Control Experiments and What I Learned From Them: A Personal Journey

Dennis S. Bernstein

This article is a personal account of my experiences in developing control experiments for the purpose of control research. The article does not address the important questions surrounding the development of control experiments for undergraduate education. Rather, the emphasis is on research, specifically, the role that control experiments can play in motivating new theoretical ideas. To stimulate discussion about these issues I organized a session for the 1997 American Control Conference entitled "Control Experiments: What Do We Learn From Them?" The reader is invited to peruse the various papers that were contributed to that session for further insights into this question.

Is "Control Experiment" an Oxymoron?

Control is a contradictory subject when it comes to experimentation. On the one hand, if there was ever a subject that cried out for hardware application, it is control. After all, the purpose of control is to *control* something, and real-world applications inspired fundamental developments by Watt, Maxwell, Routh, Nyquist, Bode, Black, and many others. On the other hand, control theory grew up as a branch of applied mathematics in the hands of mathematicians such as Wiener, Bellman, Kalman, and Pontryagin. Although other branches of engineering such as fluid mechanics and structural mechanics are major users of mathematics, none is as mathematical in style and spirit as is control. When was the last time you saw a theorem-proof format in a fluids or structures journal?

Nevertheless, when I left industry and came to the University of Michigan, the desire to actually control something gnawed at me. I saw my colleagues in the Aerospace Engineering Department building all sorts of exotic experiments and I wondered why there was no fundamental need to experiment with anything in control. In the department there had been a hardware tradition in control dating back to the pioneering development of analog simulators by the Gilbert brothers and Robert Howe. These developments, which contributed greatly to aerospace technology during the Apollo years, had long since given way to purely theoretical research. Since I knew researchers in the control community who regularly conducted control experiments (for example, Gary Balas, John Hauser, Carl Nett, Umit Ozguner, and Steve Yurkovich) I was motivated to develop some experiments of my own.

I thought about it at great length, however, and even began to question the very phrase "control experiment." If control results were essentially mathematical statements that were self consistent and provably correct, then it seemed to me that it would be scientifically pointless to build an experiment to test them. In addition, I found it difficult to think of an hypothesis that a control experiment might settle and therefore would warrant testing. I wondered whether the phrase "control experiment" was in fact an oxymoron. While a chemist or fluid dynamicist can run experiments to discover and explore new phenomena (how exciting!), it seemed to me that a controls researcher could at best hope to demonstrate how proficient they were at building hardware that served no other purpose than to mimic the assumptions of some mathematical theorem. Any residual questions could always be addressed by numerical simulation.

Or so it seemed.

There was only one path out of my dilemma. I teamed up with Pete Washabaugh, a structures experimentalist in my department, and he and I drove up to Michigan State to visit Clark Radcliffe and seek his advice. Here was someone with a lab full of interesting control experiments. Clark had his own (and good)

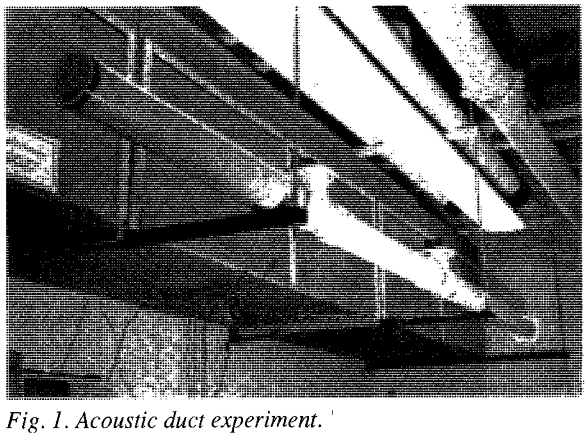

Fig. 1. Acoustic duct experiment.

The author is with the Department of Aerospace Engineering The University of Michigan Ann Arbor, MI 48109, dsbaero@umich.edu

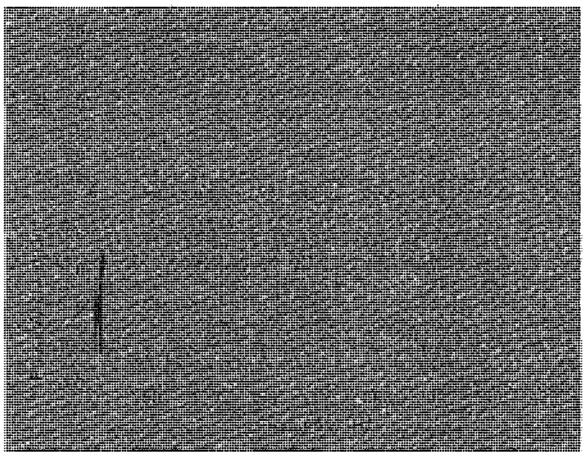

Fig. 2. Identified model of the acoustic duct using ARMAKOV/Toeplitz/ERA identification.

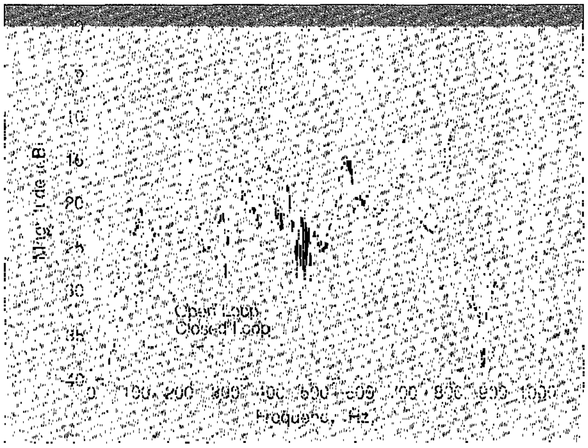

Fig. 3. Open- and closed-loop response of the acoustic duct.

reasons for building control experiments, but I had to learn these for myself.

Like a true experimentalist, Pete immediately suggested all kinds of control experiments. Although I still had nagging doubts, I went along in the hope that my questions would somehow be answered. It seemed to me that the best strategy was to suppress my concerns and proceed at full speed. For the time being I assumed, and eventually learned,

Lesson 1. *In order to understand why control experiments are valuable for control research, you must first do control experiments.*

Being in an aerospace engineering department, it was not difficult to find students who were interested in working on control experiments. In fact, I found that many engineering students thrived on them. I also found that all of those useless skills that were suppressed in my academic career (like working with tools) were suddenly useful. However, I had no idea how I would relate my theoretical work to experiments or where this would lead. Yet I had to press on.

The Acoustic Duct Experiment

Within a day of our visit to Clark's lab and inspired by some noise control experiments we saw there, Pete picked up some four-inch-diameter PVC pipe at Builders Square and some speakers and microphones at Radio Shack, and set up an acoustic duct experiment (see Fig. 1). With a stereo to serve as an amplifier, our investment was less than $200. Finding a suitable control computer was another matter, however, and this held us up for a while. A variety of audio processors and data acquisition boards seemed like they might work, but none was configured for true real-time I/O (many had A/Ds with large delays or had limited buses, for example). Over the years I had heard stories of control experimentalists stymied by the lack of processors that could do real-time control. While I wanted to do control experiments, I did not want to have to design a control computer as well. Unfortunately, it was difficult to find a salesperson who could communicate about computers and control systems; mentioning "*A,B,C*" often led to a blank stare. Pete and I spoke to several vendors including, thanks to Carl Nett's suggestion, dSPACE Inc., which was the only company we found at the time to offer a PC-based real-time control processor board. We later found that a spectrum analyzer was an essential piece of laboratory equipment, as was an assortment of scopes, multimeters, power supplies, amplifiers, and filters.

It turned out that controlling noise in a duct was an excellent first experiment because of its rich dynamics, the availability of good sensors and actuators, and low cost. In fact, an attempt to use high-quality microphones revealed that they had significant phase lag, and we thus returned to the inexpensive variety. I simply had to explain to people that the purpose of our first control experiment was to control a plant made out of, well, *air*. Even this had its advantages, since instability merely resulted in a blown fuse. We bought them by the dozen.

We began our experimental activities by learning to model the acoustic dynamics of the duct. However, it quickly became clear that the dynamics of the speakers and microphones had a major impact on the transfer functions, and it was challenging to account for the electrical-mechanical-acoustic interfaces. Meanwhile, Pete taught me that every sensor, actuator, filter, and amplifier needs to be tested and calibrated since manufacturers' specs are often misleading and are almost always incomplete. The application of mathematical results was preceded by nontrivial effort, as every component needed to be modeled, tested, and verified.

As my students and I learned to build models and as we implemented controllers, a strange process took place. Years of control theory began to assume a new dimension for me. Theoretical concepts such as gain and phase margin, poles and zeros, and sensitivity were no longer abstractions. The students measured not only the closed-loop response, but also the loop gain and gain margin in order to understand the interaction between the plant and the controller. Although gain margin and sensitivity were invisible and abstract to "non-control" visitors to the lab, we were gaining firsthand experi-

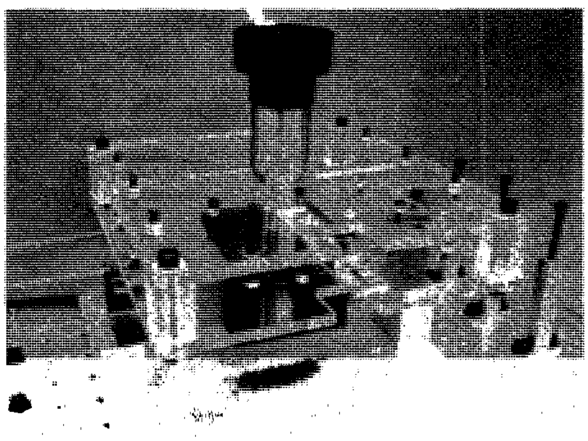

Fig. 4. Rotational/translational actuator (RTAC) experiment.

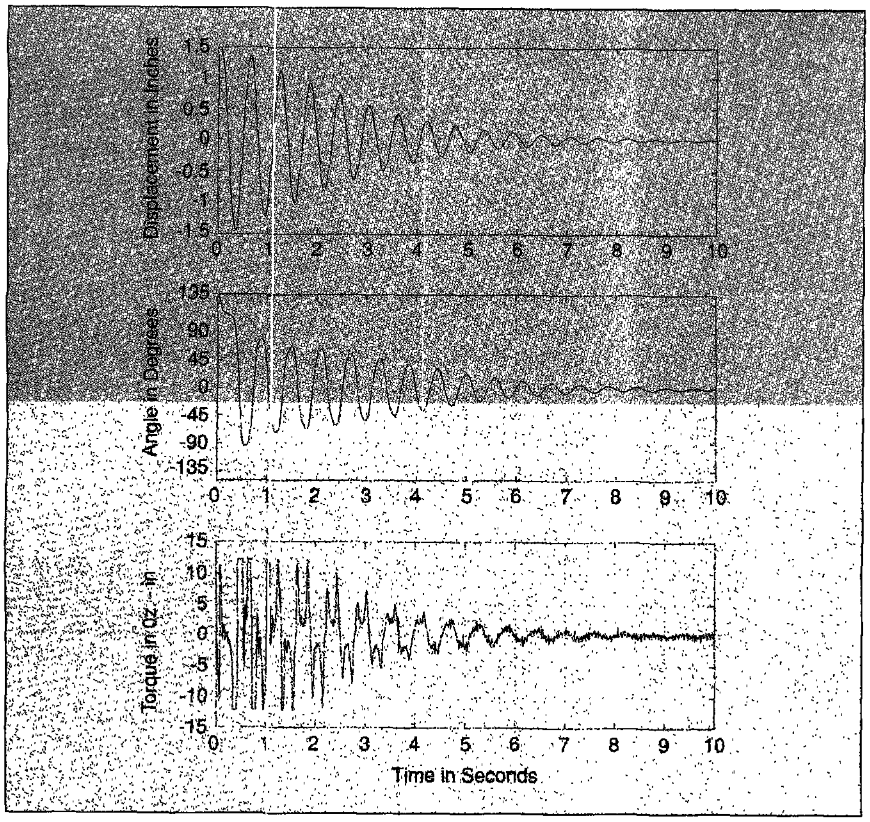

Fig. 5. Integrator backstepping control of the RTAC with saturation.

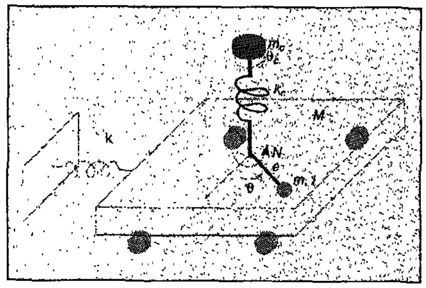

Fig. 6. Virtual absorber subsystem for controlling the RTAC.

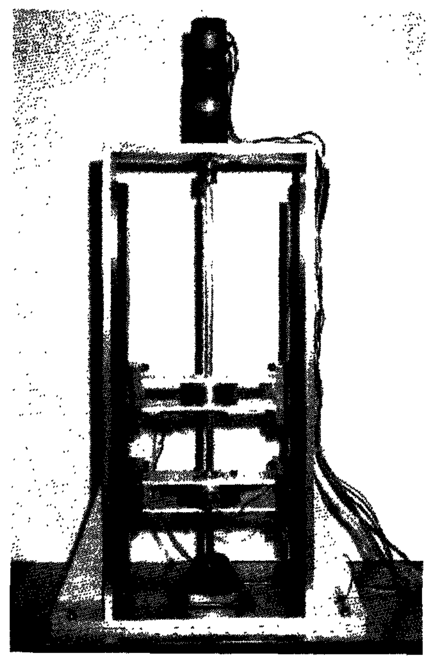

Fig. 7. Unbalanced rotating shaft experiment.

ence with their existence and meaning. We relied on these concepts to get controllers to work.

But the honeymoon was soon over as we began to notice strange phenomena that the textbooks mentioned briefly but ominously. For example, when we collected data at different disturbance levels, the plant transfer function wasn't the same. At frequencies for which the plant gain was low, such as near zeros, we couldn't even collect good data. The perfect Nyquist plots of the loop transfer function I was used to seeing in textbooks (especially those fanciful ones that wrapped around at infinity, not to mention those unobtainable ones corresponding to unstable loop transfer functions) just didn't exist in the lab. The notion of poles and zeros got fuzzier and fuzzier, and those perfect root locuses couldn't be found either. Relative degree became suspect, and I began to question my faith in rational functions. As we began to recognize nonlinearities and realized that noise was everywhere, things reached crisis proportions. The self-consistent theorems of control theory didn't seem as powerful as they used to seem. The real world was an amazingly messy place, and our ability to probe it was impeded by a fundamental fact of life given by

Lesson 2. *All real data is finite and noisy.*

Through all of this, I developed a new feeling for the meaning of an "assumption." While an assumption in mathematics always means an unquestioned axiom, an assumption in the physical world serves as an *approximation to reality*. Expertise in an area of physics or engineering is needed to determine the realism and accuracy of any given assumption. I learned to accept the fact that no mathematical assumption, whether it is of a deterministic or stochastic nature, is ever satisfied in the real world, while Lesson 2 taught me that the ability to verify the validity of any assumption is inherently limited.

Nevertheless, we eventually developed a sense of the limits to linear modeling and proceeded to construct linear dynamical models as the basis for active noise control [1, 2]. However, we soon realized that because of sensor and actuator dynamics, analytical modeling could not be trusted to provide reliable models. We therefore learned to obtain useful models (nominal plus uncertainty) of the duct dynamics using identification techniques, namely, the inverse FRF/ERA and ARMARKOV/Toeplitz/ERA methods [3-5] (see Fig. 2). With the effectiveness of these techniques and the difficulties we experienced trying to obtain precise analytical models, we learned

Lesson 3. *An ounce of identification is worth 10 pounds of modeling.*

Since acoustic dynamics are inherently linear and have high modal density, they provide an ideal testbed for linear, robust control [6-13]. Our first attempt was to colocate the measurement sensor and control actuator in order to exploit inherent stability robustness. However, this sensor/actuator arrangement led to *spillover*, which appeared as amplification of the open-loop gain by the controller in a given frequency range. (Here we are referring to *spectral* spillover, although *spatial* spillover occurs as well.) In fact, the experiment immediately focused our attention on this phenomenon for the simple reason that we *heard* it. In addition, since spillover occurred no matter

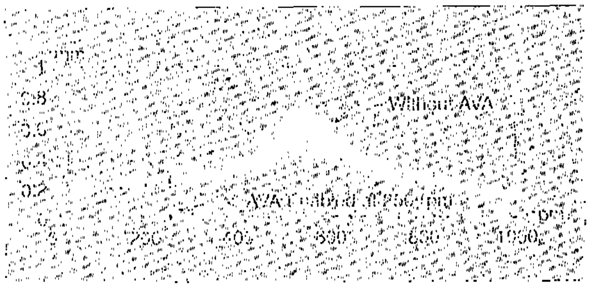

Fig. 8. Adaptive virtual autobalancing control of the unbalanced rotating shaft experiment during spinup.

Fig. 9. CAD drawing of the control-moment-gyro experiment.

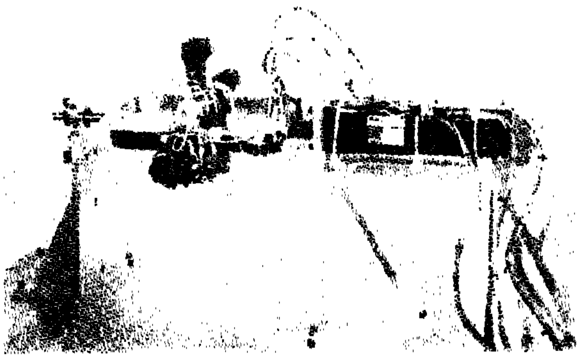

Fig. 10. Control-moment-gyro experiment.

how accurately we identified the plant, we concluded that spillover was not a consequence of uncertainty. We soon realized that the Bode integral constraint on sensitivity was at work here and that appropriate placement of sensors and actuators was essential [14]. While the feedforward noise control community had known this for a long time, we were forced to learn it from a feedback perspective. By exploiting classical results from singular LQG control that depend upon nonminimum phase transfer functions, we saw the spillover disappear (see Fig. 3). We had learned

Lesson 4. *Control experiments often focus attention on performance and implementation issues that are overlooked and difficult to capture in numerical simulation.*

The RTAC Experiment

The next experiment we built was an attempt to rein in Mother Nature. The RTAC (rotational/translational actuator) is a mechanical device with two degrees of freedom, namely, a translating oscillator and a rotational motor with a mass mounted on an eccentric arm [15] (see Fig. 4).

For the RTAC, we designed and implemented dissipative controllers, which require only the angular velocity of the arm [16, 17], as well as integrator backstepping controllers, which require full-state feedback [18]. Experiments focused our attention on two issues. First, the integrator backstepping controllers required far larger control inputs than the dissipative controllers, which thus presented saturation difficulties (see Fig. 5). However, the dissipative controllers, which are based upon the reaction of the arm to the oscillator motion, cease to be effective at low oscillation amplitudes due to stiction. On the other hand, the integrator backstepping controllers, which are full-state feedback controllers, react to the translational motion of the cart and thus are more effective at low amplitudes. By observing the effects of saturation and stiction on the various control algorithms, we learned

Lesson 5. *A control experiment can reveal whether the mathematical assumptions of control theory are realistic and can help identify which physical effects are important.*

Later, we learned to control the RTAC by manipulating the flow of energy between the oscillator and the arm as well as into and out of the plant. To do this, we developed control strategies involving virtual absorbers which are reset at various times, thereby instantaneously removing energy from the system (see Fig. 6). Although resetting a virtual absorber by zeroing out computer states would appear to have no effect on the *real* energy of the system, in fact, the true effect of the resetting procedure is to prevent the control system from *reintroducing* energy into the plant [19, 20]. Only by observing this effect in the lab did we convince ourselves that this was a viable control strategy. The development of resetting virtual absorbers was a consequence of

Lesson 6. *Control experiments can suggest new research problems and directions as well as new control approaches.*

The Rotating Shaft Experiment

Next we built an experiment to suppress vibrations due to an unbalanced rotor. This is a universal problem in rotating machinery where rotating masses induce vibrations due to imbalance. Previously, we had developed theoretical results for controlling the spinning top, first for the symmetric (balanced) case [21, 22] and then for the asymmetric (unbalanced) case [23]. To develop a realizable control experiment we needed a suitable actuator, and we were fortunate that Brad Paden provided us with a magnetic bearing. We then designed the rotating shaft experiment to allow us to implement control algorithms for counteracting the effects of mass imbalance. In our experimental setup we mounted the shaft vertically to avoid the need for shaft levitation in order to focus on imbalance compensation (see Fig. 7). In developing viable controllers for imbalance compensation we realized that accurate measurements of the inertia matrix of the shaft were extremely difficult to obtain. In fact, any realistic control strategy for compensating rotor imbalance must be effective in the presence of unknown and possibly changing inertia. The strategy we adopted was physically motivated and sought to emulate the motion of passive weights confined to a fluid-filled annulus surrounding the shaft. While mechanical devices based on this principle have been known since the 1930s, *adaptive virtual autobalancing* [24, 25] was an attempt to capture this idea as an active control algorithm. Experimentally, the algorithm worked successfully (see Fig. 8), while analysis showed (rather surprisingly) that the pas-

As the article suggests, my original concentration as an undergraduate was in mathematics, specifically, applied mathematics at Brown University. There were several reasons for this choice, not the least of which was that Brown University has an excellent program in applied mathematics. However, equally relevant was the fact that while I knew what mathematics was, having had calculus in high school, I had had no exposure to engineering as a potential academic major. I was fortunate to discover controls while I was an undergraduate at Brown, where I enrolled in a graduate-level control course taught by doctoral student Panos Antsaklis. I eventually transitioned over to engineering through the Computer, Information, and Control Engineering Program at the University of Michigan, where my advisor was Elmer Gilbert. The history of control engineering is itself a study in the interaction of applications and theory. I am excited by the possibility that control engineering can benefit immensely from the interaction of theory and experiment. Experiments can emphasize implementation issues, subject assumptions to hidden effects, and guide theoretical effort. Experiments can bring back the joy of serendipity, the unexpected discovery that pervades science and engineering but which is often lacking in the mathematical world of control theory. Finally, experiments can aid in the transition of new control ideas and technology to applications. For more details and a picture, see D.S. Bernstein, "A Student's Guide to Classical Control," *IEEE Contr. Sys. Mag.*, vol. 17, pp. 96-100, August 1997.

—Dennis Bernstein

sive mechanical device being emulated can be viewed as the embodiment of an internal model controller. This experiment reinforced Lesson 6.

A challenging aspect of the rotating shaft experiment is the fact that the magnetic bearing actuator has a permanent magnet bias which has the tendency to snap the shaft from one side to the other unless there is a minimal level of control authority to effect stabilization. To counteract this instability without a physical stiffness (which we included for this reason at the base of the shaft) we needed a good actuator model, which was difficult to obtain empirically in the presence of the instability. A deliberate sign change in the feedback loop with the shaft spinning at 1,000 rpm causes the shaft to ricochet violently and gives a graphic demonstration of the consequences of instability. This lesson had been stressed by Gunter Stein in the title of his classic Bode Lecture which taught us

Lesson 7. *Respect the unstable.*

The Actively Controlled Control-Moment Gyro

Since nonlinearities arose whenever we least expected them, we decided to build an experiment that was intentionally nonlinear and had more degrees of freedom than the RTAC. This objective led to the design of an actively controlled control-moment gyro (CMG), which consists of three rigid bodies whose dynamics involve large-angle nonlinear gyroscopic effects [26]. Attached to a spacecraft and with a rapidly spinning rotor, a CMG provides stiffness for the spacecraft and, by applying torques to the gimbals, can be used to slew the spacecraft. The CMG involves an outer gimbal and an inner gimbal, both of which are actuated by motors. The inner gimbal is controlled by a pair of matching motors to double the available torque and to balance the motor mass. Attached to the inner gimbal is a fourth motor that drives a rotor. For control experiments, we can attach various rotors to perform a wide range of control experiments involving slewing of the rotor and imbalance compensation. We can also invert the outer gimbal to stabilize a rotating spherical pendulum. While a CAD drawing illustrates the basic design (see Fig. 9), a photo of the actual testbed shows added mass and stiffness due to cabling, connectors, and other hardware (see Fig. 10). These considerations, as well as effects such as stiction, emphasize

Lesson 8. *The need for nonlinear identification is pervasive.*

While nonlinear identification is essential due to both large-angle nonlinearities and modeling uncertainty, the CMG experiment motivated us to develop an attitude control technique for spacecraft tracking that is adaptive with respect to inertia [27]. To avoid singularities in an Euler angle representation, the approach in [27] is based upon quaternions to represent attitude.

Active Control of Combustion

As a much more challenging control experiment, we undertook the problem of controlling combustion instability. It turned out that the combustion program in the Aerospace Engineering Department involved research on a 300 kW natural gas combustor (see Fig. 11) which exhibits thermo-acoustic instabilities in the form of loud rumbling when operated in certain regimes. This combustor was available for a two-month period, and we were allowed six weeks to try to control it.

Our approach to controlling the combustor was to use speakers to counteract the instability caused by the interaction of the flame and the acoustic dynamics. The trick was to insert acoustic energy sufficiently close to the flame in order to achieve the greatest possible advantage (and without melting the speaker!). To do this we exhausted numerous strategies, including (carefully) inserting a speaker directly into the natural gas fuel line (see Fig. 12). Ultimately, we realized that the design of the combustor did not accom-

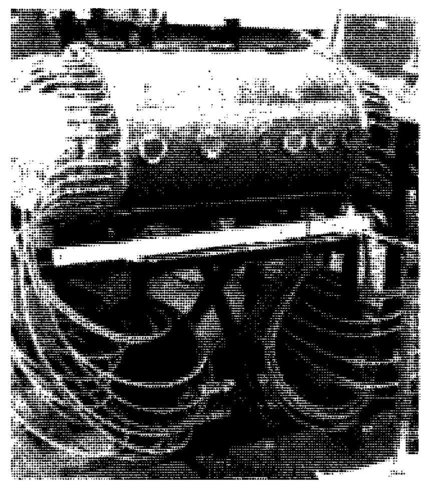

Fig. 11. 300 kW natural gas combustor.

Fig. 12. Speaker housing attached to the natural gas fuel line of the combustor.

modate acoustic control. What we really needed were fast servovalves which would entail actuator development. We fully appreciated

Lesson 9. *Control experiments allow one to practice the "outer loop" of control design, namely, the specification, design, and implementation of sensors and actuators.*

This lesson taught us that "off-the-shelf" control experiments deprive experimentalists of one of the most important aspects of control engineering.

Whereas for the acoustic duct experiment we tested controllers with machine-generated disturbance signals, in the world of fluids and flames Mother Nature creates the disturbance through complex dynamics. In this case one cannot count on mathematical assumptions such as stationarity, Gaussian, etc., to hold. This taught us

Lesson 10. *The difference between a "toy" control experiment and a "real" control experiment is whether the disturbance is of your own construction or is thrown at you by Mother Nature herself.*

Adaptive Control Experiments

Our experience with the combustor motivated us to focus on adaptive control techniques that would work in the presence of poor plant models and unknown disturbance spectra. While control technology has had a long and successful history in electrical, mechanical, aerospace, and other technological applications, adaptive feedforward algorithms were developed independently for noise cancellation. These techniques include LMS (least mean square) algorithms with FIR and IIR controllers, lattice filter techniques, and numerous variants. The theoretical foundation for these techniques varies greatly from method to method, as does their performance in practice. In contrast to classical feedback techniques, adaptive feedforward algorithms have limited modeling requirements and are robust to disturbance spectrum uncertainty. These features have been exploited in applications with good success.

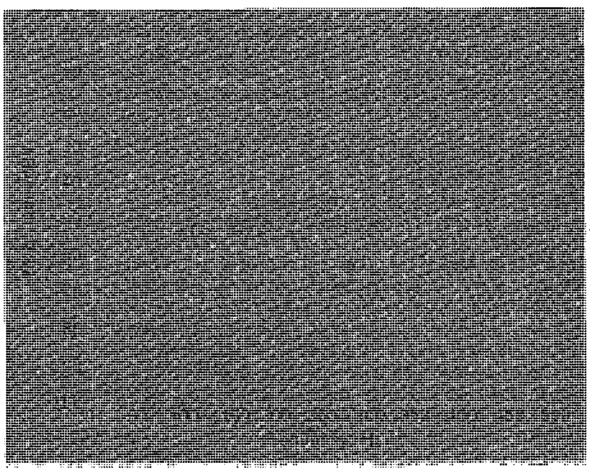

Fig. 13. ARMARKOV/Toeplitz adaptive cancellation with dual-tone disturbance.

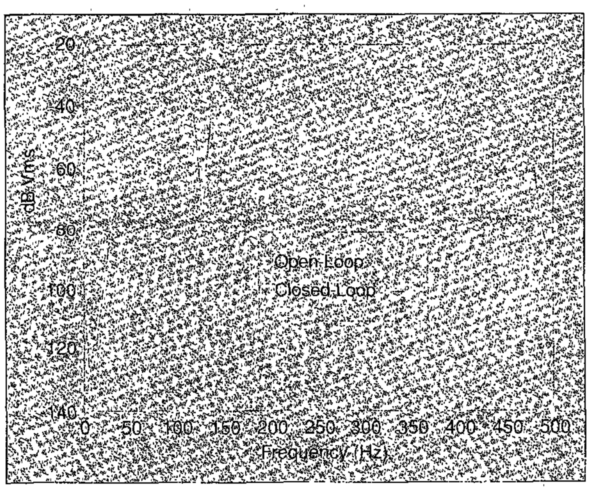

Fig. 14. ARMARKOV/Toeplitz adaptive cancellation with white-noise disturbance.

Although a rigorous theoretical foundation for adaptive cancellation algorithms is often lacking, experimental implementation of these algorithms can be used to assess their effectiveness. We therefore implemented various adaptive cancellation algorithms along with an ARMARKOV/Toeplitz-based algorithm [28-30]. This algorithm converged in the presence of single-tone, dual-tone, and broadband disturbances. Figs. 13 and 14 show the open-loop and converged closed-loop response of the ARMARKOV/Toeplitz-based algorithm for dual-tone and white-noise disturbances. Note the presence of harmonic overtones due to speaker stiffness nonlinearity, which emphasizes Lesson 5 and Lesson 8. In keeping with Lesson 10, this algorithm was also tested with random noise generated by an AM radio tuner set between stations (see Fig. 15).

Adaptive control algorithms are exciting to observe in the lab. The ideal controller would work with an initially poor model, learn and improve with age, and change when the plant changes. That is the Holy Grail of adaptive control. However, since adaptive controllers change in response to changing disturbances and plant dynamics, their behavior and reliability is difficult to ascertain by means of theory alone. Control experiments thus provide a convenient means for testing variants of adaptive control algorithms. On numerous occasions we learned

Lesson 11. *Control experiments provide a quick way to identify control methods that seem to work under real-world conditions as well as those that clearly don't.*

Thus, in a truly experimental spirit, control experiments are useful for discovering promising new algorithms. In a similar vein, we also learned

Lesson 12. *Control methods based on rigorous theory may fail for unknown reasons, while heuristic control methods may work for equally unexplained reasons.*

Both of these lessons motivate theoretical research to explain both unexpected failures and unexpected successes. In any event, our experience with hardware taught us that control experiments are an effective arbiter of whether an adaptive control method will work under real-world conditions.

Why Do Control Experiments?

The above discussion does not in any way reflect the hard work required to design, build, and operate a control experiment for control-systems research. Design of a control experiment is an iterative process that depends upon extensive analysis to size and select appropriate components. In addition, the reliable operation of the hardware components as well as all of the supporting real-time software (an often-underemphasized aspect of the control curriculum) can be a major, time-consuming task.

Ironically, although a control experiment can take months or years to build and render operational, reporting experimental results may occupy only a small fraction of a research paper, with the theoretical portion receiving top billing. However, as noted in Lesson 11, a working control experiment has the ability to reveal very quickly which control approaches are promising and which are not, thus suggesting the most fruitful research directions. *I believe that this guidance is of inestimable value to control research and technology.* In addition, we found that control experiments invariably motivated the development of new control algorithms and techniques.

It is fair to say that control as an experimental science has had far too little emphasis. Control research can be enriched in innumerable ways by proper emphasis on control experiments. A control experiment can bring out important physical phenomena that a theorist would not think of considering. We live in an excellent time for undertaking control experiments, especially because of fast processors for real-time control. Let Mother Nature be our teacher!

Finally, the most profound lesson I learned was

Lesson 13. *Control research without experiments is like music without sound.*

Acknowledgments

I learned a tremendous amount by working with students and colleagues in developing control experiments. Undoubtedly, there are many in the control community who have learned similar lessons in their own experimental activities. The goal in this account has been to help guide and encourage others who may be contemplating such activities.

These experiments have influenced my theoretical research and thinking in ways I never could have imagined. I developed an appreciation for and an interest in other areas of aerospace engineering, including structures, fluids, and combustion. I developed a rapport with my colleagues who devote their careers to these fields of research.

I am deeply indebted to my colleagues Pete Washabaugh and Vince Coppola and to our students Jasim Ahmed, Jim Akers, Dave Atkins, Sanjay Bhat, Robert Bupp, Scott Erwin, Jeongho Hong, Kai-Yew Lum, Robert Miller, Scot Osburn, Andy Sparks, Feng Tyan, Tobin Van Pelt, Ravi Venugopal, and C.-J. Wan, with whom I worked for the past six years at the University of Michigan on these control experiments. I especially thank Jim Akers for his tremendous commitment to developing the Noise and Vibration Control Laboratory in the Aerospace Engineering Department.

All opinions expressed are those of the author. I am grateful for research support provided by AFOSR, NSF, and the University of Michigan Office of the Vice President for Research.

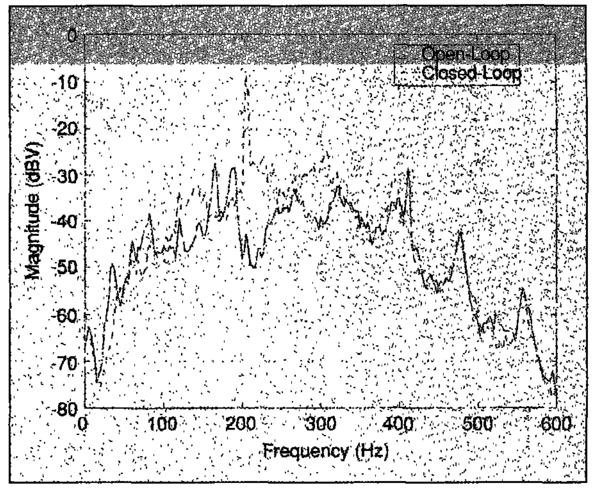

Fig. 15. ARMARKOV/Toeplitz adaptive cancellation with AM radio noise disturbance

References

[1] 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. Sys. Tech.*, vol. 4, pp. 283-291, 1996.

[2] R. Venugopal and D.S. Bernstein, "State Space Modeling of an Acoustic Duct with an End-Mounted Speaker," *Proc. Conf. Contr. Appl.*, pp. 954-959, Dearborn, MI, September 1996; *J. Vibr. Acoustics*, 1998.

[3] J.C. Akers and D.S. Bernstein, "Measurement Noise Error Bounds for the Eigensystem Realization Algorithm," *Proc. Amer. Contr. Conf.*, pp. 2566-2570, Seattle, WA, June 1995.

[4] J.C. Akers and D.S. Bernstein, "An Adaptive Toeplitz/ERA Identification Algorithm," *Proc. 37th Structures, Structural Dynamics, and Materials Conference*, pp. 1024-1034, Salt Lake City, UT, April 1996.

[5] J.C. Akers and D.S. Bernstein, "Time-Domain Identification Using ARMARKOV/ Toeplitz Models," *Proc. Amer. Contr. Conf.*, pp. 191-195, Albuquerque, NM, June 1997.

[6] J.C. Akers and D.S. Bernstein, "ARMARKOV Least-Squares Identification," *Proc. Amer. Contr. Conf.*, pp. 186-190, Albuquerque, NM, June 1997.

[7] W.M. Haddad and D.S. Bernstein, "Parameter-Dependent Lyapunov Functions and the Popov Criterion in Robust Analysis and Synthesis," *IEEE Trans. Autom. Contr.*, vol. 40, pp. 536-543, 1995.

[8] A.G. Sparks and D.S. Bernstein, "Real Structured Singular Value Synthesis Using the Scaled Popov Criterion," *AIAA J. Guid. Contr. Dyn.*, vol. 18, pp. 1244-1252, 1995.

[9] Y. Ge, L.T. Watson, E.G. Collins Jr., and D.S. Bernstein, "Probability-One Homotopy Algorithms for Full and Reduced Order H_2/H-Infinity Controller Synthesis," *Optimal Contr. Appl. Meth.*, vol. 17, pp. 187-208, 1996.

[10] F. Tyan and D.S. Bernstein, "Shifted Quadratic Guaranteed Cost Bounds for Robust Stability and Performance," *Proc. 13th IFAC World Congress, vol. G: Education, Robust Control I*, pp. 285-290, San Francisco, CA, July 1996.

[11] R.S. Erwin, A.G. Sparks, and D.S. Bernstein, "Fixed-Structure Robust Controller Synthesis via Decentralized Static Output Feedback," *Int. J. Robust Nonlinear Contr.*, 1998.

[12] W.M. Haddad, V.-S. Chellaboina, and D.S. Bernstein, "An Implicit Small Gain Condition and an Upper Bound for the Real Structured Singular Value," *Sys. Contr. Lett.*, vol. 29, pp. 197-205, 1997.

[13] R.S. Erwin and D.S. Bernstein, "Discrete-Time H_2/H-Infinity Control of an Acoustic Duct: δ-Domain Design and Experimental Results," *Proc. IEEE Conf. Dec. Contr.*, San Diego, CA, December 1997.

[14] J. Hong and D.S. Bernstein, "Bode Integral Constraints, Colocation, and Spillover in Active Noise and Vibration Control," *IEEE Trans. Contr. Sys. Tech.*, vol. 6, pp. 111-120, 1998.

[15] C.-J. Wan, D.S. Bernstein, and V.T. Coppola, "Global Stabilization of the Oscillating Eccentric Rotor," *Nonlinear Dynamics*, vol. 10, pp. 49-62, 1996.

[16] R.T. Bupp and D.S. Bernstein, "A Benchmark Problem for Nonlinear Control Design: Problem Statement, Experimental Testbed, and Passive Nonlinear Compensation," *Proc. Amer. Contr. Conf.*, pp. 4363-4367, Seattle, WA, June 1995.

[17] R.T. Bupp and D.S. Bernstein, "Experimental Comparison of Passive Nonlinear Controllers for the RTAC Testbed," *Proc. Conf. Contr. Appl.*, pp. 279-284, Dearborn, MI, September 1996.

[18] 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.*, 1998.

[19] R.T. Bupp, D.S. Bernstein, V.-S. Chellaboina, and W.M. Haddad, "Finite-Time Stabiliza-

tion of the Double Integrator Using a Virtual Trap-Door Absorber," *Proc. Conf. Contr. Appl.*, pp. 179-184, Dearborn, MI, September 1996.

[20] R.T. Bupp, D.S. Bernstein, V.-S. Chellaboina, and W.M. Haddad, "Resetting Virtual Absorbers for Vibration Control," *Proc. Amer. Contr. Conf.*, pp. 2647-2651, Albuquerque, NM, June 1997. *J. Vibr. Contr.*, to appear.

[21] C.-J. Wan, V.T. Coppola, and D.S. Bernstein, "Global Asymptotic Stabilization of the Spinning Top," *Optimal Contr. Appl. Meth.*, vol. 16, pp. 189-215, 1995.

[22] C.-J. Wan, P. Tsiotras, V.T. Coppola, and D.S. Bernstein, "Global Stabilization of the Spinning Top Using Stereographic Projection," *Dynamics and Control*, vol. 7, pp. 215-233, 1997.

[23] K.-Y. Lum, D.S. Bernstein, and V.T. Coppola "Global Stabilization of the Spinning Top with Mass Imbalance," *Dynamics and Stability of Systems*, vol. 10, pp. 339-365, 1995.

[24] K.-Y. Lum, V.T. Coppola, and D.S. Bernstein, "Adaptive Autocentering Control for an Active Magnetic Bearing Supporting a Rotor with Unknown Mass Imbalance," *IEEE Trans. Contr. Sys. Tech.*, vol. 4, pp. 587-597, 1996.

[25] 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*, 1998.

[26] J. Ahmed, R.H. Miller, E.H. Hoopman, V.T. Coppola, T. Andrusiak, D. Acton, and D.S. Bernstein, "An Actively Controlled Control Moment Gyro/GyroPendulum Testbed," *Proc. Conf. Contr. Appl.*, pp. 250-252, Hartford, CT, October 1997.

[27] J. Ahmed, V.T. Coppola, and D.S. Bernstein, "Asymptotic Tracking of Spacecraft Attitude Motion with Inertia Identification," *Proc. IEEE Conf. Dec. Contr.*, San Diego, CA, December 1997, pp. 2471-2476.

[28] R. Venugopal and D.S. Bernstein, "Adaptive Disturbance Rejection Using ARMARKOV/Toeplitz Models," *Proc. Amer. Contr. Conf.*, pp. 1657-1661, Albuquerque, NM, June 1997.

[29] T. Van Pelt, R. Venugopal, and D.S. Bernstein, "Experimental Comparison of Adaptive Cancellation Algorithms for Active Noise Control," *Proc. Conf. Contr Appl.*, pp. 559-564, Hartford, CT, October 1997.

[30] R. Venugopal and D.S. Bernstein, "Adaptive Disturbance Rejection and Model Reference Adaptive Control Using ARMARKOV System Representations," *Proc. IEEE Conf. Dec. Contr.*, San Diego, CA, December 1997, pp. 1884-1889.