

A Plant Taxonomy for Designing Control Experiments

By Dennis S. Bernstein

A control experiment is its own analog simulation.

—Anonymous

Control experiments can have a significant impact on control theory by forcing researchers to confront real-world issues that affect design tradeoffs and performance specifications. Sensor and actuator constraints, modeling and identification issues, and hardware imperfections (such as noise, drift, bias, and nonlinearity) must all be addressed for successful controller implementation.

Control experimentation, however, is not an established discipline, and there are many fundamental issues and questions that are worthy of deep and careful consideration. In particular, the definition of a control experiment is open, as are guidelines for verification and reproducibility. I will not attempt to address these important questions here. Rather, my objectives are limited to assessing plant features that are appropriate for investigating system-theoretic problems in feedback technology. The term *system-theoretic* refers to issues such as phase variation, nonlinearity, uncertainty, accessibility, and cross coupling, which transcend a specific application or hardware realization.

This discussion of control experimentation venues is based on a *plant taxonomy*; that is, a systematic classification of plant properties and the challenges they present to control experimentation and, indirectly, to control engineering practice. Plant properties and the limits they impose on achievable performance under linear control are analyzed in [1]-[5] and the numerous references given therein. A related analysis in the context of active noise control is given in [6]. Many of the issues and tradeoffs discussed in [1]-[6], as well as additional issues of relevance to control experiments and hardware implementation, are addressed here.

This article is partially motivated by [7], which lists 16 candidate plants for undergraduate control experiments. Their selection criteria are: interesting, visual, instructive, and reasonably challenging. I am especially interested in the features of these and other plants that render them worthwhile for experimentation for either education or research. Additional motivation is provided by [8], where I discussed the lessons I learned from working with students and colleagues in designing, building, and operating various control experiments. Details concerning the experiments described in [8] are given in [9]-[17].

My objective is to provide perspective on some of the issues that arise in designing control experiments for both ed-

ucation and research. I hope that this discussion will serve as useful guidance for educators designing control experiments to illuminate control issues, as well as for researchers developing experiments to complement theoretical and computational research in control.

Problem Classification

The design of a control experiment must begin with an explicit control system objective. I consider three largely distinct objectives:

- *Stabilization/dynamics modification.* This control system objective seeks to modify the dynamics of the system through feedback. A special case is that of stabilization, where the open-loop system is unstable and the control system must render an equilibrium of the closed-loop system stable in some sense (local, semiglobal, global). More general cases include pole placement and improved transient response, where the dynamics are suitably modified. These problems generally focus on the inherent dynamics of the system, with minimal emphasis on exogenous disturbances or command signals. In this class of problems, the domain of attraction usually plays an important role.
- *Disturbance rejection.* This control system objective seeks to ensure system performance in the presence of exogenous disturbances. Disturbances may be persistent or decaying, tonal or multitonal, periodic or nonperiodic, or stochastic with fixed or variable spectrum. The disturbance signal may be uncertain, although some information may be available regarding its current or future spectral or temporal behavior. In special cases, a measurement of the disturbance is available, in which case a feedforward architecture can be used. Disturbance rejection is usually an equilibrium-related notion with the objective of rejecting disturbances while remaining close to a specified equilibrium. Simultaneous disturbance rejection and command following/tracking can be considered as well.
- *Command following and tracking.* In command-following problems, the objective is to have selected plant variables follow user-specified input signals. In some cases, the form of the command-following problem is

The author (dsbaero@umich.edu) is with the Aerospace Engineering Department, University of Michigan, Ann Arbor, MI 48109-2140, U.S.A.

similar to that of a disturbance rejection problem, although there are important differences. For example, the present value of the command is known in command following, since it is generated by the user, and it is used to determine the command-following error. In tracking problems, the future of the command is known as well. Also, unlike the disturbance rejection problem, command-following objectives need not be associated with an equilibrium.

The design of a control experiment must begin with an explicit control system objective.

Plant Taxonomy

The challenges that arise in using feedback control to achieve dynamics modification, disturbance rejection, and command following depend on diverse factors. To understand the challenges that these factors impose on achievable performance, I consider several fundamental issues in isolation. Except for open-loop instability, which is discussed first because I believe it is the most severe, I do not claim that this ordering reflects relative difficulty. Nevertheless, in my experience this listing correlates roughly with the factors that render plants difficult to control in hardware experiments.

Open-Loop Instability

In my opinion, the distinction between stable and unstable plants is vastly underemphasized in the research literature. An unstable plant provides almost no opportunity for online identification, so control engineers must rely on analytical modeling and extrapolation from stable regimes. Identification of unstable plants in closed-loop operation is often discussed, but this approach presupposes a stabilizing controller, which is often unknown prior to identification. One solution to this problem is to employ adaptive stabilization; however, all adaptive stabilization methods require at least some prior modeling information [18]-[20].

Linearization of unstable plants warrants special attention. The stability of a plant whose linearization has non-repeated poles on the imaginary axis cannot be determined from linearized analysis, but rather depends on nonlinear effects. A plant with a chain of integrators or imaginary poles is unstable in a polynomial, but not exponential, sense. One example is the double integrator, which is meaningful in applications and can be viewed as inherently linear. On the other hand, a plant with one or more open right-half plane poles is more seriously unstable, but there appear to be no such plants that are truly linear. Instead, all such models appear to be linearizations of nonlinear plants whose

nonlinearities can exacerbate performance over the range of operation, no matter how small. (Think of the inverted pendulum as it moves farther and farther from the vertical equilibrium.) This point will be discussed below in the context of saturation effects.

Unstable plants are unforgiving in the sense that, once large deviations occur, saturation limits may prevent recovery. Furthermore, linearizing a nonlinear unstable plant may obscure the actual saturation recovery limits, which are invariably smaller than those of the linearized model. This point warrants investigation.

In addition to saturation, open-loop instability exacerbates virtually all other difficulties in feedback control. Therefore, in the design of a control experiment, the first and most critical decision concerns whether the chosen plant will be stable or unstable. Although the inverted pendulum is unstable, many of the difficulties associated with unstable dynamics are circumvented by low-order dynamics, high actuator authority, and feedback accessibility. These issues will be discussed in the following sections.

Feedback Accessibility

Feedback accessibility refers to the extent to which the sensors and actuators provide actuation and sensing of the plant dynamics. It is a myth of feedback technology that single-input, single-output (SISO) plants are categorically easier to control than multiple-input, multiple-output (MIMO) plants. On the contrary, multiple sensors and actuators often allow the controller to have greater access to the plant dynamics and thus achieve better performance than a single sensor/actuator pair. To clarify this point, think of the guaranteed gain and phase margins of the MIMO full-state-feedback linear-quadratic regulator (LQR), which are not shared by the SISO or MIMO output feedback linear-quadratic-Gaussian (LQG) compensator.

This point also applies to the use of much simpler controllers. Indeed, the standard proportional-integral-derivative (PID) controller is a MIMO controller since the “D” denotes the availability of a differentiated output, so that a true PID controller requires two sensors. Note that improper controllers are not considered to be controllers per se since they are not implementable except through approximation by proper compensators, which roll off at high frequency.

More specifically, what often makes a SISO plant difficult to control is its phase variation over the control bandwidth. By phase variation, I am referring to the total phase variation between zero frequency (dc) and high frequency, as well as the phase fluctuation that occurs over the control bandwidth as determined by the plant poles and zeros. Vari-

ous factors can contribute to phase variation, such as high dimensionality, nonminimum phase (right-half plane) zeros (due to noncolocation of the sensor and actuator), high relative degree (due to sensor/actuator placement), and transport delays. The impact of phase variation on controller tuning can be seen from Bode and Nyquist methods.

The main point is that the difficulty in controlling a plant is often due to the lack of accessibility to the plant dynamics by the allowable control hardware through undersensing or underactuation. In fact, SISO control (one sensor and one actuator) is the most stringent case short of open-loop control (zero sensors) and no control at all (zero actuators). Ironically, by teaching students SISO classical control techniques first and MIMO state space control methods later, we give the impression that SISO control is somehow easier than MIMO control. This is the source of the MIMO myth.

Roughly speaking, a mechanical system is fully sensed if each degree of freedom (rotational, translational, or vibrational) is force actuated and position or velocity sensed. For such a system in decoupled form, little modeling is required, and direct adaptive control methods are often effective, even in the presence of severe nonlinearities [21]. Consequently, rigid robots with fully actuated joints are not inherently difficult to control (although their kinematics and dynamics can be quite complicated). On the other hand, flexible structures and acoustic plants involve numerous degrees of freedom, and thus full actuation is rarely possible.

Some plants, however, are difficult to control even with a high degree of accessibility. The difficulty is not due to dimensionality per se, but rather is due to the extent of cross coupling among the sensors and actuators through the plant dynamics. For example, a flexible structure with a large number of sensors and actuators would still be difficult to control due to sensor/actuator coupling, where “large” means on the same order as the number of significant structural modes. Thus, the difficulty of MIMO control is not due to the *number* of sensors and actuators, but rather the manner in which they interact.

Plants that possess both large phase variation (such as those with nonminimum phase zeros, high relative degree,

or delays) and unstable dynamics are inherently difficult to control in the sense of extreme sensitivity to modeling uncertainty [1]-[6]. The classic example is the inverted pendulum on a cart with cart force control input and cart position sensing, which is unstable and has nonminimum phase zeros. This system is so difficult to control that it is suggested in [5] that sensor-actuator redesign of the plant is warranted. This analysis is thus valuable for control architecture design, as well as for identifying challenging control experiments.

An experimental difficulty associated with undersensed plants is the inability to measure or set the initial conditions. This problem arises in controlling plants with numerous degrees of freedom (for example, flexible or acoustic modes) or inaccessible states (such as an internal voltage). Without this ability it may be difficult to reproduce and validate experimental data. In some cases, this issue can be circumvented by fully instrumenting the plant for diagnostic purposes while using only a subset of the sensors for feedback.

Sensor/Actuator Limitations

Although plant accessibility as determined by the number, type, and placement of sensors and actuators has a significant effect on the ability to control a plant, the *authority* afforded by the sensors and actuators is also critical. For example,

if sensors and actuators saturate, then plant accessibility is effectively decreased. During actuator saturation, the control input is constant and the plant is effectively operating in open loop, although closed-loop control will return if and when saturation ceases.

Saturation is often the first nonlinearity encountered by the control engineer [22], and, because of its effect on plant accessibility, I discuss it independently of more general nonlinearities. Because all real actuators saturate (in both amplitude and rate), this nonlinearity can never be completely circumvented by any technological development. Saturation is a continuous nonlinearity that has little effect on the local (near equilibrium) behavior of the system. Sensor saturation can also occur and has an analogous effect on plant observability [23], [24].

In practice, it is desirable to achieve the best possible performance from the available control system hardware. As-



©1995 Master Series

suming there is no fuel or energy constraint, this goal may require that the actuators operate at or near saturation. In this case, saturation limits are not necessarily regions to be avoided, but rather are sought so as to maximize usage of the available control input. Keep in mind, however, that, for some plants, the desirable control values need not lie near the boundaries. For example, in driving your car, the extreme right and extreme left steering wheel angles are rarely used, and thus bang-bang control is not optimal. Although it is not surprising to find the best control input on the boundary of the set of admissible control values, it is far more interesting when the optimal control occurs at an intermediate point. This is the subtle case of singular control [25].

The difficulty in controlling a plant is often due to the lack of accessibility to the plant dynamics by the control hardware through undersensing or underactuation.

The distinction between stable and unstable systems is important when addressing saturation issues. If the plant is open-loop stable, saturation is an issue only when performance is quantified since the zero control is unsaturated and stabilizing. On the other hand, global and semiglobal stabilization of plants with open right-half plane poles is impossible in the presence of saturation. Therefore, maximizing the domain of attraction is often the primary objective for unstable plants. In fact, a rare disturbance of high magnitude can perturb the state and render the equilibrium unrecoverable. This problem is critical when considering the use of feedback for disturbance rejection on unstable systems.

As discussed above, linearizing a nonlinear unstable plant may obscure the actual saturation recovery limits, which are invariably smaller than those of the linearized model. Hence, maximizing the domain of attraction of the linearized plant may not be an effective control strategy.

Actuator limitations often arise in the form of actuator dynamics, which appear as gain rolloff. In practice, actuator rolloff is often modeled as a rate saturation constraint. Note that the gain rolloff model is linear, whereas rate saturation is nonlinear. The relationship between these models warrants investigation.

A philosophical problem that arises in designing a control experiment concerns the level of control authority to provide. If the available control authority is high, then the

control problem may be too simple, especially if the plant is fully actuated. If the plant is underactuated, then the experiment may be challenging even if the sensor/actuator authority is unlimited. On the other hand, if the control authority is too low, then it may not be possible to achieve significant performance improvement relative to open-loop performance, and thus the experiment may not be worthwhile.

Modeling Uncertainty

Modeling, which includes a deep understanding of the relevant physics, is essential and invaluable for control architecture design, which concerns the choice, sizing, and placement of sensors and actuators. This kind of modeling, which is usually analytical (and thus hypothetical), is often quite different in character and requirements from modeling for controller tuning, which may be largely empirical (that is, data based).

Modeling for controller tuning is made difficult by the fact that a seemingly small physical change can have a large effect on plant dynamics. For example, as lubrication dissipates, bolts loosen, and battery voltages drop, dynamic response can change dramatically. The effects of these unpredictable changes on control system performance are the responsibility of the control engineer. Indeed, one of the main reasons for implementing a feedback control system is to achieve performance in the presence of uncertainty, but not all uncertainty can be characterized or predicted.

No matter how well analytical modeling can be performed for control architecture design, some identification is usually needed for controller tuning. Real hardware abounds with manufacturing variations and imperfections. In addition, modeling a system in a piecemeal fashion is of limited usefulness for controller tuning, since components can interact dynamically in complicated ways due to spurious feedback paths and unexpected interactions such as impedance mismatch. Consequently, end-to-end identification is desirable whenever possible. Obviously, identification is only meaningful after the system has been constructed and data can be obtained.

The ability to perform identification depends on the nature of the plant, as well as on the environment. Identification of the uncontrolled plant is generally not feasible if the plant is open-loop unstable. In this case, a stabilizing controller is needed, thus requiring analytical modeling or adaptive methods. In addition, the presence of ambient disturbances can adversely affect the ability to identify and adapt.

When applying identification techniques to real systems, there is no conclusive arbiter of success; that is, the “real” system is unknowable except through the data collected and the identification method employed. Identification methods can be tested on hypothetical, simulated plants, but all real plants possess features that are unmodeled and unmodelable.

The objective of robust control is to guarantee performance for a given level of modeling uncertainty. Consequently, robust control facilitates controller synthesis with relaxed model accuracy specifications. To do this, robust control requires an accurate characterization of the model uncertainty. Thus, the control engineer must answer the question: What don't you know, and exactly how well don't you know it? Obtaining an accurate uncertainty characterization may require substantial testing and analysis, which partially defeats the original reason for robust control. In addition, modeling uncertainty is usually statistical despite the plethora of robust control methods based on deterministic uncertainty characterization.

In practice, robust control forces the control engineer to give up performance for robustness. Ultimately, robust control requires that the controller gains be decreased to account for uncertainty, thereby reducing performance. The inability of a robust controller to learn makes this tradeoff unavoidable.

Robust control combined with continual or intermittent identification constitutes adaptive control. In practice, it is unrealistic, if not dangerous, to assume that the plant is unchanging, and it is certainly advantageous to try to improve model fidelity during operation. In some applications, however, it may not be feasible to inject identification signals into the plant during operation. For example, identification during operation will certainly degrade performance. An additional obstacle is the fact that, while the control signal must have authority that is roughly comparable to the disturbance level, a useful identification signal may need to be of significantly larger magnitude. This point warrants investigation. Another issue is that many identification methods give biased parameter estimates when the input and disturbance are correlated, which occurs during closed-loop identification.

In a laboratory setting, control experiments can be modified during operation to simulate variations in plant characteristics and disturbance spectra. The challenge to the experimentalist is to implement these changes in a verifiable and repeatable manner so that the test data are meaningful.


Plant Nonlinearity

“Linearity over a range” is an oxymoron; that is, a mathematical contradiction in terms. Nevertheless, it is a useful engineering concept. However, nonlinear effects assume greater importance as performance requirements become more stringent. Many control methods consider smooth nonlinearities, which are linearizable near equilibria and which have an increasing effect over a larger range of operation. Geometric and kinematic nonlinearities in multi-body systems such as spacecraft or robotics are examples of smooth, global nonlinearities. Control theorists often assume that these nonlinearities are sufficiently well known that the functions and their derivatives can be used in transformation techniques.

On the other hand, many control applications require accurate motion over small amplitudes. In this regime, the nonlinearities tend to be nonsmooth and possibly discontinuous. Friction [26] is a common example of a nonsmooth nonlinearity. In addition, nonsmooth nonlinearities such as stiction and backlash may possess hidden, unmeasurable states. These nonlinearities often give rise to a continuum of equilibria, which have semistable behavior and which exhibit hysteresis under quasi-static operation [27], [28].

Classical control theory addresses smooth nonlinearities through absolute stability theory. Nonsmooth nonlinearities, however, are equally prevalent in the applications literature. Although large amplitude motions can sometimes be slowed down without major loss of performance (and this is done in practice, such as in robotics, to reduce the effects of geometric nonlinearities on the inertia and gyroscopic terms), lack of precision in small-amplitude applications can seriously degrade the value of a mechanical system. In particular, robots and multibody spacecraft can be repositioned more slowly if necessary, but a lack of precision in a machining operation may not be tolerable at all.

Although nonsmooth nonlinearities involve a smaller amplitude range, these nonlinearities come in a wide variety of types, they may be hidden, and they may change drastically and unexpectedly over different operating ranges; that is, they may not be repeatable. Smooth nonlinearities are difficult to identify because of the range of operation re-



©1995 Master Series

My fascination with control experiments stems from my belief that 1) serendipitous discoveries are exciting and potentially valuable (for example, penicillin, Teflon, Silly Putty, and many more); 2) you can't simulate what you can't model, and you can't model everything; and 3) the bandwidth of reality is beyond technological emulation (despite the specious illusions of *The Matrix*). There is more complexity in a cup of tea than anyone can imagine. None of these beliefs should be construed as detracting from the incredible power of intellectual abstraction and representation, which suggests that the collective human mind has a fighting chance of comprehending the subtleties of reality. In the dawning age of digital virtual reality, I think it is important to remember the distinction between what is real and what isn't.



quired to collect data. Control theorists tend to view such nonlinearities as well known because of the analytical nature of idealized models. In applications such as aircraft maneuvers over a large flight envelope, identification of global nonlinearities is difficult. In such cases, linearized models based on stability derivatives are constructed for a collection of operating points.

Control experiment design to capture nonlinearities is both easy and difficult. In fact, it is fair to say that the true challenge is to design a control experiment (or any system, for that matter) that is *linear*. Although few open-loop plants have truly linear dynamics, sensor and actuator nonlinearities invariably render the controlled plant nonlinear. Hence, every plant that can be built will exhibit nonlinearities. The difficult challenge then becomes characterizing these nonlinearities and overcoming their (usually negative) effect on performance.

For fully actuated systems, adaptive methods can tolerate some unmodeled nonlinearities, as noted in [21]. However, unmodeled nonlinearities combined with unmodeled dynamics or persistent disturbances can destabilize an adaptive controller [29].

An additional difficulty associated with nonlinear plants is the coupling between stability and disturbance rejection. For linear systems, asymptotic stability is equivalent to bounded input, bounded output (BIBO) stability. This equivalence does not hold, however, for nonlinear systems, although there are notions of input-output stability [30], [31]. Theoretical examples and experience show that a nonlinear system with an asymptotically stable equilibrium can have

unbounded response in the presence of arbitrarily small amplitude disturbances. This phenomenon represents a potentially serious impediment to nonlinear identification of stable or stabilized nonlinear systems where the disturbance signal is an excitation input for identification.

State Constraints

Many control applications involve plants whose states must avoid or are confined to certain regions or values. For example, in mechanical systems such as robotics or vehicle platooning, collisions must be avoided. These are hard constraints that must be reliably satisfied due to the high cost of violation. This problem is one of the most difficult and practically meaningful in control technology, but only indirect or limited theoretical techniques are available. Control experiments can readily be designed to capture this kind of hard constraint. In fact, this kind of plant difficulty is ideal for control experiments where the constraints can be violated under laboratory conditions without serious penalty.

Discussion

As control experiments become more widespread as a serious vehicle for advancing control theory and technology, there is a greater need to understand and articulate issues concerning the design and operation of meaningful experiments. In this article I have constructed a taxonomy of plant features to provide a systematic framework for characterizing these issues.

One such issue of extreme importance in control engineering is modeling uncertainty. This issue is subtle for numerous reasons, not the least of which concerns the knowledge possessed by the control engineer or experimentalist who is designing and implementing the controller. In this sense, controller tuning is an art, involving procedures and knowledge that are unique to each practitioner. In practice, control engineers tweak gains and employ empirical rules of thumb to tune and diagnose control systems. We need to understand how such techniques can be viewed as technically valid, reproducible experimental procedures.

In addition, control experimentation must confront the *uncertainty* paradox of control engineering, which recognizes that the inherent purpose of feedback control is to operate under conditions of uncertainty. Consequently, the level of success of a control technique is ultimately unverifiable due to that very uncertainty. This paradox can be circumvented to some extent in experiments by a combination of high-fidelity identification and “conscious ignorance”; that is, by purposely not exploiting available modeling information (control model versus “truth” model). The success or failure of a given controller implementation may ultimately hinge on effects that are unknown or unknowable.

Moreover, there really is no “truth model”—only the real system with which the controller interacts.

Additional issues of great importance are reproducibility and repeatability. Because of noise and statistical error, no experiment in any sphere of science or engineering is exactly reproducible. Only statistical data analysis can ascertain whether the effects of interest have been reliably captured. In control experiments, effects such as friction may not be repeatable from day to day, let alone among multiple copies of the same experimental setup, or among different setups entirely. This raises the related question as to how one performs meaningful experiments on systems with highly nonrepeatable behavior.

It is also important to stress that a small change to an experimental configuration may entail a substantial change in the control challenge that the experiment presents. For example, relocating a sensor or introducing a delay may render the experiment enormously more challenging. (Try regulating the temperature in your house with the thermostat located outside rather than inside.) Conversely, upgrading an experimental configuration to have more control authority or additional sensors can render the experiment significantly less challenging and possibly less relevant (or sometimes more relevant) to real-world application. Consequently, the exact configuration of a control experiment must be reported for the results to be meaningful. Analogous, but even more subtle, observations apply to the modeling information available to the experimentalist.

In practice, control engineers design plants consciously or unconsciously to minimize these challenges. A control experimentalist must *not* do the same. It is our task to design control experiments that motivate the development of tools and techniques for inherently challenging plant configurations. Although a deliberate attempt to design inherently challenging plant configurations may seem artificial in light of practical control applications, the knowledge gained will inevitably be useful when control practitioners are faced with challenges that cannot be solved by plant redesign or standard control techniques.

It should be clear from this discussion that many aspects of control experimentation, such as instability, underactuation, uncertainty, nonlinearity, and external disturbances, do not present a severe challenge when present in isolation. An obvious example is the inverted pendulum, which is controlled routinely. When more than one of these aspects is present simultaneously, the difficulty of the control problem increases substantially.

Taking a reductionist approach, it may be desirable to design a control experiment to emphasize certain challenges while minimizing or eliminating others. Although such experiments are often viewed as “toys,” their true value, as with traditional experiments in general, is to isolate the effects of interest. However, the true challenge of control engi-

neering may very well lie in addressing multiple difficulties simultaneously, rather than in isolation. This observation suggests holistic rather than reductionist control experiments involving multiple issues simultaneously. The drawback of such experiments, which might be based on scaled-down versions of industrial equipment, is that the presence of multiple control issues invariably obscures the meaning of what has been achieved. This dichotomy requires careful consideration.

In practice, if a control experiment is dangerous, difficult, or expensive to operate, then the amount of available experimentation time will be limited. Unfortunately, such a limitation undermines the reasons for doing experiments in the first place, namely, to try new ideas, refine and motivate new techniques, and benefit from unexpected results arising from real (not virtual) hardware. This kind of experience is rarely obtainable from operational systems, which are often expensive or dangerous to operate in an unconventional manner.

This discussion has focused on system-theoretic control issues rather than domain-specific technology. In many applications, the control challenge may hinge on obtaining a good understanding of the underlying physics and especially the *control physics*; that is, the physical principles that are exploited to effect control over the plant. This understanding is essential for developing effective control strategies, including the design of sensors and actuators with adequate sensitivity and authority. Unfortunately, domain-specific technology development often obscures the underlying system-theoretic issues. The ultimate goal of system-theoretic control experiments is thus to illuminate control system design principles and develop generic techniques that can be applied to a broad range of specific applications, thereby enhancing control technology at the most fundamental levels.

References

- [1] H.W. Bode, *Network Analysis and Feedback Amplifier Design*. New York: Van Nostrand, 1945.
- [2] J.S. Freudenberg and D.P. Looze, *Frequency Domain Properties of Scalar and Multivariable Feedback Systems*. New York: Springer, 1988.
- [3] R.H. Middleton, “Trade-offs in linear control system design,” *Automatica*, vol. 27, pp. 281-292, 1991.
- [4] J. Chen, “Sensitivity integral relations and design tradeoffs in linear multivariable feedback systems,” in *Proc. Amer. Control Conf.*, 1993, pp. 3160-3164.
- [5] M.M. Seron, J.H. Braslavsky, and G.C. Goodwin, *Fundamental Limitations in Filtering and Control*. New York: Springer, 1997.
- [6] J. Hong and D.S. Bernstein, “Bode integral constraints, collocation, and spillover in active noise and vibration control,” *IEEE Trans. Contr. Syst. Technol.*, vol. 6, pp. 111-120, 1998.
- [7] M.T. Hagan, C.D. Latino, E. Misawa, and G. Young, “An interdisciplinary control systems laboratory,” in *Proc. Conf. Control Applications*, Dearborn, MI, 1996, pp. 403-408.
- [8] D.S. Bernstein, “Control experiments and what I learned from them: A personal journey,” *IEEE Contr. Syst. Mag.*, vol. 18, pp. 81-88, Apr. 1998.

- [9] J. Hong, J.C. Akers, R. Venugopal, M.-N. Lee, A.G. Sparks, P.D. Washabaugh, and D.S. Bernstein, "Modeling, identification, and feedback control of noise in an acoustic duct," *IEEE Trans. Contr. Syst. Tech.*, vol. 4, pp. 283-291, 1996.
- [10] R.T. Bupp and D.S. Bernstein, "Experimental implementation of integrator backstepping and passive nonlinear controllers on the RTAC testbed," *Int. J. Robust Nonlinear Contr.*, vol. 8, pp. 435-457, 1998.
- [11] S.L. Lacy, R. Venugopal, and D.S. Bernstein, "ARMAKOV adaptive control of self-excited oscillations of a ducted flame," in *Proc. Conf. Decision and Control*, Tampa, FL, Dec. 1998, pp. 4527-4528.
- [12] J. Ahmed and D.S. Bernstein, "Adaptive control of a dual-axis CMG with an unbalanced rotor," in *Proc. Conf. Decision and Control*, Tampa, FL, Dec. 1998, pp. 4531-4536.
- [13] H. Sane and D.S. Bernstein, "Active noise control using an acoustic servovalve," in *Proc. Amer. Control Conf.*, Philadelphia, PA, June 1998, pp. 2621-2625.
- [14] E.M. Atkins, R.H. Miller, T. Van Pelt, K. Shaw, W.B. Ribbens, P.D. Washabaugh, and D.S. Bernstein, "SOLUS: An autonomous aircraft for flight control and trajectory planning research," in *Proc. Amer. Control Conf.*, Philadelphia, PA, June 1998, pp. 689-693.
- [15] K.-Y. Lum, S.P. Bhat, V.T. Coppola, and D.S. Bernstein, "Adaptive virtual autobalancing for a magnetic rotor with unknown mass imbalance: Theory and experiment," *ASME Trans. J. Vibr. Acoustics*, vol. 120, pp. 557-570, 1998.
- [16] J. Hong and D.S. Bernstein, "Experimental application of direct adaptive control laws for adaptive stabilization and command following," in *Proc. Conf. Decision and Control*, Phoenix, AZ, Dec. 1999, pp. 779-783.
- [17] R. Venugopal and D.S. Bernstein, "Adaptive disturbance rejection using ARMAKOV system representations," *IEEE Trans. Contr. Syst. Tech.*, vol. 8, pp. 257-269, 2000.
- [18] B. Martensson, "The order of any stabilizing regulator is sufficient a priori information for adaptive stabilization," *Syst. Contr. Lett.*, vol. 6, pp. 87-91, 1985.
- [19] A. Ilchmann, *Non-Identifier-Based High-Gain Adaptive Control*. New York: Springer, 1993.
- [20] D.E. Miller and E.J. Davison, "An adaptive controller which provides Lyapunov stability," *IEEE Trans. Automat. Contr.*, vol. 34, pp. 599-609, 1989.
- [21] A.V. Roup and D.S. Bernstein, "Stabilization of a class of nonlinear systems using direct adaptive control," in *Proc. Amer. Control Conf.*, Chicago, IL, June 2000, pp. 3148-3152.
- [22] D.S. Bernstein and A.N. Michel, "A chronological bibliography on saturating actuators," *Int. J. Robust Nonlinear Contr.*, vol. 5, pp. 375-380, 1995.
- [23] R. Koplon, M.L.J. Hautus, and E.D. Sontag, "Observability of linear systems with saturated outputs," *Linear Alg. Applicat.*, vol. 205-206, pp. 909-936, 1994.
- [24] G. Kreisselmeier, "Stabilization of linear systems in the presence of output measurement saturation," *Syst. Contr. Lett.*, vol. 29, pp. 27-30, 1996.
- [25] D.J. Bell and D.H. Jacobson, *Singular Optimal Control Problems*. London: Academic, 1975.
- [26] B. Armstrong-Helouvry, P. Dupont, and C. Canudas de Wit, "A survey of models, analysis tools and compensation methods for the control of machines with friction," *Automatica*, vol. 30, pp. 1083-1138, 1999.
- [27] S.P. Bhat and D.S. Bernstein, "Lyapunov analysis of semistability," in *Proc. Amer. Control Conf.*, San Diego, CA, June 1999, pp. 1608-1612.
- [28] S.L. Lacy, D.S. Bernstein, and S.P. Bhat, "Hysteretic systems and step-convergent semistability," in *Proc. Amer. Control Conf.*, Chicago, IL, June 2000, pp. 4139-4143.
- [29] C.E. Rohrs, L. Valavani, M. Athans, and G. Stein, "Robustness of continuous-time adaptive control algorithms in the presence of unmodeled dynamics," *IEEE Trans. Automat. Contr.*, vol. AC-30, pp. 881-889, 1985.
- [30] S. Sastry, *Nonlinear Systems: Analysis, Stability, and Control*. New York: Springer, 1999.
- [31] E.D. Sontag and Y. Wang, "Lyapunov characterizations of input to output stability," *SIAM J. Contr. Opt.*, vol. 39, pp. 226-249, 2001.

Dennis Bernstein teaches in the Aerospace Engineering Department at the University of Michigan in Ann Arbor. His current research interests include identification and adaptive control with application to aerospace systems. He has been a visiting faculty member at the University of Leeds and the University of Glasgow. His secondary interests include the history of technology, Native American history, and Celtic music.