

On-Line Calculation of Feedforward Trajectories for Tokamak Plasma Shape Control*

M.L. Walker, R.D. Johnson, J.A. Leuer, B.G. Penaflor

Abstract — In this paper, we report on the implementation and experimental test of an algorithm for computing feedforward coil current trajectories that can produce an approximately correct evolution of plasma shape, using a nonlinear on-line optimization method that maintains coil currents far from their limits. The voltage trajectory that produces these currents can be used as the feedforward component of plasma shape control when combined with a multivariable linear feedback control algorithm.

I. INTRODUCTION

In the past few years, much work has been done at various institutions on application of multivariable linear controllers for tokamak plasma shape control, including several design studies and experimental implementations (see [1] for a survey). A great deal has been learned about shape control through these studies. Some of the practical constraints imposed by working tokamaks must now be addressed in order to make the benefits of multivariable control available to working experimental devices. Because of the large plant size typical for tokamak plasma shape control, developers are largely constrained to use of mature linear design tools for multivariable controllers. By contrast, for many of the nonlinear problems faced, we must appeal to rather immature or even yet to be developed nonlinear control techniques.

Multiple current and voltage constraints such as limits on coil current and a constraint on a sum of currents at DIII-D (see below) imply that the range of accessible plasma equilibria is constrained. Programmed attempts to reach equilibria outside these constraints can lead to problems such as exceeding coil current limits or saturation of actuators leading to loss of control. In this paper, we describe a nonlinear on-line optimization method for computing a feedforward coil current trajectory that can produce an approximately correct evolution of plasma shape. The voltage trajectory that produces these currents can be used as the feedforward component of shape control when combined with a multivariable linear control algorithm. This online optimization calculation is able to exploit the existence of real time reconstructed equilibria and the prediction accuracy of

plasma models linearized around these equilibria to generate nominal coil current trajectories on-line that can produce an approximately correct evolution of plasma shape while remaining far from current boundaries.

II. DIII-D OPERATIONAL ERROR CALCULATION

The isoflux control method, now in routine use on DIII-D, exploits the capability of the real time EFIT [2] plasma equilibrium reconstruction algorithm to calculate magnetic flux at specified locations within the tokamak vacuum vessel. Figure 1 illustrates a lower single null (LSN) plasma that was produced using isoflux control and




Fig. 1. Illustration of a lower single null (LSN) plasma equilibrium and the isoflux control points and X point grid used for calculation of shape error.

indicates quantities relevant to the control scheme. The real time EFIT algorithm calculates the value of the poloidal flux in the vicinity of the plasma boundary very accurately. Thus, the controlled parameters are the values of flux at prespecified control points along with the X-point R and Z positions. (The X-point is the location where magnetic field equals zero. The plasma of Fig. 1 is an LSN because the X-point, or field "null", is in the lower half of the vacuum vessel.) By requiring that the flux at each control point be equal to the same constant value, the control forces the same flux contour to pass through all of these control points. By choosing this

*This work supported by the U.S. Department of Energy under Contract DE-FC02-04ER54698.

M.L. Walker is with General Atomics, P.O. Box 85608, San Diego, CA 92186-5608, USA. phone: 858-455-2483; fax: 858-455-4156; e-mail: walker@fusion.gat.com.

R.D. Johnson is with General Atomics. e-mail: johnsonb@fusion.gat.com

J.A. Leuer is with General Atomics. e-mail: leuer@fusion.gat.com.

B.G. Penaflor is with General Atomics. e-mail: penaflor@fusion.gat.com.

constant value equal to the flux at the X-point, this flux contour must be the last closed flux surface or separatrix. The desired separatrix location is specified by selecting one of a large number of control points along each of several control segments. An X-point control grid is used to assist in calculating the X-point location by providing detailed flux and field information at a number of closely spaced points in the vicinity of the X-point.

Present DIII-D operations use the isoflux control method with proportional, integral, and derivative (PID) calculations operating on the control point flux and X-point R and Z errors [3]. The resulting signals are multiplied by a gain matrix to produce commands to pulse width modulated (chopper) power supplies on many of the plasma shaping coils. The gain matrix is sparse, so most individual shape errors are corrected through the application of only a small number (often one) of coil voltage changes. Control of the X-point requires coordinated action by the largest number (4) of shaping coils.

This approach usually provides good control. However, there are several control problems that have not been solved with this simple decoupled control approach. Some of these problems are presently handled by what might be called "operator in the loop adaptive control" but this operator tuning consumes valuable experimental time. Tendencies of certain shapes to cause coil currents to exceed their limits are handled manually between plasma discharges by adjusting shapes to decrease required current levels or by adding additional choppers to increase current capability for some coils. Variability in chopper gain from coil to coil and for different equilibria is handled by operator manual gain changes. Other control problems such as shape oscillation are also adapted to by manually retuning PID gains in the controller.

III. CURRENT AND VOLTAGE LIMITS

The topic of this paper is the problem of coil current limiting. The choice to control a large number of error signals (Fig. 1) causes the control problem to be "overdetermined" in the following sense. Roughly, it means that there are more control parameters than actuator degrees of freedom. There are 18 F-coils that shape the plasma; of these, up to 16 have power supplies attached. An additional constraint on a particular sum of currents reduces the number of actuator linear "degrees of freedom" to 15. This seems to match the set of 13 independent control points plus X point R and Z values shown in Fig. 1. However, this simple dimension count does not illustrate all of the constraints. There are strong current and voltage limit constraints that imply that the 15 control parameters cannot always be simultaneously minimized.

In fact, it is easy for an operator to choose a reference shape that is incompatible with device constraints, even in

steady state [3]. Even with a relatively small number of control parameters [4], the currents required to obtain zero error for those parameters at equilibrium are affected strongly by the choice of control point locations; a poor choice can create an ill conditioned map from coil currents to control parameters, which in turn leads to large required currents. Any linear controller would be unaware of the current limits and would try to obtain the required currents, thus driving coils to their limits. This is in fact what happens when a linear controller is executed in closed loop simulation with the model plant [5].

Since this problem is primarily caused by requesting a reference shape that is not physically realizable by the device, what seems to be an obvious "fix" to the problem is to specify only reference shapes that are compatible with all constraints of the device. Plans for future reactors such as ITER [6] have generally included such carefully planned programmed shape evolutions designed to minimize the danger of limiting currents. A necessary prerequisite for this approach is the ability to compute in advance reference equilibria (e.g., the plasma in Fig. 1) that are completely compatible with the device constraints. However, these equilibrium calculations are dependent on knowledge of the plasma internal current distribution that can change in an uncontrolled manner during a plasma discharge. They are also sensitive to the accuracy of the plasma and constraint models used. Thus, some real time methods are still desirable to handle cases where coil currents or voltages are near their limits. The ITER design, for example, includes a supervisory layer that is prepared to react to violations of certain constraints.

An effort has been underway for some time at DIII-D to construct model based multivariable controllers to address the problems currently seen in operational shape control. These controllers are presently based on models derived from previous experimentally obtained equilibria and are usually tested by "piggybacking" on experimental discharges intended for physics studies. It is often the case that equilibria in these discharges use currents already very near to current limits. This implies that there is little headroom for multivariable linear controllers, which, when tested, almost immediately sends coil currents into the limits. Reducing the number of controlled parameters can alleviate this problem, but cannot guarantee adherence to constraints under all conditions and, in general, is not a preferred solution.

In this paper, we address the specific problem of providing sufficient coil current headroom for multivariable controllers to be implemented. This problem is not strictly one of control, but is driven primarily by plasma equilibrium (steady state) considerations. Many shapes actually require currents outside of the current boundaries in order to be produced exactly [3],[4].

IV. APPROACH TO SOLUTION

The work described in this paper exploits several key ideas and observations.

(1) Linear models generated by linearizing around a plasma equilibrium are remarkably accurate, as long as the equilibrium doesn't change significantly [see (5) below]. For a fixed plasma equilibrium, Faraday's law and simple force balance allow derivation of the linearized model equations for poloidal field (PF) coil current dynamics with standard ohmic current drive

$$\begin{aligned} M_{cc}^* \dot{I}_c(t) + R_c I_c(t) + M_{cv}^* \dot{I}_v(t) + M_{cp}^* \dot{I}_p(t) &= V_c(t) \\ M_{vv}^* \dot{I}_v(t) + R_v I_v(t) + M_{vc}^* \dot{I}_c(t) + M_{vp}^* \dot{I}_p(t) &= 0 \\ L_p^* \dot{I}_p(t) + R_p I_p(t) + M_{pc}^* \dot{I}_c(t) + M_{pv}^* \dot{I}_v(t) &= 0 \end{aligned} \quad (1a)$$

and outputs

$$\begin{aligned} y(t) &= y_{eq} + C_{I_c} [I_c(t) - I_{c,eq}] + C_{I_v} [I_v(t) - I_{v,eq}] \\ &\quad + C_{I_p} [I_p(t) - I_{p,eq}] \quad , \end{aligned} \quad (1b)$$

where I_c , I_v , and I_p represent currents in PF coils, vessel (more generally, passive conductors), and plasma, respectively, V_c is the vector of voltages applied to the PF coils, M_{ab}^* are mutual inductance matrices for $a, b \in \{c, v, p\}$, modified to include the plasma response, and R_a represent resistance matrices. A sufficiently accurate model of plasma boundary shape response can be derived under the assumption of rigid (R and Z only) motion of the plasma current channel [7]. These matrices are uniquely determined by the device conductor geometry and the current density distribution in the plasma. The additional subscript "eq" denotes values at the plasma equilibrium from which the model equations are derived.

The concept of plasma equilibrium is different from an equilibrium of the dynamical system (1). In the linearization (1), the equations are written with current and voltage coordinates centered at zero [7], not relative to some dynamic equilibrium condition. In the plasma/tokamak system corresponding to this model, there does not exist such an equilibrium nor even a steady state operating point [7] since, using ohmic plasma current drive, some or all of the PF coil currents must be continuously ramped downward to induce voltage to drive the plasma current I_p .

In DIII-D, the dedicated ohmic heating coils (E-coils) provide nearly uniform flux over the vacuum region so that they have almost no effect on the plasma shape. For this reason, the plasma current I_p is controlled separately from the shape using the E-coils as actuators. Replacing $y(t)$ with $e(t)$, representing errors in the parameters used to control the plasma boundary at DIII-D — flux at discrete target points along the plasma boundary and field at a target

X-point location (Fig. 1) — and neglecting contributions to the errors due to currents in the E-coils and vessel, we can write (1b) as

$$e(t) = e_{eq} + G_{eF} [I_F(t) - I_{F,eq}] + G_{ep} [I_p(t) - I_{p,eq}] \quad . \quad (2)$$

Here G_{eF} and G_{ep} denote matrices mapping F-coil and plasma currents to control errors, where the calculation is done relative to the plasma equilibrium and I_F denotes F-coil currents. We neglect the effect of vessel current variation on shape since the vessel currents due to the ohmic drive are very small (hundreds to thousands of amperes) relative to the plasma current and the multi-turn equivalent currents in the F-coils (hundreds of thousands of amperes).

(2) If there are more actuator degrees of freedom than controlled parameters, there is freedom to modify the actuators without affecting the control parameters. In the case of more PF coils than parameters to be controlled, there exists a subspace of coil current perturbations away from the equilibrium current vector that will not affect the controlled plasma shape parameters. Coil current vectors in this "shape null-space" can be added to the equilibrium current vector to move it away from current limits.

(3) Information about the plasma equilibrium is available on-line, produced by the real time EFIT reconstruction.

(4) It is now feasible to generate linearized plasma models on-line from the on-line computed equilibria. Only the output equation needs to be regenerated for the calculations considered here. This has not been implemented in real time code yet.

(5) Under the rigid model assumption, there are only three ways in which a change to the plasma can occur that is significant enough to modify the accuracy of the linearized model: (1) change of current profile, (2) change in total plasma current I_p , or (3) change in shape. A significant change in any of these implies a redistribution of plasma current [7]. The current profile represents plasma current density as a function of radius, extending outward from the magnetic axis (Fig. 1). A change in this profile typically occurs on slow timescale, so sufficiently frequent re-linearizations can maintain good model accuracy. Changes in total plasma current or plasma shape are *predictable*, in the sense that if the control works reasonably well, we know in advance roughly what their values will be and therefore, when re-linearization should occur.

(6) Reference signals for plasma current and plasma shape are piecewise linear.

We exploit these ideas to define an on-line calculation of nominal feedforward trajectories of coil current and voltage that produce (approximately) the desired shape while simultaneously maintaining the nominal current trajectory far from current boundaries.

V. ALGORITHM DESCRIPTION

Figure 2 shows the intended use of the feedforward trajectory calculation in a shape control application. The top path shows calculation of the feedforward trajectories occurring at infrequent intervals (e.g. 100 ms). Between calculation times, the time varying feedforward voltage trajectories are added to a shape control feedback voltages δV every 250 μ s to produce the total voltage demand for the power supplies on the DIII-D F-coils. Efforts described here are focused on verifying that on-line trajectories, computed in the block labeled Current Trajectory Calculation using the output equation (2), can be used to produce approximately the correct shape in a feedforward manner. Once a feedforward trajectory for coil currents is defined, it is fairly straightforward to generate the required feedforward voltages if only piecewise linear current trajectories are used.

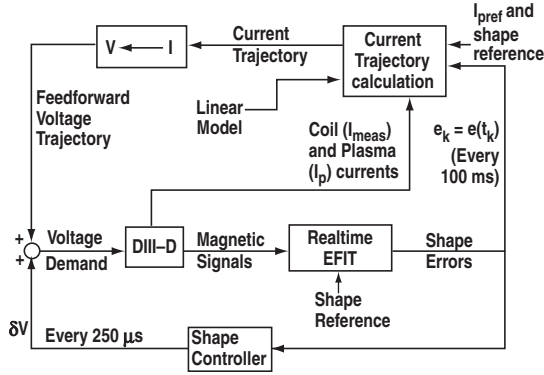


Fig. 2. Illustration of the role of on-line computed feedforward trajectories in plasma shape control.

In the following, we describe the constrained minimization algorithm used on-line to compute the trajectories. An earlier alternative method was described in [8]. In both methods, a calculation takes place at each time t_k to compute the feedforward trajectory between t_k and t_{k+1} . We use the coil and plasma currents as measured by diagnostic sensors and the error signals as calculated by the real time EFIT code to define the plasma equilibrium in (2), relative to which all calculations are done (Fig. 2). We assume that reference trajectories for plasma current I_p and shape are defined, i.e. that $I_{p,ref}(t)$ is specified, and that the desired evolution of control point locations and X-point location (Fig. 1) are given.

The method used here is a constrained optimization that seeks to find a nominal current vector I_{nom} that minimizes the error in shape, subject to the constraint that I_{nom} remains within some imposed current limits. Given an existing level of shape error $e_k = e(t_k)$ corresponding to a measured current vector I_{meas} , find the change in F-coil current ΔI_F that, when added to I_{meas} , will minimize the error at the time t_{k+1} , and such that $I_{nom} = I_{meas} + \Delta I_F$ remains within current limits. This is given by the solution of

$$\min_{\Delta I_F} \left\| W \left(e_k + G_{ep} \Delta I_p + G_{eF} \Delta I_F \right) \right\|^2 + \left\| \Delta I_F \right\|^2, \quad (1)$$

subject to

$$L - I_{meas} \leq \Delta I_F \leq U - I_{meas} \quad (2)$$

The values U and L are the imposed upper and lower current limits respectively on all 18 F-coils. The extra term $\|\Delta I_F\|^2$ with $0 < \varepsilon \ll 1$ is needed for regularization, i.e., to make the quadratic term full rank, since the number of errors is less than the number of currents. The weight W is used to account for the fact that different coils have different allowable current ranges. We use the additional term $G_{ep} \Delta I_p$, where $\Delta I_p = I_{p,ref}(t_{k+1}) - I_{p,ref}(t_k)$, to compensate for the change in flux and field due to any anticipated change in plasma current between the calculation times. The reference shape is specified through the DIII-D plasma control user interface as a sequence of control point and X-point locations, changes in which are compensated in the trajectory calculation by using the matrices G_{eF} and G_{ep} corresponding to the reference control locations at time t_{k+1} . Thus, calculations done at t_k are actually intended to reduce the error at time t_{k+1} . Changing the value of I_{nom} at each time t_k would impose a step disturbance on the control, so a method of piecewise linear interpolation is used instead (see next section).

By neglecting terms independent of ΔI_F , the objective functional may be written as

$$F = 2e_{est}^T W^T W G_{eF} \Delta I_F + \Delta I_F^T G_{eF}^T W^T W G_{eF} \Delta I_F + \Delta I_F^T \Delta I_F, \quad (3)$$

where the estimated error at t_{k+1} before correction by modified F-coil currents is given by

$$e_{est} = e_k + G_{ep} \Delta I_p \quad (4)$$

We apply a modification of a standard inequality constrained quadratic programming (IQP) algorithm designed for the problem:

$$\min_x \left(\frac{1}{2} x^T H x + f^T x \right), \text{ subject to } Ax + q \geq 0, Bx = 0, \quad (5)$$

with

$$x = \Delta I_F, H = G_{eF}^T W^T W G_{eF} + \mathbf{I}, f^T = e_{est}^T W^T W G_{eF},$$

$$A = \begin{bmatrix} \mathbf{I} \\ -\mathbf{I} \end{bmatrix}, q = \begin{bmatrix} I_{meas} - L \\ U - I_{meas} \end{bmatrix}, \quad (6)$$

and equality constraints $Bx = 0$ defined by the circuit wiring configuration of the F-coils, which is modified according to the type of plasma [e.g. lower single null (Fig. 1), double null, or upper single null] to be produced.

An alternate regularization term is $\|\Delta I_F - \Delta I_{center}\|^2$, which is intended to give preference to currents nearer to the center when the shape error is small. This regularization has not yet been implemented in the real time code.

This method is a significantly more complex calculation than that in [8] but it guarantees that the nominal current trajectory produced will have a well defined minimum distance (defined in advance by the operator) to hardware current limits for each F-coil, which can be relied on by any shape controller centered around this trajectory. In addition, it provides a means for compensating for programmed changes in shape or plasma current.

VI. IMPLEMENTATION DETAILS

The special structure of this problem was exploited to reduce the calculation time sufficiently for real time use. A total of 36 constraints are of the form $\pm I_i \geq q_i$, where currents I_1, I_2, \dots, I_{18} represent currents in the 18 F-coils. It is clear that at most 18 of these constraints can be active at any time. We use an active set method for IQP described in [10], which invokes an equality constrained quadratic programming (EQP) algorithm to operate with constraints active. To minimize $x^T G x / 2 + g^T x$, subject to $a_i^T x + b_i = 0$, $i = 1, 2, \dots, n_{active}$, where n_{active} is the number of active constraints, we use the EQP algorithm

$$\begin{aligned} Gx^* - A^T \lambda^* &= -g \\ -Ax^* &= b \end{aligned}$$

where A (whose rows are the a_i^T , $i = 1, 2, \dots, n_{active}$) consists of only active constraints, x^* is the solution vector, and λ^* is the Lagrange multiplier. This method allows exploitation of the fact that, when entering EQP, most constraints are of the form $I_i = \pm b_i$, so that a number of elements of x^* are already defined. By eliminating those variables first, a matrix equation is generated and solved that is of reasonable size (between 19 and 22 variables) independent of the number of active coil limit constraints. In addition to this special structure, it was not necessary to iterate the IQP to convergence, since three iterations were enough to ensure the calculated solution was feasible, i.e. satisfied all constraints, and produced a good reduction in error. The resulting algorithm executed in roughly 120 μs , less than the 250 μs sampling time routinely used for shape control. This time is actually much faster than needed, since the calculation is used only to produce a feedforward trajectory and thus occurs only infrequently during the plasma discharge.

The computed nominal currents are difficult to use as a nominal trajectory, because they are not continuous across successive calculations. A smoother piecewise linear trajectory $I_0(t)$ is produced on-line by interpolating between the smoothed value at the recompute time and the newly computed nominal value I_{nom} ,

$$I_0(t) = \lambda I_{nom}(t_k) + [1 - \lambda] I_0(t_k)$$

for $t_k \leq t < t_{k+1}$, where t_k is the time at which recalculation of I_{nom} takes place, $\lambda = m_k(t - t_k)$, $\alpha(\tau) = \tau$, $0 \leq \tau < 1$, and $\alpha(\tau) = 1$, $\tau \geq 1$, and the rate of change of current m_k [on the order of $1/(t_{k+1} - t_k)$] is chosen to avoid voltage saturation. The times $t_k, k = 1, 2, \dots, n$ at which to recompute $I_{nom}(t_k)$ are determined in off-line calculations after the desired evolution of the shape and plasma current for the discharge are programmed. They are chosen such that the nominal value I_{nom} is recomputed whenever the shape or plasma current reference changes significantly. Additional calculation times are presently added at intervals to have more frequent calculation points during algorithm testing. These additional calculation times can also serve the purpose of adjusting for changes in the equilibrium profiles, if the linear models are generated online from real time reconstructed equilibria. The initial value of $I_0(t_0)$ is computed to be the vector of currents closest to I_{meas} that can be maintained in steady state by the set of power supplies connected to the F-coils.

VII. RESULTS OF EXPERIMENTAL TESTS

This algorithm was implemented in the DIII-D plasma control system (PCS) and tested experimentally in a sequence of three plasma discharges. The test interval in these discharges extended from 2 to 4 s. The plasma current reference was constant during the test interval. Plasma shape requested only a 4 cm change in X-point radial position between 2 and 3 s, followed by a significant overall shape change request in the interval between 3 and 4 s. The nominal trajectory calculation was updated every 100 ms between 2 and 3 s, but updated more frequently after that because of the requested shape changes. In these experiments, seven control segments were used, $\epsilon = 1 \times 10^{-9}$, and $W = \text{diag}([5, 1, 5, 1, 5, 1, 1, 10, 10])$. In order to provide a means for direct evaluation of the effectiveness of shape compensation of the computed current trajectories, feedback control of current was used to force coil currents to closely follow the computed nominal trajectories. No feedback of shape was used between trajectory calculation times. A maximum of 500 A/s was imposed on the absolute value of coil current derivatives.

Figure 3 shows two representative shapes achieved during plasma discharge 122129, one before the test interval while the plasma was under full shape control and the other during the period between 2 and 4 s when only the feedforward trajectory method was used. The shapes are nearly identical to the eye, although differences can be seen in traces of error signals (Fig. 4). The shape is held approximately steady for the first second, when the shape reference is unchanging. From 3 to 4 s the reference shape is changing, and the




Fig. 3. Plasma shapes achieved in discharge 122129 before the test interval (1981.75 ms, solid) when under full shape control and during the middle of the test interval (2500 ms, dashed) when using only feedforward trajectory calculation.

trajectory calculation cannot follow the fast shape change, limited primarily by the maximum allowed current derivatives, but the shape does move in the correct direction.

A second test used lower current limits raised and upper limits reduced by an additional 700 A to force several current constraints to be active during the calculation. Figure 5 shows the result of this test in discharge 122131, compared with the result for the less constrained test of discharge 122129. Four of the 18 F-coil currents (2A, 3A, 9A, and 2B; see Fig. 1) were actively constrained by this new limit, which caused a relatively small distortion of the shape.

While we expected that the plasma shape should remain in the vicinity of reference shape, we did not expect it to be well controlled. There was no mechanism in the test algorithm for responding to disturbances or fast reference changes. Performance was actually better than expected, in the sense that the plasma shape was well maintained and followed fairly closely a slow change in reference for the X-point position.

Lack of test time precluded testing all facets of the algorithm. In particular, plasma current was not varied so the ability to compensate for an I_p change was not tested. Similarly, calculation of the proper feedforward voltages given the nominal current trajectory was not tested, although a previous simulation [8] has shown that this also works. More recently, successful feedback control experiments were conducted using the full scheme shown in Fig. 2.

VIII. CONCLUSIONS.

In this paper, we have described methods for computing on-line feedforward trajectories based on real time computed equilibria and linearized models derived from those equi-

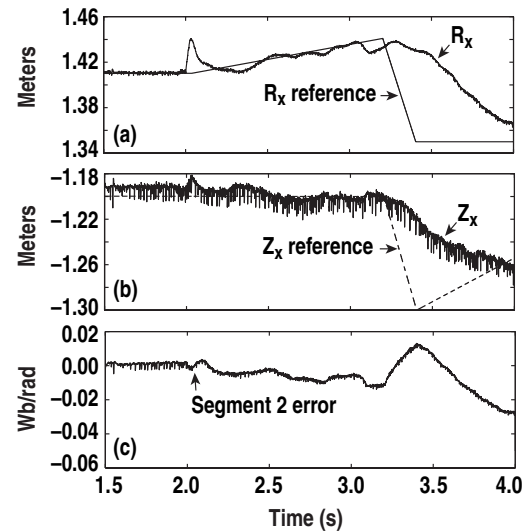


Fig. 4. Traces of signals corresponding to equilibria shown in Fig. 3. (a) and (b) show the reference and actual location of the X-point R and z position, respectively. (c) shows the flux error for one segment (number 2) of the 7 total control segments. (Only three signals are shown due to lack of space.)




Fig. 5. Shapes produced by feedforward trajectories with two different constraints at 2500 ms. Shot 122129 (solid) enforces lower (resp. upper) constraints 100 A above (resp. below) the actual current limits. Shot 122131 (dashed) enforces lower (upper) constraints 800 A above (below) the actual limits.

libria. These methods have been verified experimentally to provide a means to define a nominal trajectory that can be used as the coordinate origin for multivariable shape controllers and provide uniform headroom for control. This approach addresses the two problems of finding improved equilibrium currents for a given desired shape and adapting to nonlinear changes in the plasma. The use of nominal trajectories that provide good open-loop prediction of voltages required to produce the given shape also provides a means for reducing required gains for linear shape controllers.

This initial work has verified that the output equation (2) can be used to generate feedforward coil current trajectories on-line. Subsequent experiments successfully used this nominal trajectory calculation as part of a full multivariable control. In the present experimental tests, we used a single linearized model computed off-line for an identical coil configuration and a similar plasma to that used in experiment. In the long term, the necessary model objects can be computed on-line using straightforward matrix multiplications [7].

REFERENCES.

- [1] G. Ambrosino and R. Albanese, "A survey on modeling and control of current, position, and shape of axisymmetric plasmas," *IEEE Control Systems Magazine*, to be published in October 2005.
- [2] J. R. Ferron, M. L. Walker, L. L. Lao, "Real time equilibrium reconstruction for tokamak discharge control," *Nucl. Fusion*, vol. 38, no. 7, pp. 1055–1066, July 1988.
- [3] M. L. Walker, D. A. Humphreys, E. Schuster, "Some nonlinear controls for nonlinear processes in the DIII-D tokamak," in *Proc. 42nd IEEE Conf. on Dec. Cont.*, Piscataway, 2004, Paper FrE14-4, pp. 6565–6571.
- [4] D. A. Humphreys, et al., "High performance integrated plasma control in DIII-D," presented at the 23rd Symp. on Fusion Technology, Venice, Italy, September 20–24, 2004; accepted for publication in *Fusion Engin. and Design*.
- [5] M. L. Walker, et al., "Next-generation plasma control in the DIII-D tokamak," presented at the 22nd Symp. on Fusion Technology, Helsinki, Finland, September 9–13, 2002, *Fusion Engin. and Design* vol. 66–68, p. 749, 2003.
- [6] ITER Web site [Online]. Available: <http://www.iter.org/>
- [7] M. L. Walker, D. A. Humphreys, "Valid coordinate systems for linearized plasma shape response models in tokamak," submitted to *Fusion Sci. Technol.*, 2005.
- [8] M. L. Walker, D. A. Humphreys, R. D. Johnson, J. A. Leuer, "Nonlinear methods for current limit constraint satisfaction," *Fusion Sci. and Technol.*, vol. 47, pp. 790–795, 2005.
- [9] M. C. Bartholomew-Biggs, "Algorithms for general constrained nonlinear optimization, in algorithms for continuous optimization, the state of the art," Emilio Spedicato, Ed. NATO ASI Series C: Math. and Phys. Sci., vol. 434, pp. 169–208.