

Some Nonlinear Controls for Nonlinear Processes in the DIII-D Tokamak

M.L.Walker, D.A. Humphreys, and E. Schuster*

General Atomics, P.O. Box 85608, San Diego, California 92186-5608

*University of California, San Diego, La Jolla, California 92121

Abstract. In the past few years, much work has been done at various institutions on application of multivariable linear controllers for tokamak shape control. A great deal has been learned about tokamak plasma shape control through these studies. Now some of the practical constraints imposed by working tokamaks must be addressed in order to make the benefits of multivariable control available to working experimental devices. In this paper, we discuss several nonlinear processes and constraints at DIII-D which must be dealt with as part of the implementation of a routine operational controller. We describe partial solutions to the problems already implemented and point out some characteristics required of future solutions to the remaining problems.

1. Introduction

In the past few years, much work has been done at various institutions on application of multivariable linear controllers for tokamak shape control, including several design studies [1-4] and even some experimental implementations [5,6]. A great deal has been learned about tokamak plasma shape control through these studies. Now some of the practical constraints imposed by working tokamaks must be addressed in order to make the benefits of multivariable control available to working experimental devices. In this paper, we discuss several of the nonlinear processes and constraints at DIII-D which must be dealt with as part of the implementation of a routine operational controller. We first describe the present DIII-D operational controller and some of its characteristics which we seek to preserve in upgrading to more advanced control. Next, the present state of multivariable controller development on DIII-D, including some of the nonlinear problems already solved, will be discussed. Finally, the consideration of several remaining nonlinear problems will be discussed.

Due to the large plant size typical for tokamak plasma shape control, we are largely constrained to use of mature linear design tools. By contrast, for many of the nonlinear problems faced, we must appeal to rather immature or even yet to be developed nonlinear control techniques.

2. DIII-D Operational Error Calculation and Control

The isoflux control method, now in routine use on DIII-D, exploits the capability of the real time EFIT [7] plasma equilibrium reconstruction algorithm to calculate total magnetic flux at specified locations within the tokamak vacuum vessel. Figure 1 illustrates a lower single-null (LSN) plasma which was controlled using isoflux control and indicates quantities relevant to the control scheme. The real time EFIT algorithm can calculate 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. 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 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.




Fig. 1. Isoflux control points and X-point grid used for calculation of shape error.

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 (Fig. 2). 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.

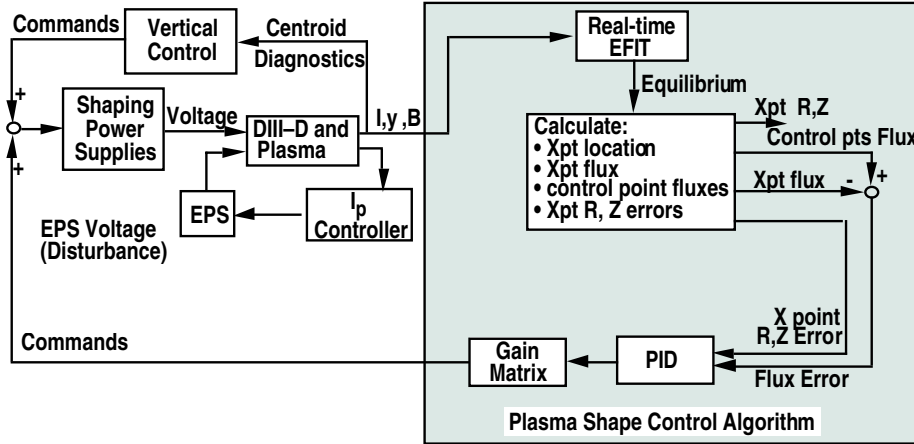


Fig. 2. Present operational PID-based controller.

This approach usually provides good control, but can require significant operator tuning and use of valuable experimental time. In addition, there are several control problems which 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". 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.

3. Practical Constraints and Nonlinear Systems

There are several problems, some closely related, which make practical implementation of shape controllers on DIII-D challenging: (P1) vertical stability control and shape control share the same actuators, i.e. the shaping coils (F-coils) F2A, F2B, F6A, F6B, F7A, and F7B (Fig. 1); (P2) limitations on actuator voltage (from choppers) imply that commands to choppers often saturate, particularly with large or fast disturbances such as programmed changes of shape; (P3) computational speed is insufficient to do both stabilization and shape control with the same controller; (P4) multiple current and voltage constraints such as limits on coil current and a constraint on a sum of currents 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. (P5) Changes in linearized response of the nonlinear plasma during the course of a discharge can cause difficulties in control by linear controllers. (P6) Shape control power supplies (choppers) are extremely nonlinear. (P7) The isoflux error calculation is nonlinear, significantly so in the event of nontrivial programmed motion of the X-point.

Items (1) and (2) taken together constitute a nonlinear problem for sufficiently large disturbances or reference

changes. During disturbances, commands produced by the vertical stability algorithm are large and oscillatory while commands accepted by the chopper supplies are clipped at ± 10 V. The result is that low frequency components of the resulting commands to choppers are determined primarily by time in residence of the commands at the values ± 10 V, rather than by the shape control commands. In these cases, control of the boundary shape near those coils degrades. This is a particular problem for the outer coils because these are the only coils which provide adequate actuation for shape control at the outer plasma boundary.

4. Multivariable Control Algorithm Development

An effort has been underway for some time to construct model based multivariable controllers to address the problems currently seen in operational shape control of DIII-D. Algorithm work has focused on addressing each of the problems listed in Section 3. The problem of nonlinear choppers (P6) was addressed previously by constructing closed loop controllers for the chopper power supplies [8] using a nonlinear output inversion. The issue of sharing of actuators by vertical stabilization and plasma shape control (P1) has also been considered by other developers of multivariable control. In the original proposed ITER design, a separate fast coil with a separate power supply was added to the initial design [1] and the controller was designed to provide simultaneous shape control and stabilization with a single sampling rate [a much slower rate than is necessary for DIII-D, hence (P3)]. For the modified ITER-FEAT design, a rather elegant proposed solution [9] uses an additional fast power supply connected in parallel across multiple control coils to provide extra voltage for stabilization using these coils. In this case, vertical and shape control are performed on separate time scales [10,11]. In an implementation on TCV [5], separate coils were used for stabilization and for shape control, again on separate time scales.

The problem of loss of shape control actuator authority for 6 out of 18 coils was a particular problem for multivariable control development on DIII-D. A nested control method has been implemented to address the sharing of

actuators between vertical and shape control. A linear controller was constructed which simultaneously stabilizes and provides control of the vertical control coil currents on a fast time scale. Figure 3 shows the closed loop system comprised of the DIII-D plant and stabilizing controller. This closed loop system is stable and the 6 coil currents F2A, F2B, F6A, F6B, F7A, and F7B are approximately controlled to be equal to a set of input reference values. As a result, this closed loop system can act as an inner control loop for the shape control. It provides as input actuators the 6 coil current reference signals and (up to) 12 of the original (up to) 18 coil voltages. By integrating control of the vertical control coils into the stabilizing controller, conflicts between shape controller use of these coils and vertical stabilization are eliminated.

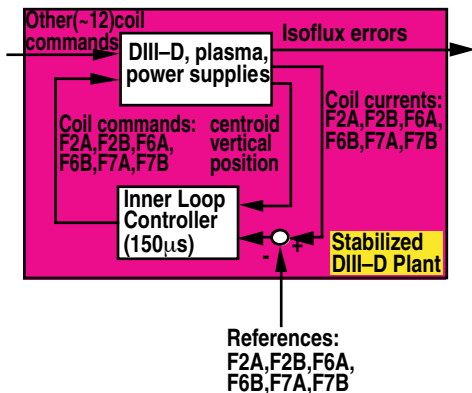


Fig. 3. Block diagram of DIII-D tokamak and plasma stabilized by a fast controller which also regulates currents to specified input reference values.

After the inner loop was tested experimentally and verified to work as expected, an outer loop linear shape controller was designed and tested in simulation [12]. Figure 4 illustrates the shape control loop wrapped around the now stabilized DIII-D plant. A linear controller was developed using the H-infinity method described in [13]. Nonlinear modifications of the basic controller include logic to avoid current limits [6] and a variable loop gain to prevent controller windup [14].

The tradeoffs which have been made should be noted here. Shape control response speed for the 6 vertical control coils was degraded in order to make those coils responsive to shape control commands at all times.




Fig. 4. Outer loop shape control loop around DIII-D plant after stabilization by inner current control loop.

(Vertical control response appears to be much smoother while still well stabilized.) Closing of the current control loop was responsible for some of the slower response, but a more significant cause is the present need to be conservative in use of these actuators in the linear shape controller design. Requests for too fast a current change can cause voltages driving these coils to saturate in response. Once saturated, they are no longer able to respond effectively to high frequency disturbances and instability can result. Thus, we have a new (nonlinear) problem - how to use these coil current actuators to their fullest potential without risking causing the system to go unstable. Some work on a solution to this problem is presented in [15,16].

5. Current and Voltage Limits

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 which 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 which imply that the 15 control parameters cannot always be simultaneously minimized.

In fact, it is easy for an operator to choose a reference shape which is incompatible with device constraints, even in steady state. Figure 5 illustrates the relationship between currents necessary to maintain the programmed shape specified during plasma discharge number 99339 at 1800 ms with the coil current limits imposed by various hardware constraints. The largest curve represents the current perturbations from actual equilibrium currents attained at 1800 ms which are needed to attain a zero steady state error for all control points and the X-point for the reference equilibrium specified during that shot. These

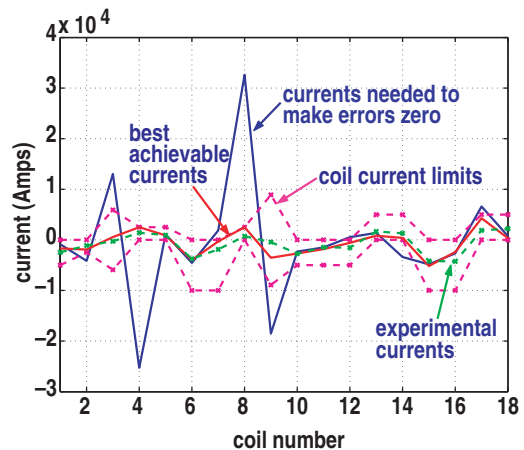


Fig. 5. Currents necessary to obtain zero steady state error lie outside the coil current limit boundaries.

violate the hardware current limits (shown by upper and lower dashed lines) for several coils. The result of a calculation which minimizes the mean-squared shape errors constrained by the coil current limits is also shown.

Even with a small number of control parameters, the currents required to obtain zero error for those parameters at equilibrium is 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. In the figure above, any linear controller would be unaware of the current limits and would try to obtain the required currents, thus driving at least 6 coils into their limits. This is in fact what happens when a linear controller as illustrated in Fig. 4 is executed in closed loop simulation with the model plant [12].

The issue of number of controlled parameters has been considered before. The choice was made in Alcator C-Mod [17] to use an overdetermined system for reasons of flexibility. A similar choice was made at DIII-D. In devices such as ITER, where certain parameters (e.g. antenna coupling, divertor strike point locations, portion of scrape off layer intersecting top of machine) must be rigidly controlled, a smaller number of controlled parameters seems appropriate. In one ITER scheme, 6 independently driven PF coils and a segmented ohmic coil stack control a total of 6 gaps and the plasma current. The location and size of these coils have been carefully designed to produce the desired envelope of plasma shapes.

Limiting of coil currents is presently handled at DIII-D through operator intervention between discharges, by modifying the control or even by adding an additional power supply so as to increase the current limit. Problems of this type must be expected in an experimental device, where new and untested equilibria are continually created according to the needs of the experimental program. However, a worthwhile goal is to provide a "soft landing" for these cases, to include a graceful degradation of reference tracking in the face of impending current limits, operator methods for defining how this degradation should take place, and a warning rather than a premature end of the plasma discharge in the event of incompatible equilibrium difficulties.

Since this problem is primarily caused by requesting a reference shape which is not physically realizable by the device, what seems to be an obvious "fix" to the problem is to specify only reference shapes which are compatible with all constraints of the device. Plans for future reactors such as ITER 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 which are completely compatible with the device constraints. These equilibrium calculations are dependent on current profiles however which, in experimental devices, can change in an uncontrolled manner during a plasma discharge. They are also sensitive to the accuracy of the plasma and constraint models used. The models of voltage constraints imposed by the DIII-D poloidal power systems, for example, will never be completely accurate because of the extreme nonlinearity of the chopper power supplies.

Thus, some real time methods are still desirable to handle cases where coil currents or voltages are near their limits. Even the ITER design includes a supervisory layer which is prepared to react to violations of certain constraints.

An ideal real time solution to this problem would have certain characteristics. It would:

- Guarantee that the control would never cause the system to attempt to violate device constraints
- Impose no constraints on system performance in the absence of proximity to those constraints
- Provide a systematic design procedure, with guarantees for performance and stability
- Allow for operator specification of the relative importance of control of individual error signals.

The reason for the last objective is that some quality of control must be sacrificed as constraints are approached, in order to prevent actually violating the constraint. Experimentalists need to have some mechanism for specifying which control parameters should be sacrificed first in these situations.

Probably the first effort to address this problem in tokamaks was for an application to the ITER design [18] where reference signals were modified in response to coil current proximity to current limits. The modified values were based on off-line calculation of the minimum "safe" reference change which would ensure currents did not exceed their limits, based on the closed loop steady-state gain from reference signal to coil currents. During on-line control, the reference signals were interpolated between operator specified references and the off-line precomputed "safe" values based on proximity of coil currents to their proscribed limits.

A method which was implemented on DIII-D [19] involved feedback of coil current "errors" in order to prevent coil current limit faults from ending plasma shots. Coil current reference signals were constructed as heavily filtered versions of measured coil currents whenever currents were not near limits. A coil reference signal was modified by nonlinear logic to produce a large error if the corresponding current approached a limit. In at least three cases during experimental testing, F-coils came close to limiting values, at which point the algorithm increased the coil current error and prevented those currents from causing faults. Although this appears to provide an adequate solution, the external nonlinear logic used to handle coil limits is presently tuned via simulations because there is essentially no theory which provides guidance for how to specify this nonlinear process to maintain stability and avoid limit conditions. If improperly tuned, limit cycle oscillations can occur. In addition, this approach provides little direct control over where the inevitable degradation of boundary control will occur.

Alternative methods have also been proposed or are planned for implementation. A method which uses an SVD reduction of the output dimension for the mapping from coil currents to control parameters is planned for the "extreme shape controller" to be implemented on the JET tokamak [20]. The idea is to find a reduced order best approximation

to the reference shape specified by the tokamak operator with a shape more compatible with the actuator degrees of freedom of the device.

Another proposal is to apply current regulation logic to individual coils whose current approaches a current limit in order to prevent them from attempting to violate that current constraint. The idea is to change the limiting constraints from being limits on the states of the system to being modifications of the input (voltages) of the system, which defines a more traditional anti-windup problem. Unfortunately, for multivariable systems, the "best" solution to this problem is also not completely resolved. For SISO systems, a number of techniques are available; for example, a theoretically sound SISO antiwindup method [21] was implemented on DIII-D in individual chopper voltage control loops with excellent results [8]. This method falls short however for MIMO systems, where it is necessary to worry about the proper "directionality" of the actuator commands; the current DIII-D MIMO implementation uses an ad hoc variation of a method proposed in [14].

In the past few years, an increasing number of methods have been proposed in the control research community for dealing with actuator and plant limitations. The concept of a reference or command governor has been proposed by many researchers (e.g. [22-25]). A reference or command governor is a nonlinear memoryless transformation of the system state and reference input to a constrained reference signal, which has the property that the constrained value is the same as the input reference in the absence of constraints.

The reference governor in [23] is based on the original idea of an error governor in [25], which multiplies all errors by a constant <1 to prevent constraints from being active. This is similar to the approach in [14]. The idea is to maintain the directionality of the error vector by multiplying errors by a scalar gain ≤ 1 . Other methods (e.g. [22,24,25]) propose use of on-line optimization to compute a modified reference signal which is near to the original reference but which causes predicted future values of states and inputs to satisfy their constraints.

The combination of improved optimization algorithms and installation of a new, faster realtime computer system on DIII-D [12] makes it feasible to consider on-line optimization methods. A simple version of this on-line optimization approach which has been implemented in simulation for evaluation with DIII-D shape control seeks an on-line constrained replacement e_c of the original computed error signal vector e_0 via the constrained minimization problem

$$\min_{e_c} \|W(e_c - e_0)\|^2, \\ \text{subject to } \begin{cases} I_{\min} \leq I_{eq} + G_{eI}(0)e_c \leq I_{\max} \\ \text{sign}(e_c) = \text{sign}(e_0) \end{cases}$$

where I_{eq} is the present time equilibrium current vector, $G_{eI}(0)$ is the steady state open loop gain from errors to F-coil currents, and W is a diagonal weighting matrix which

can be used to increase relative importance of matching certain errors over others. Figures 6 through 8 show a comparison of two simulations, one in which the linear shape controller is executed without consideration for limits on coil currents and one in which the on-line optimization is performed. Figure 6 shows the current in F5B going to its constraint value of 0 at about 0.52 s for the simulation without limit prevention logic (dashed lines). In actual




Fig. 6. Coil currents from simulations with (solid) and without (dashed) current limit prevention logic.

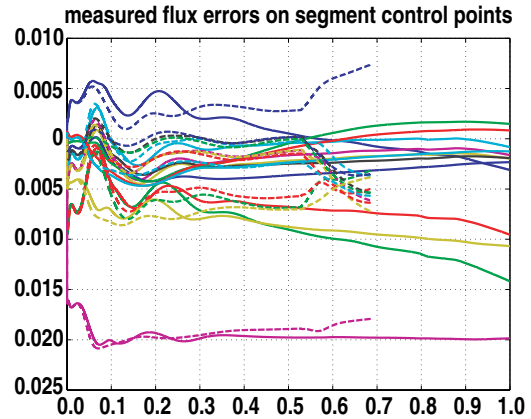


Fig. 7. Control point flux errors from simulations with (solid) and without (dashed) current limit prevention logic.




Fig. 8. X-point R and Z from simulations with (solid) and without (dashed) current limit prevention logic versus target X-point location (dash-dot).

operation, this behaviour would have caused a chopper power supply to fault and ended the plasma discharge. In contrast, the currents in the simulation which included the constrained optimization calculation (solid lines) all remain well away from their current limits. This method is a form of error governor which does not attempt to keep the error vector direction the same as the original, but instead modifies the direction so as to maintain smaller errors on the more important (larger weighted) control quantities.

The fact that nothing comes for free is illustrated in Figs. 7 and 8 which show the tracking of the operator specified references by the two simulations. In both figures, the simulation without limit prevention (dashed lines) shows a slightly better tracking of reference signals prior to 0.52 s, at which time a serious degradation of control results. The figures show that the tracking by the simulation with limit prevention is comparable to the original simulation initially, then gets slightly worse as the limiting coil gets close to the current boundary. At that time, a slow increase in error can be seen in the figures as the constrained optimization solution starts to "hide" the larger errors from the controller to prevent coils from attempting to reach their current limits.

6. Summary

A great deal has been learned about tokamak multivariable plasma shape control through various control studies and prototype implementations. In order to make the benefits of multivariable control available to working experimental devices, some of the practical nonlinear constraints imposed by working tokamaks must be addressed. This paper has summarized a number of nonlinear control issues which must be addressed in order to provide the advantages of multivariable control for DIII-D experimental operations.

A summary of some of the problems with the present operational control was provided. These problems have motivated the work on advanced control. In providing advanced controllers for experimental devices, we must be cognizant of the intended users of these controllers. Tokamak physics operators will not accept reduced performance simply to implement an advanced controller. There must be clear advantages shown. For example, they are likely to accept a small reduction in some facet of performance if the tradeoff is a large gain in some other area, such as those discussed in this paper. They will also demand that some influence over these tradeoffs be made available to them at the control console.

An overview of the present operational control and of the status of the work in progress on a replacement multivariable controller was provided. One nonlinear problem, due to ineffective sharing of actuators, has been solved but in so doing has raised a secondary nonlinear problem of limiting actuators. Current work focuses on methods for handling nonlinear current and voltage limit constraints, including the problem of ensuring robust stabilization of the plasma vertical instability in the face of possible loss of actuation due to saturation.

Acknowledgment

This is a report of work supported by the U.S. Department of Energy under Contract Nos. DE-AC03-99ER54463 and DE-FG03-01ER54615.

References

- [1] Albanese, et al., Plasma Current, Shape, and Position Control in ITER, *Fusion Tech.* **30** (1996) 167.
- [2] A.S. Sharma, et al., "Modeling and Control of TCV," Centre de Recherches en Physique des Plasmas Report LRP 715/01 (2001).
- [3] M.L. Walker, D.A. Humphreys, J.R. Ferron, "Control of Plasma Poloidal Shape and Position in the DIII-D Tokamak," Proc. 36th IEEE Conf. Decision and Control, San Diego, CA, 1997, 3703-3708.
- [4] G. Ambrosino, M. Ariola, Y. Mitrishkin, A. Pironti, A. Portone, "Plasma Current and Shape Control in Tokamaks Using H-infinity and m-synthesis," Proc. 36th IEEE Conf. Decision and Control, San Diego, CA, 1997, 3697-3702.
- [5] M. Ariola, et al., "A Modern Plasma Controller Tested on the TCV Tokamak," *Fusion Tech.* **36**, 126-138 (1999).
- [6] M.L. Walker, et al., "Implementation of Model-based Multivariable Control on DIII-D," *Fusion Eng. and Design* **56-57**, 727-731 (2001).
- [7] J.R. Ferron, et al., "Real Time Equilibrium Reconstruction for Tokamak Discharge Control," *Nucl. Fusion* **38**, 1055-1066 (1988).
- [8] M.L. Walker, et al., "Status of DIII-D Plasma Control," Proc. 16th IEEE/NPSS SOFE (1995) p. 885.
- [9] P.L. Mondino, et al., "Studies to Increase the Elongation in ITER: Impact on Passive and Active Plasma Vertical Position Stabilization Systems," *Fusion Eng. and Design* **56-57**, 699-704 (2001).
- [10] A. Beghi, et al., "Plasma Vertical Stabilization in ITER-FEAT," *Fusion Eng. and Design* **56-57**, 783-788 (2001).
- [11] A. Portone, "Design of the ITER-FEAT Plasma Control System by a Frequency Separation Approach," *Fusion Eng. and Design* **56-57**, 789-794 (2001).
- [12] M.L. Walker, et al., "Next-Generation Plasma Control in the DIII-D Tokamak," Proc. 22nd Symp. Fusion Technology, Helsinki, Finland, 2002, to be published in *Fusion Eng. and Design*.
- [13] D.C. McFarlane, K. Glover, "Robust Controller Design Using Normalized Coprime Factor Plant Descriptions," *Lecture Notes in Control and Information Sciences* **138**, Springer-Verlag, 1990.
- [14] R.A. Hyde, "H-infinity Aerospace Control Design. A VSTOL Flight Application," *Advances in Industrial Control Series*, Springer-Verlag, 1995.
- [15] E. Schuster, M.L. Walker, D.A. Humphreys, M. Krstic, "Plasma Vertical Stabilization in Presence of Coil Voltage Saturation in the DIII-D Tokamak," Proc. 2003 Amer. Control Conf., Denver (2003).
- [16] E. Schuster, M.L. Walker, D.A. Humphreys, M. Krstic, "Anti-windup Scheme for Plasma Shape Control with Rate and Magnitude Actuation Constraints in the DIII-D Tokamak," this conference.
- [17] I.H. Hutchinson, et al., "Plasma Shape Control: A General Approach and its Application to Alcator C-Mod," *Fusion Tech.* **30**, 137-150 (1996).
- [18] G. Ambrosino, et al., "A Control Scheme to Deal with Coil Current Saturation in a Tokamak," *IEEE Trans. Control Sys. Tech.* **9**, 831 (2001).

- [19] D.A. Humphreys, M.L. Walker, J.R. Ferron, J.A. Leuer, "Initial Implementation of a Multivariable Plasma Shape and Position Controller on the DIII-D Tokamak," Proc. IEEE Conf. on Control Applications, Anchorage, Alaska, 2000.
- [20] G. Ambrosino, et al., "Specifications for the Extreme Shape Controller," private communication.
- [21] R. Hanus, et al., "Conditioning Technique, a General Anti-windup and Bumpless Transfer Method," *Automatica* **23**, 729-739 (1987).
- [22] A. Casavola, "Reference Governor for Constrained Uncertain Linear Systems Subject to Bounded Input Disturbances," Proc. 35th IEEE CDC, Kobe, Japan (1996) pp. 3531-3536.
- [23] E.G. Gilbert, I. Kolmanovsky, K. TinTan, "Nonlinear Control of Discrete-Time Linear Systems with State and Control Constraints: A Reference Governor with Global Convergence Properties," Proc. 33rd IEEE CDC, Lake Buena Vista, Florida (1994) pp. 144-149.
- [24] A. Bemporad, A. Casavola, E. Mosca, "Nonlinear Control of Constrained Linear Systems via Predictive Reference Management," *IEEE Trans. Aut. Contr.* **42**, 340-349 (1997).
- [25] D. Angeli, A. Casavola, E. Mosca, "Command Governors for Constrained Nonlinear Systems: Direct Nonlinear vs. Linearization-Based Strategies," *Int. J. Robust Nonlin. Control* **9**, 677-699 (1999).
- [26] P. Kamasouris, M. Athans, G. Stein, "Design of Feedback Control Systems for Stable Plants with Saturating Actuators," Proc. 27th IEEE CDC, Austin, Texas (1988) pp. 469-479.