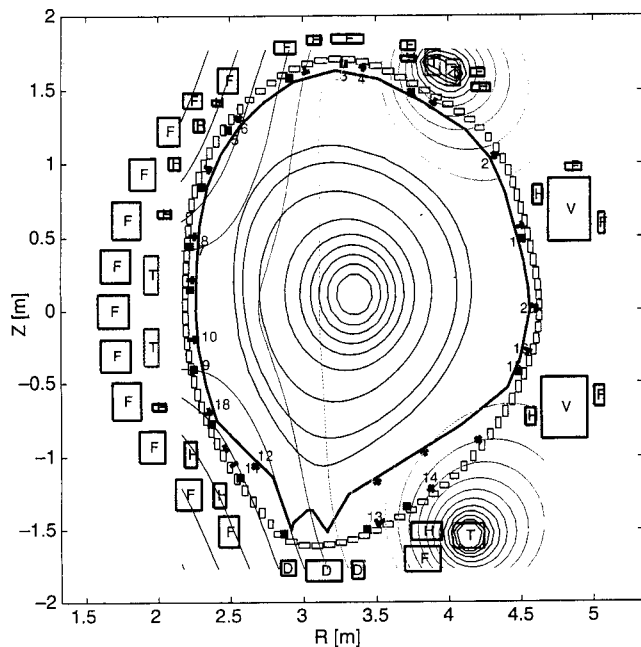


Figure 1. A cross-section of the JT-60U tokamak, showing the five PF coil sets labelled D, F, T, H, V, the diagnostics with the flux loops shown as red squares and the magnetic probes shown as short blue and magenta lines indicating their orientation. The vessel segments are shown as black rectangles and the limiting surfaces are shown as a solid blue line. The vacuum field contours for the triangularity coil set (T) are shown, together with the equilibrium flux surfaces for an Ohmically heated plasma (E35023).

JT-60U by a combination of PF coil currents. Figure 1 indicates the vacuum flux contours produced by a constant current in the T coil set (responsible for controlling the plasma triangularity). The plasma equilibrium surfaces are also shown for one of the discharges used in this work (E35023).

JT-60U is equipped with a full set of magnetic diagnostics, also indicated in figure 1. Fifteen flux loops are positioned on the vacuum vessel. Nineteen tangential PF probes are positioned close to the vacuum vessel or the divertor baffle plate and 18 normal PF probes are similarly located.

The JT-60U control system is non-linear and operates by controlling five control parameters with the five coil sets. These controlled parameters are chosen to be the vertical plasma position (Z), the radial plasma position (R), the triangularity (δ), the plasma current (I_p) and the height of the X-point from the divertor (X_p) and they are primarily controlled by the H, V, T, F and D coil sets respectively. These parameters are estimated in real time by a set of non-linear expressions optimized by regression from the input-output examples in an equilibrium database [7].

The basic electromagnetic model of JT-60U was generated from information taken from construction blueprints and other machine design descriptions.

4. Plasmaless experiments

4.1. Low-frequency calibration

The first stage in producing the validated model of JT-60U is to ensure that the purely electromagnetic description (or plasmaless model) is accurate. As such, an initial

calibration was performed to check the DC characteristics of the model against the JT-60U data.

The data used for this were pre-existing plasmaless pulses, in which each PF coil-set current is ramped up, held constant and then ramped down, in turn. The effective PF coil resistances were obtained from the flat-top currents and the steady voltages, while the effective inductance was obtained from the ramp-up and ramp-down of the coil-set currents. The values obtained from this data were slightly different from the nominal values of the model, mostly because the connection leads to the power supplies are not accounted for in the nominal model. This approach to correcting the nominal values is sensitive to small offsets in the voltage measurements which can have a significant effect on the estimated resistance. However, these pulses do allow checking of the flux-loop and PF probe positions under conditions where the currents in the passive structures are negligible. The almost DC flux loop responses agreed with the predicted values to better than 1% and the PF probe disagreement had a distribution width of about 1.5%. The responses of a small number of probes to particular PF coil voltages were well outside this normal distribution of agreement and these input-output responses were not used in our subsequent analysis. Probes which are close to perpendicular to the PF variation are particularly sensitive to this calibration. This check validates the excellent consistency between the PF current and the static diagnostic calibrations, but does not provide any check on the dynamic input-output responses of the electromagnetic model.

The vacuum vessel loop resistance was fixed at $160 \mu\Omega$, which was deduced from prior estimates made on JT-60U.

4.2. Identification experiment design

The next step was to proceed with the system identification of the plasmaless model of JT-60U. This allows checking of the dynamic response of the electromagnetic model, including the interactions between the PF coil sets and the vacuum vessel. The nominal model was used to predict the transfer function of JT-60U and it was decided that the frequency range of interest would be about 1–50 Hz, although it was also necessary to consider the upper resolution limit of the diagnostic systems. The range of interest for establishing a model for developing a feedback controller is essentially bounded on the lower side by the frequency at which the determining equations are almost static, when all induced currents have decayed. At such a low frequency the tokamak behaves as if there were no passive structure and the shape control can be derived very simply. On the higher side, the range of interest is bounded by the frequency at which the power supplies do not respond or the response of the controlled parameters is small. At this high end, the phase of the response is important. The five frequencies chosen as suitable were 1.25, 3.2, 7.5, 19 and 43 Hz, selected so that none overlap with any first or second harmonics, or with the mains frequency harmonics. A measurement is defined by the stimulation of each of the five coil sets with a sinusoidal voltage signal. During a single measurement each coil was simultaneously excited by a different one of the five different frequencies. The frequency used at each coil was then changed to another of this set of frequencies for the next measurement. A preliminary closed-loop model

was used to estimate the response to the command currents so that the power supply voltage demand limits were respected. Each measurement lasted about 1 s and it was possible to perform three experiments in a single pulse, completing the five measurements required in two pulses.

Assuming a linear response, analysis of the signals allows determination of the contribution of each coil set to each diagnostic signal (i.e. a frequency of 1.25 Hz in a diagnostic signal must be due to the 1.25 Hz components in the PF coils).

The data analysis follows the method described in [6]. The frequency components in each signal are estimated by a least squares fit of the data to a basis set of sine and cosine waves at the five experimental frequencies, plus first- and zero-order terms to remove any measurement drifts or offsets. This fit proves extremely accurate with minimal noise or harmonics left as a residual. From the sine and cosine coefficients the magnitude and phase of the signal at each frequency are calculated and define a complex amplitude for each frequency component of the signal. For each frequency ω_n , we define a complex column vector of inputs $\mathbf{u}(\omega_n)$ as the vector of the complex amplitudes of the voltages (inputs) applied to each of the five coil sets. Similarly we define a complex column vector of all of the diagnostic signal responses (outputs), $\mathbf{y}(\omega_n)$. The measurements are grouped together such that $\mathbf{U}(\omega_n) = [\mathbf{u}_1(\omega_n) \ \mathbf{u}_2(\omega_n) \ \cdots \ \mathbf{u}_k(\omega_n)]$, where the subscript on \mathbf{u} refers to each of the k separate measurements. We similarly define an output matrix $\mathbf{Y}(\omega_n)$. The transfer function $\mathbf{G}(\omega_n)$ is defined by the input–output mapping, $\mathbf{Y}(\omega_n) = \mathbf{G}(\omega_n)\mathbf{U}(\omega_n)$. We estimate the plant transfer function $\mathbf{G}(\omega_n)$ through the simple relationship

$$\mathbf{G}(\omega_n) = \mathbf{Y}(\omega_n)\mathbf{U}^{-1}(\omega_n). \quad (4.1)$$

An important indicator of the quality of the experiment is the condition number of the matrix $\mathbf{U}(\omega_n)$, since the matrix must be inverted.

The signal-to-noise ratio of the diagnostic responses was excellent, due to the continuous excitation of the system and the pulse length available on JT-60U. By excellent, we mean that the variance between two separate measurements of amplitude or phase was always less than the variation from frequency to frequency. The amplitude of the response at the first and second harmonics of the driven frequencies and at intermediate frequencies was verified to be negligible compared with the amplitudes at the driven frequencies, confirming the linearity of the response and the adequacy of the signal-to-noise ratio. The result of these first experiments is a set of transfer functions (the complex responses of amplitude and phase) at the five frequencies, between the five PF coil-set voltages and all other signals. RZIP models these responses, so model and experiment can be directly compared.

5. Optimization of the plasmaless model

The results of the experiment and the predictions from RZIP were found to be in excellent but imperfect agreement. We therefore attempted to use the information in the experiments to improve the fit of the model to the data. This was done by quantifying the disagreement between the model and the data using a cost function and changing certain model parameters to minimize the disagreement.

Another optimization approach had been developed on TCV [8] which applied techniques of triangulation to find the globally minimal set of changes to all the physical properties of the TCV model to reduce the discrepancies of the measurements at very low frequency. This was not considered possible on JT-60U because the coil sets contain multiple coils and allowing all coils, diagnostics and gains to be variable would lead to an ill-defined optimum with only five different voltage inputs. Although the alternative TCV approach led to a simple perturbation of the device geometry (i.e. the machine implied by the modifications is still a tokamak), the method we have used on JT-60U does not necessarily lead to a physically consistent solution. However the very small modifications made, principally to the PF coil-set parameters, make us confident that the final model is extremely close to the real device. Note also that we have chosen not to modify the positions of the flux-loops and PF probes, since these had been established very precisely and since the DC calibration described in section 4.1 showed satisfactory agreement.

5.1. Choice of cost function

There is no universally applicable cost function because the data are noisy and the best cost function for model optimization for control purposes may be less suitable for investigation of the physical assumptions of the model.

There are two sets of data that can be used for comparison with the model (the transfer functions and the direct diagnostic measurements) and two approaches to quantifying the model-data error (the maximum model error, which we refer to as the H_∞ model error, and the root mean square of the error, referred to as the χ^2 error). The input to the cost function is the set of differences (ε) between the experimental data and the model predictions. The differences must be normalized to take into consideration different quantities and orders of magnitude in the measurements. Quantifying the agreement in terms of the χ^2 is a popular metric while minimizing the H_∞ model error has a specific link to robust controller design. The permutations permit four possible approaches, described in more detail in appendix A. Once this choice of cost functions is made, the optimization can be handled in several ways. In principle, it can be fully automated, but we found that many parameters showed almost no influence on the cost function. Rather than make large changes to these parameters to reduce the cost function by small amounts, we chose to leave the model parameters close to their nominal values unless there were clear indications to modify the values. In this way the relative ordering of the physical effects was respected, and the model parameters remained at their nominal values unless there was significant reason to change them.

5.2. Single-parameter scans of the plasmaless data

A single-parameter scan of selected electromagnetic model parameters was made in a specific order and those parameters which appeared to generate the largest improvement to the cost functions were manually modified (a visual gradient descent).

The order selected was:

- the self-inductances of the PF coil sets,
- the PF coil resistances,
- the mutual inductances of the PF coil sets,
- the vessel eigenmode description.

Figure 2 shows the single-parameter scans of the four cost functions for all the elements of the 5×5 PF coil-set mutual inductance matrix (maintaining symmetry) and for each of the five PF coil-set resistances. The self-inductances and resistances were scanned over $\pm 12\%$ of their central values. The mutual inductances were scanned over ± 0.12 in the coupling coefficient $k_{ij} = M_{ij}/\sqrt{M_{ii}M_{jj}}$. The central values in the figure are those of the tuned model and so the cost functions are fairly centred around zero variation. Simply projecting the single-parameter scans was justified by inspecting several parameter pairs and finding little correlation in the error surface, as well as converging to the same solution after different attempts. This complicated figure demonstrates three important features which we will consider in turn:

- There is a considerable difference in the sensitivity of the four cost functions to variations in the different parameters.
- There is a considerable difference in the behaviour of the cost functions themselves, creating difficulty defining a unique optimum.
- the level of consistency between cost functions for the inductance calculations is greater than that for the resistances.

The values of the minima for the two χ^2 cost functions are consistent to within 1 or 2% for all of the inductance matrix elements. This 1–2% precision indicates quantitatively the potential of this technique. The H_∞ cost functions derived

from the measurements (blue ‘○’) and from the transfer functions (green diamonds) do not always agree with each other so well, do not always agree with the χ^2 minima and are much broader than the χ^2 minima for some cases. This difference of behaviour led to some confusion during the initial optimization. This is now understood in terms of the presence of outliers in the set of model–experiment differences. The H_∞ cost function is dominated by the worst agreement between the model and the data. Using this cost function to optimize the model will necessarily move the model to minimize this maximum difference and the optimum is therefore dominated by any outliers, whether valid or invalid. The minimum is also less well defined since making the greatest number of points slightly better or slightly worse has no effect on this cost function. As such, this cost function is excessively conservative for our physics purpose. The χ^2 cost functions are also sensitive to the outliers, but their effect is minimized by the large number of data points showing good agreement. The broader H_∞ cost function also implies that the exact tuning of the model is less critical in the presence of statistical disturbances.

The cost function based on the direct measurements is better behaved than that based on the transfer function. The latter is found via an inversion of the voltage matrix U in equation (4.1), which leads us to conclude that an experiment designed with a diagonal matrix U should show no difference between the two cost functions and would be optimal. Our result implies that if the matrix U is not diagonal, then the effect of outliers on G will tend to be worsened. In the present plasmaless experiment, the condition numbers of the 5×5 $U(\omega_n)$ matrices for the five driving frequencies were 8.7, 8.1, 5.5, 7, 16, in order of increasing frequency. These numbers are assumed to be low enough to avoid excessive noise propagation




Figure 2. Single-parameters scans of the cost functions for different elements of the PF coil set mutual matrix M (only the upper half of the symmetric M matrix is shown) with the 5 PF coil resistances P (equation (2.1) defines these matrices). The cost functions χ^2 shown are for measurements (red \times), transfer functions (magenta $+$), H (measurements) (blue \circ) and H (transfer functions) (green \diamond). The horizontal scale is the relative modification of each individual parameter. The vertical scale is the relative change to the cost function from its minimum in each scan.

in the inversion of the $U(\omega_n)$ matrices when evaluating the transfer functions. The coupling between the coils was due to the presence of current feedback control loops on each PF coil set. These acted to oppose the currents induced at the frequency of the voltages applied to other coil sets.

When varying the parameters to optimize the model, we chose to centre the χ^2 cost functions. The tuned model is as such a physics-oriented model on the basis of the previous argument, searching for an optimum which represents the whole data set. Figure 2 illustrates that the H_∞ model is not identical to this best fit model. For most parameters, the difference is small, but for some it is significant. The observation that the tuning required to optimize the χ^2 cost function is much closer to the nominal model also suggests that the H_∞ tuning is being biased by inadequacies in the data. We conclude that the choice of the best model for physics purposes should be based on the χ^2 cost function, but that the resulting model is not necessarily the best model for robust controller design. This result also implies that the experiments required to derive a best worst-case error model are in fact more delicate to perform, since any measurement with excessive noise will punish that cost function but will be averaged over all available measurements in the 'physics' model. When deriving an optimal model for controller design, the whole set of available data should not be used. Rather, only the input-output transfer functions to the controlled variables should be considered, since only these particular transfer functions are important to the controller.

The experiments performed did not include frequency components low enough to allow the detailed resolution of the coil resistances, resulting in a difference in precision between the M -matrix elements and the P -matrix elements. The coils behaved almost inductively for all chosen frequencies, leading to flat minima for the P -matrix elements, even with the 10-times expanded vertical scale of the resistance cost functions in figure 2. In retrospect, for identification purposes, the frequency range should have covered a lower range. However the frequencies covered are those most important for controller design.

5.3. Results of the optimization

It was found that the nominal model fits the data better if the coil set self-inductances are modified by +23%, -4%, +2%, +1% and +24% for PF coil sets D, F, T, V and H, respectively. The two large positive corrections are understood as correcting for significant external coil impedances in the cables, power supplies and connections for the D and H coils and the remaining corrections can be attributed to small errors in the DC calibration.

The optimal PF coil-set resistances were found to be quite different from the values calculated by the DC calibration. However as mentioned before, the DC results were sensitive to offsets in voltage measurements and the AC tuning shows a flat minimum.

Only some values in the PF coil-set mutual inductance matrix, involving the H coil set, needed to be adjusted to show an improved fit to the experimental data as most of the coil sets are only weakly coupled by design. The largest change to a coupling coefficient was 2.7%, between the D and H coils. Those mutual inductances were small in the nominal model

and so large relative modifications are not unreasonable. This series of successive tuning of three sets of parameters did not have to be iterated.

5.4. Finer tuning

We next checked whether the nominal poloidal probe angles could be tuned. The errors between model and experiment were checked for the lowest frequency (presumably least perturbed by the shell current uncertainties) as the angle of the probes were moved from their nominal values to minimize the error according to the χ^2 metric. Thirteen of the 39 probes used were uncorrected. Fourteen probes were corrected by 1°–2°, 11 probes were corrected by 3°–5° and one probe was adjusted by 8°. Only those probes obviously perpendicular to the field of a particular PF coil set were considered for correction, since for this scenario the correction is the least sensitive to gain errors.

Finally, the vessel eigenmodes were corrected to take into account any differences between the nominal filament model and the physical vacuum vessel. The approach was the same as for the PF coil impedances and the adjustments to the inductance and resistance were small, ranging from -5% to +10%.

5.5. Quality of the model

The experimental results comprise the set of the 5 frequency responses of 57 validated signals (15 fluxes, 19 tangential magnetic field probes, 18 normal magnetic field probes and five PF coil-set currents) for each of the five PF coil-set voltage inputs. Each response has an amplitude and a phase with respect to the driving voltage. Figure 3 shows a set of four experimental responses (+) and the modelled responses using the nominal model (black solid line), the model with the PF coil-set impedances adjusted (blue dashed line) and with the poloidal probe angles adjusted and the vessel eigenmodes adjusted (red dotted line). These four responses are chosen to divide all of the input-output responses into a small number of categories.

The top left response (the response of the tangential probe #3 to the T coil-set voltage) is the most common quality of result (160 out of a total of 285 responses), with the models barely differing from each other and in good agreement with the data by eye. This category corresponds to responses which are dominated by a direct coupling between the driven coil and the particular diagnostic. The small modifications which were made mostly to the mutual inductances and to the dynamic part of the magnetic model have no visible effect on these input-output responses.

The top right response (tangential probe #4 to T coil set) shows a case of poor agreement (12 out of 285 responses) in which the experimental data lie relatively far from the nominal and adjusted models, although the smoothness of the experimental data suggests that the measured responses are more accurate than the difference with respect to the model. These responses were considered as suspect and were rejected from the model tuning. More information would be required to identify the probable cause of these 12 systematic disagreements.

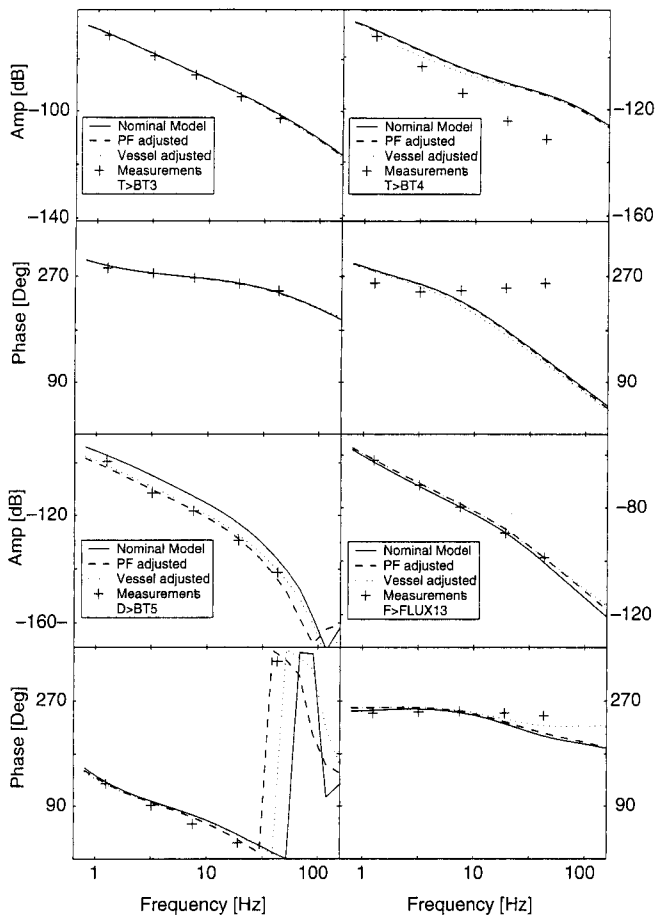


Figure 3. Bode plot showing the amplitude and phase of the plasmaless dynamic response of specific diagnostic signals to selected coil voltages (indicated Coil > signal). These four examples illustrate: no change (top left, 160 cases), poor agreement both before and after tuning (top right 12 cases), improvement by tuning the PF coil-set impedances (bottom left 103 cases) and improvement by tuning the vessel eigenmodes (bottom right 10 cases).

The lower left response (tangential probe #5 to D coil set) shows an improvement by eye to the agreement when adjusting the PF coil-set parameters and was found in 103 out of 285 cases. This category of responses can be explained by an important effect of the mutual inductance between the driven coil and other coils to which the particular diagnostic is sensitive. Modifying the mutual inductance will, in this case, affect the input–output response significantly.

Finally, the lower right response (flux loop #13 to F coil set) shows an improvement when the vessel eigenmodes are adjusted, representing only 10 out of the 285 cases. This category of responses is sensitive, at higher frequencies, to the vessel modelling and the changes to the vessel eigenmodes allows the driven flux to penetrate more or less rapidly to improve the agreement between the model and the experiment.

We conclude that by eye and in terms of the cost functions, over 30% of the transfer functions are improved by this optimization approach, which calibrates the external impedances of the PF coil set. The modifications to the vessel eigenmodes gave little visible improvement except for a few input–output cases.

Following this model calibration or model tuning method, we have obtained our optimized plasmaless electromagnetic model.

5.6. Possible reasons for discrepancies between the model and the data

The basic electromagnetic model of JT-60U was generated from the nominal parameters. Good agreement was obtained when initially comparing our experimental plasma responses with our modelling. However, some specific inconsistencies beyond the apparent error distribution led us to re-examine the assumptions made in the electromagnetic model.

One possible correction involves the assumption that a current would be uniformly distributed across the cross-section of a coil. The JT-60U PF coils are constructed from horizontal stacks of plates which are the width of the coil and there are non-centred cooling pipes in the coils. This implies that the current centroid is not necessarily centred in the coil cross-section. The nominal model was adjusted to take account of these effects. The skin effect in the PF conductors was also considered but discarded as small. If present it would have had a tendency to make the apparent coil inductance a function of the driving frequency, which was not found. The coupling of the PF coils to any metallic structures other than the vacuum vessel was also ignored. Other candidates for modifying the nominal electromagnetic model are the metallic casing of the toroidal field coils and any non-axisymmetric structures near the magnetic probes.

Other small corrections can be attributed to three dimensional eddy currents affecting the localized pick-up probes, to the distribution of the calibration gains, to the finite size of the magnetic probes and so on. All these effects can be lumped together as model imperfections since the model does not perfectly represent the data. The tuned plasmaless model is therefore one which represents reality better than the nominal model. If the plasmaless model is not tuned, we might be tempted to modify plasma response terms to compensate for these effects and draw incorrect conclusions.

6. Experiments with plasma

This section describes the measurement of the linearized plasma equilibrium response of two particular JT-60U equilibria with and without additional heating.

6.1. Ohmically heated plasma

The magnetic flux surfaces of one Ohmic equilibrium are illustrated in figure 1 (black lines). The plasma conditions were: $I_p = 1.2$ MA, $\kappa = 1.33$, $\delta = 0.28$, $q_{95} = 3.9$, $\beta = 0.06$, $l_i = 1.2$, $L_p = 5.54 \mu\text{H}$ and $n_e = 2 \times 10^{19} \text{ m}^{-3}$. The identification frequencies were chosen to be 3, 7, 16, 35 and 80 Hz. We increased the frequencies from those used in the plasmaless experiment simply because we expected the region of interest of the transfer function in the presence of a plasma to be at the highest frequencies and we had previously found adequate signal-to-noise at 43 Hz. However, as mentioned in section 5, the frequency range would ideally have been extended lower as well. Rather than modulate the power supply demand signals, the reference signals for the five control parameters used on JT-60U were modulated. These parameters are: the plasma current (I_p), the major radius (R), the vertical position (Z), the triangularity (δ) and the X-point height (X_p). The amplitudes of the control parameter variations were chosen

to provide a significant signal on the measurements, without excessively perturbing the operation, and chosen to reduce with frequency to avoid excessive voltage demand signals. Typical values for the control parameter excursions were: 10–50 kA for I_p , 2–5 cm for R , 2–5 cm for Z , 0.05–0.1 for δ and 3–5 cm for X_p . These parameters are estimated in real time by a set of non-linear expressions optimized by regression from an equilibrium database. For this work, during which the plasma was modulated in quasi-stationary conditions, a locally linearized form of these estimator expressions was used. This linearized estimator was successfully validated against the experimental values. The data required were obtained in JT-60U discharges E35009 and E35023.

For the five frequencies, the condition numbers of the five $U(\omega_n)$ matrices were 31, 21, 9.7, 13, 44, somewhat higher than for the plasmaless experiments but assumed to be still sufficiently low for the inversion of the matrices.

The resulting experimental response measurement matrix $G(\omega_n)$ is again large, including the complex response of all the magnetic probes, flux loops, poloidal currents and the linearized estimators of the control parameters, a total of 62 variables. These responses can be directly compared with the responses predicted by different models, as in the plasmaless case.

6.2. NBI heated plasma

Two more sets of experiments were carried out in the presence of 12 MW of NBI heating (discharges E35561, E35574). The aim was to determine the effects of decreasing the edge resistance of the plasma and possibly increasing the magnitude of the currents driven by the control transients. The plasma equilibrium was similar ($I_p = 1.19$ MA, $q_{95} = 2.9$, $\kappa = 1.44$, $\delta = 0.23$, $l_i = 1.01$, $\beta_p = 0.41$, $L_p = 5.34 \mu\text{H}$) and the experimental technique was the same. The increase in the electron temperature at the edge was smaller than hoped for since the discharge remained in L-mode with only $T_e = 2.5$ keV on axis, although the value of β_p was significantly greater. The result cannot therefore be considered definitive concerning the importance of edge conductivity. However, the experiments were analysed identically and provide substantial confirmation of the results of the purely Ohmic discharge.

7. Optimization of the plasma model

In this section the techniques of section 5 are used to optimize the RZIP plasma model with respect to the two JT-60U equilibria described above.

7.1. Comparison of nominal models and experimental data

The RZIP models were created for the Ohmic and NBI heated data and compared with the experimental responses for the probes. The modelled growth rate of the vertical instability for these plasmas is about 3 s^{-1} . The PF coils have a significant stabilizing influence in JT-60U, since the growth rate obtained by only considering the stabilization by the currents induced in the vacuum vessel is about 60 s^{-1} . The agreement between the model and the experimental data for the 285 transfer functions to the direct measurements can be divided into the four classes

illustrated in figure 4. This figure shows the amplitude and phase of the plasmaless data, Ohmic data and NBI data, and their respective models generated with the tuned plasmaless model and assuming zero plasma resistance.

In the first case (class A, top left) the experimental data with and without plasma are indistinguishable from one another and the model predictions coincide with each other and the data. This class of measurement is generally dominated by a particular coil-set current and the plasma barely perturbs the measurement. The data will apparently agree well with all plasma models, as long as the plasmaless model is accurate.

In the second case (class B, bottom right), the plasma data and the plasmaless data are very different in both amplitude and phase, and the plasma and plasmaless models agree well with their respective data. These measurements are sensitive to the plasma model.

The third category (class C, top right) is one in which the plasmaless model is reasonably close to the data, but the plasma model is not as close to the plasma data as in the good cases.

The fourth category (class D, bottom right) corresponds to data in which the plasmaless data and plasma data are similar, as are their respective models, but show serious disagreement. These cases are candidates for examining the calibration and geometrical data of the diagnostics themselves, or indicate a

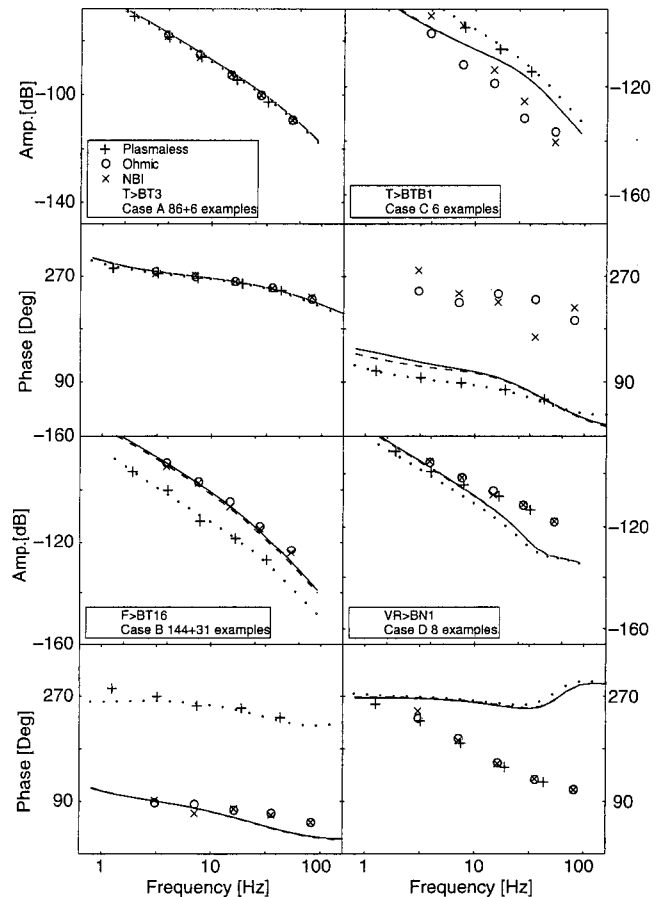


Figure 4. Comparison between the nominal RZIP plasma models and the experimental data for 4 representative cases. The amplitude and phase of the tuned plasmaless models and of the data are black (dotted line and +). The tuned Ohmic model and the Ohmic data are blue (solid line and O) and the NBI data and model are red (dashed line and x).

serious shortcoming of the plasmaless model. The fact that only a few of these cases were found, and that they appeared to be obvious outliers compared with the general level of agreement leads us to suspect the diagnostic rather than the electromagnetic model of the tokamak. There were no data in which the raw measurements differed significantly between the Ohmic and NBI plasmas and indeed the Ohmic and NBI models were always relatively close.

Inspecting the total data, we can approximately divide all the measurements into these four classes, bearing in mind that the transitions are essentially continuous between them. The result of this inspection is:

- Class A: 86 good cases, 6 more noisy cases.
- Class B: 144 good cases, 31 more noisy cases.
- Class C: only 6 cases.
- Class D: only 8 cases.

7.2. Optimizing the plasma model

Next we evaluated the parameter sensitivity of the plasma-related matrix elements of the RZIP model for two plasma discharges. We chose to adjust only those transfer functions which had been seen by eye to vary significantly in the presence of plasma and which also showed good agreement in the plasmaless case. This procedure avoided using the plasma terms in the M and P matrices to correct any residual insufficiencies in the plasmaless model. We followed the same approach as for the plasmaless model and figure 5 shows the variation of the two χ^2 cost functions for the Ohmic discharge after approximate tuning for different plasma-related matrix elements of the M and P matrices of equation (2.13). The cost functions are much less well defined than in the plasmaless case, demonstrated by the range of the horizontal axis scale (relative correction) changing from 0.5 to 1.5 and the reduced variation of the χ^2 measure in spite of the axis compression. The H_∞ model showed the same behaviour. The nominal model modifications made to obtain this figure are detailed in table 1. The subscript z refers to the state $z(I_p^0)$, R to $R(I_p^0)$, I to I_p , and the PF coil name (D, F, T etc) to the coil current state. So for example M_{II} refers to the bottom right-hand element of the M matrix in equation (2.13).

The radial field second derivative, M_{zz} , only had to be slightly corrected, taking the open growth rate of this equilibrium from 2.82 s for the nominal model to 3.3 s for the tuned model. The curvature of the cost function implies a localization of the optimum to better than $\pm 20\%$. The plasma inductance, M_{II} is more precisely determined, $\pm 3\%$, and unchanged with respect to the nominal model. The radial position parameter, M_{RR} , had to be increased by about 15%. The $z-I_p$ coupling, M_{zI} is predicted to be zero in the model and any non-zero value worsened the agreement. The $R-z$ coupling, M_{zR} , did not have to be modified. The mutual inductances between the PF coil sets and the plasma current, $M_{I,DFTVH}$, were accurately estimated by the model and did not have to be tuned. The plasma resistance, P_{II} , was estimated to be $8.8 \mu\Omega$ by the tuning method, compared with $0.37 \mu\Omega$ estimated from the experimental slow rate of poloidal flux consumption. The strong change suggests that care must be taken in making the common assumption of negligible plasma resistance. The radial derivative of the plasma resistance,




Figure 5. The variation of the model agreement cost functions obtained during single-parameter scan of the plasma-related matrix elements in the RZIP model, for the NBI heated discharge. Subscript z refers to $z(I_p^0)$, R is $R(I_p^0)$ and I is I_p . The cost functions χ^2 shown are for the measurements (red \times) and the transfer functions (magenta $+$). The horizontal scale is the relative modification of each individual parameter. The vertical scale is the relative change to the cost function from its minimum in each scan.

Table 1. Tuning of the most important plasma-related elements of the RZIP model.

Matrix elements	Significance	Nominal value/tuned value
M_{zz}	Vertical instability coefficient	$5.1 \times 10^{-7} / 5.6 \times 10^{-7}$
M_{RR}	Radial position coefficient	$-7.6 \times 10^{-7} / -9.2 \times 10^{-7}$
M_{II}	Plasma self-inductance	$7.46 \times 10^{-6} / \text{unchanged}$
M_{RI}	$R-I_p$ coupling	$2.0 \times 10^{-6} / 1.810^{-6}$
M_{zI}	$z-I_p$ coupling	$0 / \text{unchanged}$
M_{zR}	$R-z$ coupling	$2.4 \times 10^{-8} / \text{unchanged}$
$M_{I,DFTVH}$	Mutual inductances between PF currents and plasma	unchanged
P_{II}	Plasma resistance	$0.37 \mu\Omega / 8.8 \mu\Omega$
P_{RI}	Radial derivative of the plasma resistance	$0.15 \mu\Omega / 0.03 \mu\Omega$

P_{RI} , is non-zero from the tuning, but remains small, around $0.03 \mu\Omega/m$, compared with the predicted value for a rigid distribution shift if the plasma area and temperature remain constant, suggesting that this is not a suitable assumption.

The significant difference between the measured effective plasma resistance during the AC stimulation experiments and the effective DC plasma resistance from the rate of poloidal flux consumption is in agreement with a simple skin-effect model. The current profile was approximated as 20 concentric rings and the local conductivity was assumed proportional to the measured temperature profile to the power of 1.5. The response of the plasma current to modulations of the surface voltage is then explored as a function of frequency. The effective plasma inductance varies by only $\pm 2.5\%$ over the frequencies used in the experiment and remains close to the ‘slow’ inductance. However, the effective plasma resistance varied by a factor of 10, from 6 to $58 \mu\Omega$, in agreement with the tuned value of about $8.8 \mu\Omega$, compared with the ‘slow’ value of $0.37 \mu\Omega$. The RZIP model cannot be readily modified, by virtue of its structure, to take into consideration the diffusive plasma skin effect and an effective resistance is the best approximation, valid over a range of frequencies of interest.

Those individual measurements which are relatively insensitive to the presence of the plasma do contain valid information on the plasma parameters, but are dominated by the vacuum field of the PF coil-set currents. Instead of looking at the flux probe and magnetic field probe measurements, we can examine the control parameters to see if the separation between plasmaless or plasma models and data is increased.

8. Closed-loop comparisons

8.1. Closed-loop simulation

It is interesting to see whether the validated open-loop tokamak model, when combined with a model of the controller, can successfully simulate the closed-loop response in the time-domain. The controller used in the flat-top regime controls the vertical instability, and is designed to track the five parameters, R , Z , I_p , X_p and δ . These parameters are determined in real-time by the state estimator from a combination of normal and tangential poloidal magnetic field probe measurements, the D and T coil currents and the plasma current. For the closed-loop simulation we constructed an RZIP model linearized about the flat-top of Ohmic discharge E35023, at 3.9 s. The estimator was linearized about the same point and incorporated into the C matrix. The optimal grey-box model was used.

A simulation of shot E34993 was made, this shot being an Ohmic shot similar to the others but not used in the identification experiments or the grey-box modelling. As such it is a completely independent data set, suitable for comparison. The equilibrium values of the coil voltages, coil currents, and control variables, were defined as the average values between 3.7–3.9 s, before stimulation of the reference signal began.

Figure 6 shows schematically how the simulated and measured data were generated. Figure 7 shows the results of the simulation compared with the experimental data. The reference signal is shown as a thin red line, the actual experimental response as a thick green line, and the simulation as a noisier thin black line. Linear trends and offsets have been removed for clarity, although they were small. The closed-loop simulation of the estimator parameters response is extremely good. The response of the coil voltages and currents are similarly well modelled.

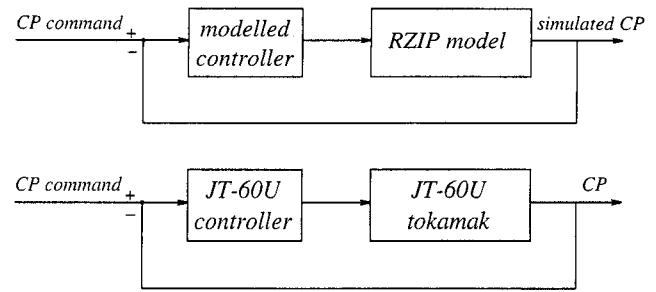


Figure 6. Block diagram of how the real and modelled data were generated.




Figure 7. A simulation of shot E34993. The thin red line is the command signal for the control parameter, the thick green line is the experimental response and the thin black line is the simulated response. Offsets and linear trends have been removed.

8.2. Controller tuning

Given the validity of the closed-loop simulation, we were able to perform control studies based on this model. To check the efficacy of the nominal controller, a simulation was made using square-wave excitation in the control variable reference waveforms. Some cross-coupling between different control variables was observed, which was ameliorated by simply adjusting the off-diagonal terms in the controller matrices, which are zero in the presently used JT-60U controller. The results of this exercise are shown in figure 8. The undesired transient cross-coupling has been significantly reduced for some cases, especially the influence of the triangularity change on the X-point height and plasma current, as well as the plasma




Figure 8. The simulated controller and the altered controller. The Control Parameter command is the thin red line, the response of the nominal controller is the thick green line and the response of the altered controller is the thin black line.

current change on the X-point height and the vertical position. None of the cross-couplings was worsened.

9. Estimating structured model uncertainty for the plasma equilibrium response model

One aim of developing an accurate plasma equilibrium response model is to provide a platform for optimal robust feedback controller design. Modern controller design methods provide optimal performance while guaranteeing stability over a prescribed range of uncertainties in the underlying model. Specifying the model uncertainty is therefore an essential part of the model based controller design process [9]. It is possible to quantify the model uncertainty in two ways, as unstructured uncertainty or as structured uncertainty. The former approach, taken in the H_∞ robust controller design process, is described in detail in appendix A, and essentially takes the maximum error of the transfer function in the frequency domain.

From physical considerations we can estimate the accuracy of elements in the M matrix of the dynamic model. The A and B state space matrices are not simply related to the underlying circuit equations (equation (2.1)), due to the required inversion of the M matrix. This inversion effectively propagates any uncertainties in the M matrix (equation (2.3)) throughout the A and B matrices, and therefore through the transfer function.

Table 2. Estimates of appropriate model coupling coefficient uncertainties. The uncertainty is defined by the relative change in the coefficient required to make the cost function change by 20%.

Equation element	Uncertainty
PF coil self-inductance	2%
PF coil mutual inductance coupling coefficient (absolute change in the coupling)	0.006
PF coil resistance	20%
Vessel eigenmode inductance, resistance	10%
PF to vessel mutual inductance (absolute change in the coupling)	0.05
Plasma inductance	10%
PF/vessel to plasma mutual inductance	10%
Vertical field weighted curvature	25%
Shafranov factor (Bv/I_p term)	25%
Radius-current coupling	50%
Plasma resistance (with respect to an estimate considering skin effect)	100%
Radial derivative of plasma resistance	200%
Radial-to-vertical motion coupling	50%

The work presented in chapters 6 and 8 gives us new insight into assessing structured model uncertainty. In fact, the χ^2 variation as different elements of the underlying model structure are varied, figures 2 and 5, provide us with exactly what is needed. The second derivative of the χ^2 variation tells us how well our set of experimental measurements can tie down particular coefficients, under realistic conditions. Specifically, comparison of figures 2 and 5 tells us how the accuracy of those coefficients which are independent of the plasma compares with the accuracy of the plasma-related coefficients. Inspection of these two figures reveals that assigning a general model uncertainty to the determining matrix coefficients would underestimate our ability to estimate the plasmaless coefficients and would seriously overestimate the precision of the plasma-related coefficients. Using this new information could, in theory, permit the design of a feedback controller which does not need to assume that the PF coil inductances, for example, are as badly known as the vertical instability growth rate, which can evolve significantly during a discharge, especially during equilibrium perturbations.

The model uncertainty estimated in this way is independent of the presence of a supporting experiment and can be used to investigate the effectiveness of a proposed diagnostic set in determining the model. Assuming the accuracy of the nominal model, we can calculate the responses for an ideal experiment, and then see how well defined the cost function minima are. This gives an understanding of how strongly the various model parameters affect the diagnostic measurements.

As support for these qualitative remarks, we have inspected all the groups of coefficients in the determining matrices and suggest an approximation of their uncertainties, shown in table 2. These suggestions take into account the present experiment's imprecision in estimating the coil resistance, for example, and include the possibility of spatial inhomogeneity for the plasma terms. The important point is the difference between the assumed precision of the structure-related parameters and the precision of the plasma parameters.

These estimates can be directly used for studying the effect of the uncertainties in the determining equation coefficients on the closed-loop feedback performance, which is an approach presently under study for the ITER tokamak.

10. Conclusions

This paper has extended previous work on linearized model development and comparison with experiment by applying to the JT-60U tokamak a systematic procedure to identify the plasma equilibrium response to the PF coil voltages. The required response was predicted with a high accuracy by a state space model derived from first principles. The original TCV work on experimentally identifying the plasma equilibrium response has been repeated on the larger and hotter JT-60U tokamak with a quite different PF coil system. The identified model has been compared with an improved derivation of the rigid current displacement model, which respects the symmetry of the determining equations. The modelling approach used, standard in analytical mechanics, is applicable to all linearized plasma equilibrium response models and has a structure which is independent of the physical assumptions made to describe the plasma response. Thus all linearized plasma equilibrium response models that have the same choice of states are members of the same model class, but have differing coefficients. This implies that adjustment of these models to fit experimental data, giving a tuned or calibrated model, provides the tuned response of any of the other models with the same choice of states. Starting with a pure rigid current displacement model is therefore simple and effective.

The approach used for tuning the plasmaless model of JT-60U has proven to be powerful although care must be taken to define a suitable cost function. The results indicate an extremely wide range of sensitivity to variations in different coupling coefficients in the plasmaless model. The same approach was used for tuning the plasma terms to agree with the experimental results for Ohmic and NBI plasmas, again obtaining an estimate of the precision obtainable. The variation in model sensitivity to the model parameters is important when defining the assumed uncertainties in a model used for feedback controller design and is the subject of current work on the ITER design.

The assumption of unstructured uncertainty results in an excessively conservative cost function for some terms, when compared with the demonstrated accuracy of the model. Other terms, particularly the plasma-dependent terms, should be considered as less certain.

Understanding the incorporation of experimental findings into a better description of the system uncertainties will be the subject of future work. It is tempting to speculate that the model optimization approach described in this paper would be especially advantageous when applied to the less well-modelled case of varying-saturation iron cored tokamaks such as JET.

The creation of a tuned plasma response model will allow design of tokamak plasma controllers with improved performance and robustness characteristics, for instance elimination of cross-coupling between certain plasma control parameters.

Acknowledgments

The authors wish to acknowledge the support of the Swiss National Science Foundation (JBL), the British Council and

JISTEC (AS), the UK EPSRC (DJNL, JPW and AS) and JAERI for partly funding this work. They are grateful for the technical assistance of the JT-60U tokamak divisions at Naka and to the Naka Fusion Research Establishment management for making this work possible. JBL and AS are grateful to the JAERI staff for the hospitality extended during their stay at JT-60U. Useful discussions with A. Portone of the ITER JCT at Naka led to us revisiting aspects of the previously published version of the RZIP model.

Appendix A. Estimating modelling uncertainty

Appendix A.1. The choice of cost function

Parametric or grey-box modelling allows certain model parameters to be adjusted until the discrepancy between observed and predicted data is minimized. In this sense the resulting grey-box model is optimal. Of course it becomes necessary to quantify the model–data discrepancy (by means of a cost function) so that minimization is possible. The appropriate choice of cost function will depend on the intended model application.

For our particular case we are considering a system that has outputs measured in many different units at different magnitudes, from mT to MA. Consequently it is necessary to define a cost function that is equally responsive to all outputs regardless of unit. There are many possible ways of achieving this. Also the plant response varies in magnitude significantly over the experimental frequency range. A suitable cost function should equally weight the response at all frequencies (given a sufficiently good signal-to-noise ratio) as different significant physical effects are important at different frequencies.

Ultimately there are two intended applications for the grey-box model. Firstly we intend to use it to refine our understanding of the importance of various physical effects in the tokamak. Secondly, we require that the model is suitable for robust controller design. Fitting the whole available data is useful for the former application; worst-case bounds on the model error are useful for the latter purpose.

These two applications suggest different curve-fitting approaches. A cost function based on a statistical approach is described in section A.2, and is considered suitable for the first application of physical modelling. A contrasting cost function is described in section A.3, suitable for robust controller design, employing a number of concepts from modern control theory.

Appendix A.2. Model error estimation for physical modelling

A first quantification of the model error is derived simply from the distances between the data points and the model approximation, for all available data points. These distances have to be normalized to an estimate of the likely error, which is difficult to provide. We consider that the root mean square of the difference between the model and data should provide a simple estimate of the uncertainty, avoiding the frequent cases where either the model or the measurement is close to zero. This weighted error has an upper bound of 2.0, corresponding

we wish to find \hat{G} such that $\|\Delta\|_\infty$ is minimized. If G has full column rank (almost everywhere), there exists a spectral factor $M \in \mathfrak{RH}_\infty$ with $M^{-1} \in \mathfrak{RH}_\infty$ such that $G^*G = MM^*$. Here we use the adjoint system $G^* = G^T(-j\omega)$.

For any \hat{G} , defining $\Delta = (\hat{G} - G)(G^*G)^{-1}G^*$ gives $\hat{G} = (I + \Delta)G$ and

$$\|\Delta\|_\infty = \|(G - \hat{G})M^{-1}M^*G^*\|_\infty \leq \|(G - \hat{G})M^{-1}\|_\infty,$$

so we have a conservative estimate of the model error $\|\Delta\|_\infty$.

We calculate M simply via the singular value decomposition of G^*G .

$$G^*G = Y\Sigma V^*,$$

with Y, V unitary, and since G^*G is symmetric, $Y = V$.

Thus we can write (since Y is unitary and Σ is positive definite)

$$\begin{aligned} G^*G &= Y\Sigma Y^* \\ &= Y\Sigma^{1/2}\Sigma^{1/2}Y^* \\ &= Y\Sigma^{1/2}Y^*Y\Sigma^{1/2}Y^*. \end{aligned}$$

Comparison with $G^*G = MM^*$ gives $M = Y\Sigma^{1/2}Y^*$.

Thus a suitable choice of cost function for robust controller design, that will not be biased towards any particular frequency range, is given by

$$\|\Delta\|_\infty \leq \|(G - \hat{G})M^{-1}\|_\infty.$$

Output scaling

Since the plant under consideration has outputs in different units, we introduce a scaling factor S on G such that each element of SG is of the same order.

We calculate S by

$$S_{ii}(j\omega) = \left(\frac{1}{n} \sum_{j=1}^n G_{ij}^2(j\omega) \right)^{-2},$$

$$S_{ij} = 0 \quad i \neq j$$

$$G(j\omega) = \begin{bmatrix} G_{11}(j\omega) & G_{12}(j\omega) & \dots & G_{1n}(j\omega) \\ G_{21}(j\omega) & & \ddots & G_{2n}(j\omega) \\ \vdots & & & \ddots \end{bmatrix}.$$

This calculation is predicated on the assumption that our plant inputs are all of the same order.

If output scaling is to be considered, the matrices SG and $S\hat{G}$ are to be used throughout instead of the plant and plant model respectively. If this is done the effect of the choice of output units on the model error estimate will be reduced.

References

- [1] Albanese R. and Villone F. 1998 *Nucl. Fusion* **38** 723
- [2] Hofmann F., Moret J.-M. and Ward D.J. *Nucl. Fusion* **38** 1767
- [3] Coutlis A., Bandyopadhyay I., Lister J.B., Vyas P., Albanese R., Limebeer D.J.N., Villone F. and Wainwright J.P. 1999 *Nucl. Fusion* **39** 663
- [4] Villone F., Vyas P., Lister J.B. and Albanese R. 1997 *Nucl. Fusion* **37** 1395
- [5] Vyas P., Villone F., Lister J.B. and Albanese R. 1998 *Nucl. Fusion* **38** 1043
- [6] Coutlis A., Bandyopadhyay I., Lister J.B., Vyas P., Albanese R., Limebeer D.J.N., Villone F. and Wainwright J.P. 2000 *IEEE Trans. Control Syst. Technol.* **8** 646
- [7] Kimura T., Kurihara K., Kawamata Y., Akiba K., Takahashi M. and Terakado T. 1997 *Fusion Technol.* **32** 404
- [8] Moret J.-M., Buhlmann F., Fasel D., Hofmann F. and Tonetti G. 1998 *Rev. Sci. Instrum.* **69** 2333
- [9] Ariola M., Pironti A. and Portone A. 1999 *Fusion Technol.* **36** 263
- [10] Wainwright J.P., Copsey D., Limebeer D.J.N., Haines M.G. and Portone A. 1997 *Nucl. Fusion* **37** 1697
- [11] Goldstein H. 1980 *Classical Mechanics* 2nd edn, section 2.5 (Reading, MA: Addison-Wesley)
- [12] Green M. and Limebeer D.J.N. 1995 *Linear Robust Control* (Englewood Cliff, NJ: Prentice-Hall)
- [13] Stewart G.W. 1973 *Introduction to Matrix Computations* (New York: Academic)

to the case where the model and data are equal in amplitude and π apart in phase. One cost function is therefore defined as:

$$\chi^2 = \left(\frac{2(d_i - m_i)^2}{d_i^2 + m_i^2} \right),$$

where d_i is a data point and m_i is the corresponding model prediction.

For the data and the model we can either take the raw data measurements, or the transfer functions. Before evaluating the cost function, outliers are removed, defined as points with a normalized error greater than 1.5.

Appendix A.3. Model error estimation for robust control

The infinity norm $\|\cdot\|_\infty$ is well understood and commonly employed in control literature [12] as a useful method of quantifying multi-input multi-output modelling uncertainty. Here we establish its relation to the singular value decomposition to demonstrate its significance as the worst case energy gain of a system. The infinity norm is then applied to the estimation of modelling uncertainty, with regard to its interpretation as the worst case energy gain.

The singular value decomposition

For any $m \times p$ complex matrix Q , there exist $m \times m$ and $p \times p$ unitary matrices Y and U such that

$$Q = Y \begin{bmatrix} \Sigma & 0 \\ 0 & 0 \end{bmatrix} U^*, \quad (\text{A.1})$$

where $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_r)$ with $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r > 0$, $r \leq \min(m, p)$. Expression 2.1 is the singular value decomposition of Q . A proof of existence can be found in [13].

Letting u_i and y_i be the rows and columns of U and Y , we can express Q in terms of the dyadic expansion

$$Q = \left(\sum_{i=1}^r \sigma_i y_i u_i^* \right).$$

Since U is unitary, $u_i^* u_j = \delta_{ij}$, u_j is mapped by Q into

$$Q u_j = \left(\sum_{i=1}^r \sigma_i y_i u_i^* \right) u_j = \sigma_j y_j.$$

We can therefore regard the matrix Q as a linear mapping from vector space C^p to C^m defined by

$$Q : C^p \rightarrow C^m : u \rightarrow Q u.$$

The set of columns of the singular-vector matrices U and Y respectively define orthogonal bases for the domain C^p and range C^m of Q . For this choice of bases, the mapping Q takes the j th basis vector u_j of C^p to a vector lying in the direction of the j th basis y_j of C^m . If we restrict Q to the one-dimensional complex subspace spanned by u_j , the corresponding singular value σ_j can be regarded as a gain factor for the restriction map $Q|_{u_j}$ [12].

The maximum singular value $\bar{\sigma}$ is clearly just

$$\bar{\sigma}(Q) = \sigma_1(Q).$$

The infinity norm as the worst case energy gain

Suppose we have the transfer function

$$y = G u, \quad G \in \mathbb{R} H_\infty^{+, p \times m}.$$

We can define the infinity norm of G as

$$\|G\|_\infty = \max_{u \neq 0} \frac{\|G u\|_2}{\|u\|_2},$$

where we use the standard $L^2[0, \infty]$ norm, where if $u(t)$ is a square integrable time function, then

$$\|u(t)\|_2 = \int_0^\infty u^*(t) u(t) dt < \infty$$

represents the energy in the signal $u(t)$. By Parseval's theorem we have

$$\|u(t)\|_2 = \|u(j\omega)\|_2.$$

Intuitively therefore the infinity norm is the worst case energy gain.

It is possible to show [12]

$$\|G\|_\infty = \sup_\omega \bar{\sigma}(G(j\omega)).$$

Model uncertainty

We can use these definitions to quantify and interpret an unstructured model error. We can represent the modelling uncertainty E in an additive way,

$$G(j\omega) - \hat{G}(j\omega) = E(j\omega),$$

where G is the 'true' system and \hat{G} is the model of the system. A useful measure of the severity of this error is the infinity norm, as it measures the worst case energy gain due to the modelling uncertainty:

$$\|E\|_\infty = \frac{\|(G - \hat{G})u\|_2}{\|u\|_2}. \quad (\text{A.2})$$

As such the infinity norm of the additive model error is a candidate choice of cost function for grey-box modelling, since it assumes no structure for the modelling uncertainty and treats the modelling uncertainty in a consistent and intuitive way. However, this representation tends to weight the low-frequency errors more heavily, since the higher gains at low frequency lead to larger absolute errors.

Alternatively a relative or multiplicative form can be used,

$$\hat{G}(j\omega) = (I + \Delta(j\omega))G(j\omega).$$

We quantify the model error with the infinity norm of Δ , which describes the magnitude of the relative model error in a worst case energy gain sense.

The relative form of model error is advantageous for our application since it is not biased towards errors at low-frequency.

Calculation of the model error

Given the model error representation

$$\hat{G}(j\omega) = (I + \Delta(j\omega))G(j\omega),$$