The Theory/ Practice Gap
special section

On Bridging the
Theory/Practice Gap

Dennis S. Bernstein

> "Mind the gap. Mind the gap."
>
> — *Synthesized warning heard in the London Underground.*

In a recent survey [1], 27 researchers provided inputs concerning the state of research in systems and control and were asked to give their opinions on the major challenges facing the control community. Among the challenges identified was the need to bridge the gap between theory and practice.

In this article I speculate on some reasons for the existence of the gap and provide concrete suggestions for bridging it. Specifically, I am interested in the following questions:

1. What is the evidence for the existence of the gap?
2. What is the extent of the gap?
3. What is the significance of the gap for systems and control research?
4. What factors have contributed to the gap?
5. What technical research problems are pertinent to bridging the gap?
6. Why is it important to bridge the gap?

First, a few disclaimers are in order. My perspective on these questions is from aerospace engineering and reflects my experiences in academia and industry. Furthermore, this article is not intended to be either a defense or a critique of "academic" research in control technology for aerospace engineering or any other branch of engineering.

According to my dictionary, "academic" means "very learned but inexperienced in the world of practical reality." This is not an accurate description of academic researchers. Although all but one of the survey respondents [1] hold positions in academia, many have extensive industrial and government experience. Furthermore, several of the respondents in [1] are strongly in favor of aggressive efforts to bridge the gap. I am encouraged by their views to present some concrete suggestions of my own.

Finally, the views I present are intended to suggest how the academic side might contribute to bridging the gap between theory and practice. However, there is much the industrial side can do as well to solve this problem. I believe there is a corresponding burden on control practitioners to articulate their needs and provide guidance and feedback to the research community.

Evidence for the Gap

I will comment only briefly on the first three questions. Although many control researchers and practitioners would probably agree that a theory/practice gap exists, the extent of the gap is the subject of much debate and can only be estimated by anecdotal evidence. Characterizing and quantifying this gap would require nontrivial effort and is beyond the scope of this article. Although there are researchers who are quick to point out successes of modern systems and control theory research, I personally believe that the gap on the whole is large and warrants serious introspection by the research community.

The significance of the gap for systems and control research is a complex and subtle issue. Here I note that (1) basic research has always benefited from the influence of applications, while (2) in the long run, the most important developments are those of

---

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

0272-1708/99/$10.00©1999IEEE

basic research that have deep and long-lasting influence, leading to the conclusion that (3) a balance between theory and applications is essential. Unfortunately, the time constants for basic research and useful application are generally quite different. In fact, new ideas can require a long time from conception to exploitation, while the time pressures of applications can divert attention from potentially valuable long-term solutions.

Next, I'll discuss some factors that have contributed to the gap's existence (question 4). Then I'll examine various systems and control issues and their relevance to bridging the gap (question 5). Finally, I'll end by commenting on the importance of closing the gap (question 6).

Why Does the Gap Exist?

To shed some light on the existence of the gap, it is worthwhile to speculate on some of the reasons for its existence.

1. Do control engineers *need* modern systems and control theory?

Generally, new technology is used in practice only when there is a clear cost or performance benefit. There are two sides to technology development, namely, technology push and market pull, where "market" can refer to either commercial or military/government applications. In aerospace technology, the latter is usually the technology driver. Although necessity is often called the mother of invention, the extent to which necessity contributes to bringing about truly new developments is a murky question. (Some inventors believe that necessity is actually the *daughter* of invention [2], p. 125.) Despite clear needs, new technologies often require a generation (or more) to mature to the point of practical usefulness. The reality is surely a combination of both push and pull.

I personally place the pointer to the side of technology push: Engineers often invent the possible and look for opportunities they can develop and exploit. The most important and fundamental developments are often inwardly motivated. After all, the Wright brothers were not funded by the Air Force, and there was no pressing need for manned flight. Nor was the transistor developed out of necessity. In fact, its potential wasn't even recognized at first in the U.S. (although it was in Japan). As another example, the mathematician G. Hardy took satisfaction in the uselessness of number theory, yet few technologies are hotter today than data encryption. On the other hand, some of the most important conceptual developments have arisen from the desire to address real problems. Fourier analysis arose from heat flow problems for which Fourier was ridiculed, and thermodynamics arose from the boring of cannon [3]. Yet physicists are often called on to contribute to revolutionary technologies such as radar and nuclear weapons.

In control engineering practice, "need" is a nebulous concept. Rarely do engineers consciously design and build a system that truly cannot be controlled with existing concepts. In many applications, the plant can be redesigned and additional sensors and actuators can be implemented to render the control problem more manageable. Fully actuated mechanical systems are a good example of this kind of design. Other brute force solutions can be devised as well. For example, performance specifications can be lowered, and hardware and manpower budgets can be increased (or, more commonly, projects can be canceled), thus eliminating opportunities for new and innovative, and therefore risky, ideas. At the same time, the opportunities afforded by conceptual advances in control are difficult to grasp. It can be difficult for researchers to demonstrate and quantify the advantages of a new idea in control where the design process is complex and there are numerous tradeoffs that interact in intricate ways.

2. What are the risks of new control methods?

In many applications, especially in aerospace, the control system is critical to system operation. This is a two-edged sword. On the one hand, control system technology is essential and therefore will command high priority when it is required. On the other hand, if the control system fails, the system may be lost, and thus applications in which human lives or great cost is at stake call for well-tested methods over novel techniques. Risk reinforces the inertia associated with the acceptance of new control technology.

Nevertheless, there are many potential applications of control technology that are not easily designed away. These include flight control with unusual configurations (canted tail fins, oblique wings, tailless aircraft) and nontraditional control applications (compressor and combustor control, active vibration control, high angle of attack flight). However, each of these applications entails risk and thus, despite potential cost and performance advantages, must undergo extensive development before it can be transitioned to practice.

3. How do publications contribute to the gap?

Bridging the gap, from the researcher's perspective, requires that new ideas be communicated to engineers who are in a position to apply them. I'll focus on publishing, where there are three main avenues, namely, journal papers, conference papers, and books.

Much effort in academia is devoted to publishing in journals, which are strictly limited to demonstrably new ideas. Consequently, journal papers are extremely terse and are generally written for other researchers, not practitioners. Authors of journal papers are rarely given much space to provide background and self-contained, pedagogical exposition that could render their papers more readable by practitioners and nonspecialists. This avenue is also slow, typically taking three years from submittal to publication.

Conference publications are more timely but are even more terse than journal papers (they often serve as summaries of results), and they are not widely available to engineers who do not attend the conference. However, CD ROM proceedings are much more accessible than bulky hard-copy versions and should allow larger distribution. I personally hope that once paper proceedings are eliminated, authors can be allowed more publishing space to develop new ideas.

Books provide an opportunity for researchers to expand and illustrate new concepts at great length, which is not possible in journals. Few books written by academic researchers, however, consider sufficient engineering detail to be directly usable to practicing engineers. Some "how to" control books are available on specialized topics such as motor servos, but few of them truly advance the state of the art in control practice (nor is it their intent to do so).

4. How can publications be rendered more usable?

A control engineer considering a new control algorithm for a potential application needs to know the answers to the following questions:

1. What problems does the control algorithm address (stabilization, disturbance rejection, tracking, etc.)?

2. To what class of plants does the algorithm apply?

3. What modeling information is needed to design and tune the controller?

4. What is the structure of the control algorithm?

5. How is the controller tuned?

6. What sensor/processor/actuator hardware is needed to implement the controller?

A satisfactory control paper will provide answers to all these questions in a clear and accessible manner. If this information is buried in the paper, the reader may have to expend considerable effort and thus may become discouraged from further considering the method. For example, if the control law is given in transformed variables, the potential user must unravel the transformations to determine the roles of the various parameters. Although this information may be a detail to the writer, it is vital to the practitioner.

As a further example, a practitioner may be interested in a special case of a general procedure. The reader may not have time to recognize the applicability of the procedure or to understand how the result specializes to useful, special cases. I personally like the "closed-loop" or "sandwich" model of paper writing: concrete, abstract, concrete. That is, motivate the paper with a concrete or specific problem, work out the theory in an abstract or general context, and, finally, return to the concrete or specific problem. In general, a paper is most useful to a control practitioner when it provides and demonstrates operational procedures for implementing the methods under conditions that are explicitly stated.

The following remarks are intended to emphasize the relevance of various topics and issues that I believe have significance for bridging the gap between theory and practice. Many of these topics and issues have been extensively studied by the systems and control community, while others have not. My objective is to emphasize those aspects that may have some bearing on the gap.

Bridging the Gap: Modeling Issues

5. Don't trivialize stabilization.

Although it seems trivial to say so, unstable plants are much more difficult to control than stable plants. Yet unstable plants are often viewed simply as linear plants with one or more open right-half-plane poles. (A plant with a chain of integrators or imaginary poles is also unstable, but less seriously.)

I believe the distinction between stable and unstable plants is vastly underemphasized in the research literature. An unstable plant provides almost no opportunity for on-line identification and therefore must rely heavily on analytical modeling and extrapolation from stable regimes. 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.

6. Distinguish between modeling for control architecture design and modeling for controller implementation.

Control architecture design and controller tuning are strongly interrelated, but they are effectively distinct tasks in control engineering practice. Control architecture design refers to the selection of sensors and actuators that need to be specified to achieve a control objective. The design of the control architecture and associated hardware usually depends on a solid understanding of the relevant physics along with detailed analytical modeling. In fact, analytical modeling at this stage in the control engineering process is extremely cost-effective since it reduces the need to fabricate and test multiple prototypes.

On the other hand, control architecture design is often only loosely coupled with controller implementation, that is, the choice of the control algorithm and its tuning (parameter settings). In fact, modeling for controller implementation usually requires information that is distinct from the information needed for control architecture design in both type and detail. For example, although finite-element modeling and computational fluid dynamics may provide important information for sensor and actuator design and placement, these modeling techniques usually cannot provide the type of detail needed for controller implementation, such as plant phase at crossover.

The distinction between modeling for control architecture design and modeling for controller implementation clarifies the role of distributed parameter models in control design. Such models provide a starting point for the former but have little relevance for the latter.

None of these remarks are intended to minimize the importance of either analytical or data-based modeling in control engineering. In fact, both kinds of modeling are extremely important, and they are the responsibility of the control engineer. However, it is important to recognize what modeling information is needed and knowable at each stage of the control engineering process. At early stages in control architecture design, the modeling may be largely analytical and hypothetical, whereas controller implementation must be strongly linked to the behavior of a specific hardware realization.

7. Reduce the dependence of analytical and data-based modeling for controller implementation.

As discussed above, analytical modeling is essential and valuable for control architecture design, but it has serious shortcomings for controller implementation. Although control architecture design often consumes the bulk of control engineering effort, advances in control algorithms can reduce the need for both analytical and data-based modeling for controller implementation.

The objective of robust control is to achieve performance for a given level of modeling uncertainty; however, robust control fails to reduce the dependence on either analytical or data-based modeling significantly. Although robust control methods do not require a precise model of the nominal plant dynamics, they do require that all uncertainty be quantified, and the construction and verification of such a detailed uncertainty model may require substantial analysis of prediction models and test data.

The main drawback of robust control is that it treats uncertainty as a static quantity, which forces the control engineer to sacrifice performance for robustness. Ultimately, robust control requires that the controller gains be decreased to account for uncertainty, thus reducing performance. The inability of a robust controller to learn makes this tradeoff unavoidable.

Finally, control engineering must accept the possibility that any given plant can change in an unexpected and unpredictable manner during operation. In fact, seemingly small physical changes can have a large effect on plant response. For example, as lubrication dissipates, bolts loosen, mass distribution changes, and components wear, the plant dynamics may change significantly. These unpredictable changes are the responsibility of the control engineer. Indeed, a major reason for implementing a feedback control system is to achieve performance in the presence of uncertainty, and not all uncertainty can be characterized or predicted.

8. Exploit identification for controller implementation.

No matter how well analytical modeling can be performed, some identification is always needed. Real hardware abounds with unmodelable effects and high sensitivities. In addition, modeling a system in piecemeal fashion is of limited use for controller implementation, since components can interact dynamically in complicated ways due to spurious feedback paths and unexpected interactions. The need for identification and hardware testing is crucial, and end-to-end identification is desirable whenever possible. Obviously, identification is only meaningful after the system has been constructed and data are 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 that case, a stabilizing controller is needed, which may require analytical modeling or adaptive methods. In addition, the presence of ambient disturbances can limit the ability to identify and adapt. In this case, identification accuracy may be low and the results of the identification may be nonrepeatable. If ambient disturbances can be eliminated, identification is much easier. (Engines can be turned off, whereas turbulent wind noise around a flying aircraft cannot.) Identification and adaptive stabilization in the presence of exogenous disturbances presents a severe challenge to control engineering. For this problem, the control engineer is forced to rely more heavily on analytical modeling.

Numerous issues of theoretical and practical significance are associated with identification. Since identification is difficult in the presence of fast and slow dynamics, a delta-operator identification theory would be useful [4]. Choosing good identification signals, especially in the presence of ambient disturbances, is a problem of practical interest. In some sense a good identification signal is "far" from a disturbance signal. Coding ideas may be useful in this regard.

Nonlinear identification is largely an open area of research with considerable practical importance. Since all real systems are nonlinear, it is overly simplistic to apply linear identification methods and expect that any such method will produce a meaningful linear model. I suspect that difficulties observed with linear identification methods are due to unmodeled nonlinearities as much as sensor and disturbance noise. We must also recognize and admit the possibility of systems that have nonrepeatable behavior due to sensitive dependence on initial conditions, ambient disturbances, and complex dynamics.

Finally, statistics has been underutilized by the control community as a whole for analyzing identification and performance data, although statistical analysis has been seeing increasing interest in process control. The analysis of any data without careful statistical analysis is naive at best.

Bridging the Gap: Control Issues

9. Respect the distinction between continuous time and discrete time.

A cruel fact of control engineering life is that most of the systems we need to control operate in continuous time while the controllers we implement on digital computers operate in discrete time. It can be difficult to reconcile the continuous and the discrete; sometimes they behave like oil and water. Consequently, the interface between continuous- and discrete-time systems is a tricky business, and it can have a significant effect on control system performance. The literature abounds with transformations between continuous-time and discrete-time dynamics such as Tustin's, exponential, and bilinear. However, these are merely conveniences that don't address stability and performance in a reliable manner.

The continuous/discrete gap is here to stay, since there is no revolution in analog controller technology on the horizon. Furthermore, even if we could implement continuous-time controllers, our identification methods operate on discrete-time data to produce discrete-time models. Identification in continuous time is not a serious prospect. So with analytical modeling and classical control in continuous time and with identification in discrete time, it's no surprise that the control literature often appears schizophrenic.

It is tempting to believe that for sufficiently fast computers, discrete-time systems can be treated as continuous-time systems. However, there are fundamental distinctions between discrete-time and continuous-time systems. For example, discrete-time control has an inherent bandwidth limitation imposed by the sample rate. A delay in continuous time is an irrational exponential function, whereas in discrete time it is rational (one nice benefit of discrete-time models). In addition, a nilpotent linear discrete-time system has finite settling time behavior, whereas a linear continuous-time system cannot settle in finite time. (A time-invariant continuous-time system that settles in finite time necessarily has non-Lipschitzian dynamics [5].) Finally, the behavior of the system between sample instants can affect closed-loop performance. If the sample interval is short, the intersample behavior should be benign. Whether this effect can be ignored in practice is an open question.

There are fundamental obstacles in sampled-data control that must be treated carefully. First of all, sampling and reconstruction devices, which provide the interface between the continuous-time and discrete-time worlds, have time-varying dynamics with inherent limitations. Arbitrarily fast sampling is an unreasonable expectation since faster hardware merely encourages engineers to consider ever faster plants or more computationally intensive control algorithms. Furthermore, fast sampling can cause numerical problems with poles aggregating near 1. The delta operator provides a practical solution to this problem [4]. Similarly, zero-order-hold signal reconstruction is a time-varying operation that produces spurious harmonics. Suppressing these effects is often required through additional filtering.

Aliasing is a problem that arises due to sampling. Folding of signals and noise is an unavoidable effect of aliasing, and it is rarely accounted for explicitly in control theories. Aliasing also causes phase shifts at lower frequencies that can destabilize a system. It is important in practice to determine sample rates and design anti-aliasing filters with minimal phase lag to suppress these effects.

Nonlinear systems are difficult to translate cleanly into discrete time. For example, finite escape time can occur in continuous time, but it has no direct counterpart in discrete time. Capturing nonlinear physics in discrete time is a nontrivial challenge, especially since our training and intuition are based in continuous time. Exact discretization of some continuous-time models is discussed in [6].

10. Distinguish real-time computing from off-line computing.

Many of the respondents polled in the survey [1] discussed the underutilization and potential exploitation of recent advances in computational power. One application of powerful computers is to solve very large order Lyapunov and Riccati equations. This is reminiscent of the "big drum" approach: Primitive tribes wishing to communicate with the outside world might conceive of ever larger drums. These are examples of misguided technology scaling.

The usefulness of any computation must be evaluated in light of the accuracy of the underlying data. Most measurements are good to only about 0.1%, and many parameters (not to mention physical effects) are significantly more uncertain. (A stern warning on the unreliability of data is given in Chapter 27 of [7].) Massive computation based on erroneous data or hypothetical models may have qualitative value for insight, but the actual numbers produced will likely have little connection with reality. (If insight is the goal, this is not a problem; it's just important to keep these goals distinct.) Computing with uncertain data has been largely a neglected topic in the scientific community, although the robust control community (to its credit) has given it serious attention.

In control engineering, large-scale computing is relevant for plant and control architecture design, which is largely a qualitative and hypothetical process. Such computing is performed off line and often occurs before the plant and control system are constructed. This computing is usually performed for the sake of modeling, but it is suspect as a viable approach to controller implementation. On the other hand, controller implementation can be enhanced by the capability for real-time, on-line computing. Identification and performance assessment during control system operation for adaptive control is one of the main beneficiaries of significant real-time computing capability.

There is no real tradeoff between on-line and off-line computing. They are distinct tasks that use different kinds of information for different purposes. Off-line computing is based on static and usually limited information about the system, whereas online computation has continual access to data from the true system and its infinitely rich physics as it behaves in possibly unpredictable ways.

11. Always recognize saturation.

Often the first nonlinearity encountered by the control engineer is saturation. (Here I am referring to amplitude saturation. The second nonlinearity encountered is rate saturation.) It is a universal nonlinearity that will never be circumvented by any technological development. Saturation is a linearizable nonlinearity that has a global impact on the system but has little effect on the local behavior of the system.

A control engineer who has invested in control system hardware is often interested in achieving the best possible performance from the chosen hardware. Whether or not fuel or energy constraints are present, this goal may require that the actuators operate at or near saturation levels. Hence saturation limits are not necessarily regions to be avoided, but rather may be sought so as to maximize use of the available control input.

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

Saturation may render linearization misleading for unstable plants. Specifically, linearizing a nonlinear unstable plant may obscure the actual saturation recovery limits, which are invariably smaller than those of the linearized model.

12. Recognize limitations due to sensor noise.

It is important to stress that all real signals are corrupted by noise, and this noise limits the achievable performance. Noise can arise from the sensors and all associated electronics, and its characteristics are rarely known prior to implementing the control system hardware. In particular, the noise may be due to details of grounding and shielding, whose effects are difficult to predict before the plant and controller have been constructed.

Feeding back control signals feeds back the sensor noise as well. Therefore, if the disturbance is narrow band but its spectrum is not known in advance, a control engineer might be inclined to use a controller with broadband gain. However, feedback in a frequency range where the plant disturbance is not present will amplify sensor noise. Thus, the presence of sensor noise forces the control engineer to limit control gains and bandwidth. This design issue is often ignored in control design theories since the sensor noise spectrum is rarely known from analytical modeling. Furthermore, in LQG theory, this constraint is difficult to handle because narrow-band noise gives rise to a singular estimation problem. In any event, real noise is surely more insidious than idealized noise models.

13. Emphasize the distinction between smooth and nonsmooth nonlinearities.

Although linearity over a range is an oxymoron, it is nevertheless useful. However, nonlinear effects assume greater importance as performance requirements become more stringent. Many control methods consider smooth nonlinearities, which are linearizable near equilibria and have an increasing effect over a larger range of operation. Geometric nonlinearities in robotics are the prototypical examples of such nonlinearities. Control theorists tend to think of these nonlinearities as being well known and amenable to global transformation techniques.

On the other hand, many control applications are of a precision nature where the objective is to produce highly accurate motion over small amplitudes. In this regime, the nonlinearities tend to be nonsmooth (that is, not linearizable) and possibly discontinuous. Friction is a common example of a nonsmooth nonlinearity. In addition, nonsmooth nonlinearities may possess memory as well, for example, stiction and backlash or hysteresis. Hysteretic nonlinearities are usually semistable subsystems [8] with multiple equilibria where subsystem convergence is fast relative to the remaining system dynamics. The memory characteristic is merely the trajectory-dependent set of equilibria that the subsystem converges to during quasi-static operation. The phase lag nature of such nonlinearities renders them potentially destabilizing even at low signal amplitudes.

Classical control theory discusses both smooth and nonsmooth nonlinearities with an emphasis on the former through absolute stability theory. However, nonsmooth nonlinearities seem to be more prevalent in applications. In fact, while large-amplitude motions can often be slowed down without major loss of performance (and this is often done in practice), lack of precision in small-amplitude applications can seriously degrade the value of the system. In other words, large robots can be operated more slowly if necessary (although it may not be desirable to do this), but a lack of precision in a machining task may not be tolerable at any speed of operation.

In general, nonsmooth nonlinearities are easier to identify because the amplitude range is smaller. However, these nonlinearities come in a wide variety of types, they may be hidden, and they may change drastically and unexpectedly over different operating ranges. On the other hand, smooth nonlinearities are difficult to identify because of the range of operation required to collect data. Control theorists tend to view such nonlinearities as well known because of the analytical nature of classical mechanics. In applications such as flight control over a large envelope, identification of global nonlinearities can be extremely difficult.

Final Observations

14. Remember the transients.

Control theorists have a fixation with equilibrium-related behavior. We seek the steady state because it is easy to characterize and provides a safe haven. Lyapunov stability theory, which continues to provide a rich hunting ground [8], [9], has spoiled us. In engineering practice, it is often the transients that matter. Collision avoidance is a good motivating example. But dealing with transients is not easy. As S. Ulam once said, "The infinite we can do immediately; the finite takes a little longer."

15. Feedback entails risk.

Most engineering disciplines are open loop in the sense that errors are not amplified. A 20% error in the strength of a structural member remains just that, and a 50% margin will compensate for the error quite nicely. (The Hoover Dam was designed with a 3x safety factor for those of you living downstream. As far as I know, it is not stabilized by feedback loops.) In contrast, feedback affects dynamic behavior, and "small" errors can produce arbitrarily large undesirable effects (such as instability). Attempts to guard against this sensitivity assume that modeling uncertainty is known, yet the control system must have the ability to cope with unexpected changes as well. Since control systems are often critical to operation with significant losses in the event of failure, the ability to cope with unexpected changes is the responsibility of the control engineer.

16. There are no details in control engineering.

There are no "details" in control engineering, since even the most insignificant "detail" may prove to be important. All engineering ultimately hinges on details, because real systems must be built from real, imperfect (not mathematical) components and must operate under real (nonideal) conditions. This point holds all the more for control engineering, a complex technology that depends on many interrelated aspects. The smallest "details" such as noise, quantization, drift, bias, crosstalk, roundoff, aliasing, nonlinearities (local or global, smooth or nonsmooth, memoryless or nonmemoryless), saturation (amplitude or rate), delays (known or unknown, fixed or variable), model errors, sensor/actuator dynamics, state constraints, and system changes can adversely affect control system operation. The gap between theory and practice can be narrowed by systems and control theory that recognizes the importance of these issues (and surely many others) and addresses them in a meaningful and useful way.

17. Reduce the dependence on modeling.

I believe that one of the main culprits in the theory/practice gap is the modeling dependence imposed by many control methods. In fact, the modeling requirements imposed by model-based control methods constitute a severe impediment to the applicability of modern control theory. Analytical modeling is essential for control architecture design, but it must be used with care for the purpose of controller implementation.

The success of PID tuning methods relative to modern control methods is a reminder of this dependence, while model predictive control based on identification is equally successful for the same reason. Therefore, it seems that the extent to which a control method is used in practice is proportional to its modeling requirements, making this issue a key factor in the existence of the gap. The first step toward remedying this problem is to distinguish between modeling for control architecture design before system construction and modeling for controller implementation (usually through identification) after system construction.

18. Why bridge the gap?

I have left this question for last because it is the most fundamental and most difficult. Thus far, I have suggested that there are pragmatic reasons for closing the gap between theory and practice. The transition of new ideas and techniques to applications ultimately justifies the cost of basic research. Although there certainly have been successes in the application of modern ideas to technology, the penetration of modern ideas in many applications seems to be fairly limited. Serious attention to technological needs and constraints is essential for understanding and correcting this state of affairs.

On the other hand, it is important to keep in mind that basic research is meant to be high risk in terms of payoff. Control engineers with project deadlines rarely have the luxury of pursuing unconventional ideas with uncertain return. That is the role of academia, where researchers can instantaneously shift research directions and pursue new ideas without management approval (unlike the usual case in industry), or pursue a novel idea for years until it is sufficiently developed to have technological value. What is largely lacking in the academic setting is motivation from real applications. I believe that exposure to such motivation, even to a limited extent, can have a significant impact on increasing the relevance of "academic" research to engineering practice.

References

[1] V. Blonder, M. Gevers, and A. Lindquist, "Survey on the state of systems and control," *European J. Contr.*, vol. 1, pp. 5-23, 1995.

[2] F.D. Reynolds, *Crackpot or Genius? A Complete Guide to the Uncommon Art of Inventing*, Barnes & Noble, 1999.

[3] H.C. Von Baeyer, *Maxwell's Demon: Why Warmth Disperses and Time Passes*, Random House, 1998.

[4] R.H. Middleton and G.C. Goodwin, *Digital Control and Estimation: A Unified Approach*, Prentice-Hall, 1990.

[5] S.P. Bhat and D.S. Bernstein, "Continuous finite-time stabilization of the translational and rotational double integrators," *IEEE Trans. Autom. Contr.*, vol. 43, pp. 678-682, 1998.

[6] J.E. Marsden, G.W. Patrick, and W.F. Schadwick, eds., *Integration Algorithms and Classical Mechanics*, American Mathematical Society, 1996.

[7] R.W. Hamming, *The Art of Doing Science and Engineering*, Gordon and Breach, 1997.

[8] S.P. Bhat and D.S. Bernstein, "Lyapunov analysis of semistability," *Proc. Amer. Contr. Conf.*, pp. 1608-1612, San Diego, CA, June 1999.

[9] V.I. Vorotnikov, *Partial Stability and Control*, Birkhauser, 1997.

Author's Note

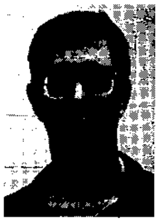

I'd like to say that my thinking on this issue was influenced by my experience in industry (Harris Corporation, Aerospace Systems Division, 1984-91) and academia (University of Michigan, Aerospace Engineering Department, 1991-present). But I think it would be more accurate to say that the greatest influences on my career were (1) the Erector Set I got when I was seven years old and (2) the time I spent fixing my car when I was in college.

Although tinkering can be valuable, what amazes me most is the ability of the human intellect to solve real-world problems through abstraction, and yet I see that much of what engineers do is rooted in empiricism. I think that a lot of engineering is actually an art, but today's art often becomes tomorrow's science. And what is surprising about this knowledge is its compressibility. An insight or breakthrough that took a lifetime to achieve may become common knowledge for the next generation. The greatest intellectual achievement of all time (the alphabet) is taught to preschoolers, while the second-greatest (the calculus) is taught to high school students.

My latest favorite quote is by Thomas Edison: "We don't know a millionth of one percent about anything." I think this is worth keeping in mind at the close of the millennium.