Hubble Trouble, Oil, and Bubble

SLOW CONTROL

We've all heard the legend that it's possible to boil a live frog without the frog jumping out of the pot. The trick is to use quasi-static control, where the temperature is increased sufficiently slowly that the frog doesn't detect the conditions leading to its demise. Since the cold-blooded frog can presumably feel temperature, what he seems to be suffering from is the inability to sense low-frequency *changes* in temperature. This is a little bit like watching the grass grow—we know it grows, but it grows so slowly that we can't detect the change; our eyes lack the resolution needed to estimate derivatives near dc. In addition to the rate at which the water temperature increases, the frog seems to lack a critical response threshold. Again, this is like trying to decide when your lawn is tall enough to warrant mowing; in the absence of a threshold, there's a tendency to let it grow taller than you might wish. As a practical application, the frog-boiling paradigm is often used as a metaphor for political entities that diminish the rights of individuals at a rate that is too small to detect and such that no threshold is apparent.

NO CONTROL

Unlike slow control, no control means not doing anything. In some cases, doing nothing is the best strategy, especially when the alternative can potentially cause harm. Medical practice strives to cure illness, but in many cases the body is self-healing and may recover without intervention. In fact, medical practice can cause harm; adverse reactions to medicine occur frequently.

ILLUSORY CONTROL

Illusory control refers to controls that are believed to be doing something but in fact have no effect. In control systems lingo, a system may be uncontrollable, although we might believe that the output is driven by the input. In medicine, this situation corresponds to placebo control, which gives the patient the expectation that the medicine prescribed by a doctor will have a positive effect. Medical doctors may prescribe placebos when "real" medicines are not available or are not warranted.

HUMAN-FREE LOOPS

Every engineered system needs maintenance—that is, maintenance by humans, who designed and built the system and now must keep it running. Components degrade with age and use, dirt accumulates, and fuel is exhausted. By modeling and predicting these trends, we perform periodic, preventive maintenance. But maintenance is costly and takes a system out of operation. In addition, there are applications—such as spacecraft—in which maintenance is impossible. In the spirit of automation, self-maintenance may allow a system to monitor its own degradation and perform its own maintenance. A machine could order its own parts and then hire someone to do the installation. Better yet, why not design a system so that it requires no human maintenance at all? The troubled Hubble notwithstanding, satellites must be designed this way. Off-shore wind turbines provide another motivation. Of course, no system can operate forever, but a zero-maintenance design philosophy could

Moriba Jah, Dennis Bernstein, and Fred Leve at the Air Force Research Laboratory in Albuquerque, New Mexico.

Digital Object Identifier 10.1109/MCS.2011.941442
Date of publication: 14 July 2011

provide the foundation for increased reliability and operational life. The potential degradation of every component becomes suspect, and each design decision is evaluated according to its maintenance needs. Of course, we don't need to interpret the "zero" in "zero maintenance" literally; we could certainly tolerate *some* maintenance, and we wouldn't refuse to repair a system that's fixable. Rather, zero-maintenance can be viewed as a design philosophy that influences the design process from the beginning by considering anything and everything that contributes to the need for maintenance. In the end, we can hope for lower cost and higher reliability.

Warren Dixon and Dennis Bernstein.

In some cases, doing nothing is the best strategy, especially when the alternative can potentially cause harm.

TRADITION

In the early days of the Conference on Decision and Control (CDC), proceedings were bound in oversized paperback books. When the last CDC proceedings was printed, conference attendees were given the option of five massive volumes or one slim CD-ROM. Over the years, the sheer size of these volumes necessitated a limit to the length of papers, which sort of converged to eight pages. In fact, papers submitted to CDC can be up to eight pages long, but here's thing: Publishing the full eight pages costs an extra US$300. If you don't want to pay this amount, then you can trim your submitted paper to fit the free-page limit of six pages. Why the discrepancy? Historically, the page charges motivated authors to shorten their papers and thus limit the massiveness of the printed proceedings. Truncation was rationalized on the grounds that the extra submitted pages allowed the reviewers to check proofs that would later be excluded from the published version. At the same time, the income from the overlength charges helped (and still helps) to augment the proceeds of the conference—which are helpful but not essential to IEEE Control Systems Society (CSS) operations. In my view, none of these explanations justifies the persisting gap between the submission constraint and the extra-cost threshold. The Society does well financially without the extra charges, and attendees ought to be reading what the reviewers reviewed. If space is an issue, we can always move up to DVDs or jump drives. The main thing is that the proceedings should include the material that's reviewed, without added costs.

TALL TAILS

Examples of 80-20 distributions abound. In one city in New Jersey, 1% of the patients account for one-third of all the medical expenses. I probably consult 2% of the books in my office 90% of the time. What accompanies distributions like these is the fact that the tails have more area and thus represent a larger number of extreme and rare outcomes than we might expect to see. Probability textbooks condition us to believe in the ubiquity of the Gaussian distribution due to the central limit theorem, where sums of independent increments—like the problems on a quiz—add up to a normal distribution, whose tails fall off exponentially. But more and more applications reveal that a surprising amount of content resides in the tails of the distribution, which typically has an infinite second moment. In fact, the central limit theorem requires that the independent random variables have finite variance, but this is often not the case in many applications. For example, stock market fluctuations and rainfall are better modeled by random variables with a power-law distribution, and therefore have infinite variance; in these cases the central limit theorem does not apply. Heavy tails are the bane of perfectionists, who are obsessed with fixing every last problem, despite the fact that the required effort grows exponentially with decreasing benefit. So why don't we shed the 20% of our tasks that consume 80% of our effort?

You never know when a dusty book may have the answer you're looking for.

We'd be only 20% less effective with a whopping one fifth of the effort—a huge gain in efficiency. The problem, however, is that it might not be so easy to determine which cases reside in the tail and what benefit might be reaped from the extra effort. And you never know when a dusty book may have the answer you're looking for—or the patient at the margin might be you.

CONFERENCE2JOURNAL

So, you've written and presented a nice paper at a CSS conference, and now it's time to submit it to *IEEE Transactions on Automatic Control* (*TAC*) or *IEEE Transactions on Control Systems Technology* (*TCST*). Is that OK? The good news is that there is no copyright impediment, since *TAC* and *TCST* are CSS journals that reside within IEEE. But the bad news is that IEEE frowns on the use of "previously published material." What's flawed about this policy is that the word "published" is used without distinction. In fact, publishing in the CDC proceedings is vastly different from publishing in *TAC* or *TCST*. The former involves one rather quick round of reviews, while the latter usually involves multiple rounds of intense scrutiny and verification. By not making this distinction, IEEE shows a lack of appreciation for the value that its transactions bring to the quality of scientific discourse. The fine print on the inside back cover of *TAC* advises authors to be sure that their submission "adds value relative to its conference version," where this "value" can consist of proofs or numerical results "omitted" from the conference version. What this guidance misses is the fact that the "value" added to conference publications is not due to additional numerics or extended proofs. Rather, it is the value that the scrutiny per se brings. Otherwise, why do we bother at all with transactions?

COSTLY PURSUITS

What happens as resources become scarcer? First, price goes up. Second, substitution occurs, where consumers choose less-expensive alternatives. These are immutable laws of economics. The problem, however, is that, whereas the first phase happens relatively quickly, the second phase can be long in coming. In the interim, prices go up, which drives a more exhaustive quest for the remaining resources, which in turn become increasingly more expensive and environmentally destructive. A depressing example is the tar sands in Alberta, Canada, where the separation of bitumen from dirt consumes a fourth as much energy as the oil delivers, destroys forests and water supplies, and leaves vast areas of toxic sludge that will persist for centuries. Political opposition to the lack of oversight of the environment and public health are all but impotent in the face of the growing power of the advocates for profit. Apparently the frog is boiling too slowly.

Dennis S. Bernstein

Natural Challenges

The widespread availability of computers, the requirement for environmental impact statements and cost-benefit ratios, and the dawn of mathematical models all arrived on the scene simultaneously in the final quarter of the twentieth century. Scientists in the 1960s and 1970s assured bureaucrats that the computer would make it possible to predict the outcomes of natural processes accurately. We don't know how to do it right now, they said, but fund us and we'll figure it out. There are still some scientists who claim successes—undaunted by several decades of the failure of certain mathematical models to provide the accurate answers that society needs.

At the beginning of the twenty-first century, predictive models of processes on the surface of the earth have come into widespread use. The recognition of complexity and chaos seems not to have diminished the still-rising star of modeling. Every year hundreds of cost-benefit ratios roll off the presses for federal engineering projects involving beaches, rivers, lakes, and groundwater flow. Engineers who have found great success in the use of models to predict the behavior of steel and concrete have applied modeling to the natural environment just as if nature were made up of construction materials with well-defined properties.

— Orrin H. Pilkey and Linda Pilkey-Jarvis,
Useless Arithmetic, Why Environmental Scientists Can't Predict the Future,
Columbia University Press, 2007.