Process Control Case Histories

An Insightful and Humorous Perspective from the Control Room

GREGORY K. MCMILLAN

MOMENTUM PRESS

Momentum Press, LLC., New York

Process Control Case Histories: An Insightful and Humorous Perspective from the Control Room

First published in 2010 by
Momentum Press®, LLC
222 East 46th Street, New York, N.Y. 10017
www.momentumpress.net

ISBN-13: 978-1-60650-176-4 (softcover)
ISBN-10: 1-60650-176-3 (softcover)

ISBN-13: 978-1-60650-177-1 (e-book)
ISBN-10: 1-60650-177-1 (e-book)

DOI forthcoming

Cover Design by Jonathan Pennell
Interior Design by Scribe, Inc. (www.scribenet.com)

First Edition March 2010

10 9 8 7 6 5 4 3 2 1

Printed in Taiwan ROC.

Acknowledgment

These articles are the sum result of the influence of many people. I have been fortunate in having interesting and stimulating friends with exemplary attributes. Special thanks go to Steve Sanders and Stan Weiner for humor and individuality, Terry Tolliver for logic and detail, the late Vernon Trevathan for insight and perspective, and my daughters Christine and Laura for creativity and artistry.

Preface

I first recognized technical writing can be boring when I fell asleep proofreading one of my works, and I was only on the second page. I considered advertising my books as a sleep aid but was afraid of lawsuits from misuse. I envisioned some person reading my book while operating heavy equipment and suing me for an improper warning label. I then looked around my book shelf and realized I had bought a lot of technical books and subscribed to a lot of technical journals that I never read despite good intentions. I thought—what good was writing a technical piece no matter how note worthy, if it was rarely read? I had to change for my own sanity and the sanity of the readers.

The result was "pH Control: A Magical Mystery Tour" which was published after much soul searching (and perhaps some doubts) by InTech in September of 1984. It broke new ground in that it used some humor to create enough interest for people to get through the whole article and absorb what I though were some important points based on actual startups. Since then, InTech has published eleven more such articles all based on personal field experience. This approach combining humor and actual experience was subsequently used in a series of ISA books coauthored by me and Stan Weiner and a "Control Talk" column that has been running monthly in Control magazine since 2002.

Process Control Case Histories—An Insightful and Humorous Perspective from the Control Room is a collection of my InTech articles that started it all along with two new pieces written especially for this book. This collection is my favorite book because it offers an insightful description by using humor of important experiences gained in plants. I did slip in one article that is not funny. Some people may take issue that there is more than one because as Mr. Rogers says "Everyone in the world is different from everyone else. We look different, we smell different, we sound different, and we have different human thoughts inside us." Thank goodness, think how boring it would be otherwise.

Contents

pH Control: A Magical Mystery Tour

I do not pretend to be the world's champion pH problem solver. But, I have helped rescue over 50 foundering pH control systems in the past 5 years and am still rational enough to tell you about it. Most people hardly notice that I twitch at the mention of hydrogen ion concentration.

WHY IT'S A PROBLEM

Why is pH control a problem? After all, you've got an odd but simple scale of measurement from 0 to 14 dimensionless units, measuring electrodes that have been around long enough to be well understood and readily applied, and instrument vendors who must have seen every possible application by now.

Rangeability and sensitivity

One basic source of difficulty is that—as countless articles, technical papers, and textbooks point out—the pH scale corresponds to hydrogen ion concentrations from 10^0 to 10^{-14} moles per liter. No other common measurement covers such a tremendous range.

Another intrinsic constraint is that measuring electrodes can respond to changes as small as 0.001 pH, so instruments can track hydrogen ion concentration changes as small as 5×10^{-10} moles per liter at 7 pH. No other common measurement has such tremendous sensitivity.

The implications of such great rangeability and sensitivity can be illustrated by considering a continuous feedback neutralization system for a strong acid and a strong base. The reagent flow should essentially be proportional to the difference between the hydrogen ion concentration of the process fluid and the set point. A reagent control valve must therefore have a rangeability greater than 10,000,000:1 for a set point of 7 pH when the incoming stream fluctuates between 0 and 7 pH. Moreover, uncertainties in the control valve stroke translate directly into pH errors, such that stick-slip of only 0.00005% can cause an offset of 1 pH for a 7 pH set point.

The situation is like playing golf. The distance from the tee to the green represents rangeability and the ratio of the hole diameter to this distance is analogous to sensitivity. For an application requiring a strong base to neutralize a strong acid or vice versa, the tee would be about 1,000,000 yards from the green and the hole would be about 3-½ inches in diameter. A hole-in-one is impossible. And using the same size control valve at each stage would be like hiring a gorilla to drive the ball to the green in one stroke, then finding that it tends to overshoot the hole on the putt.

How is it even possible to control a process under these conditions? Rangeability and sensitivity limitations can be overcome by approaching the set point in stages, using successively smaller control valves with high performance positioners.

The real world

A host of other constraints adds to the difficulties of pH control. These range from the necessity of wetting the electrodes—with consequent susceptibility to leakage and attack by the fluid—to long delays introduced by the need to mix large volumes of process material with small amounts of reagent. Even with a good understanding of measurement and control concepts, these real-world effects introduce an element of magical mystery into pH.

SOME TYPICAL PROBLEMS

No pH complications are really typical. And the systems that are easy to implement don't get referred back to those of us whom *InTech* refers to as the noodnicks from Central Engineering. But the installations I will describe are typical of those I have encountered recently and illustrate the types of problems that you can expect.

To avoid arguments with our Legal Department about proprietary information, I will not mention any places or names; I'd even prefer that you forget my name when you finish reading. Also, to help guide you through my tribulations, and those you will encounter yourselves, I've prepared Table 1, enumerating what you can think of as the Facts of Life. Commit this Table to memory; you'll be quizzed on it in the morning.

Where's the tank?

An application involved a strong acid waste flow, to be neutralized by a strong basic reagent. I was called in because the pH was swinging from 0 to 14 despite efforts to tune the controllers, manually manipulate the reagent, and regulate the influent flow. When I got to the plant, I gazed over the horizon and didn't see any tanks. I suddenly realized that I had a **major** problem.

Figure 1-1a shows the original control system. This used a ratio controller to proportion reagent to acid waste flow upstream of an in-line mixer. A separate pH controller was used in a loop on a sump. The system designers did not realize that the flow measurement error and the flow control valve stick-slip must both be less than 0.00005% to stay within 1 pH of the 7 pH set point. They assumed that disturbances would be small since the change in waste composition was slow and its flow was fixed by a controller. Their design team did not know *Fact of Life #1*.

A system involving a strong acid and a strong base normally requires three stages of control to hold a solution to within 1 pH of 7 pH (Ref 1). Since cost was stressed as a factor, I kept the existing mixer and sump as one stage and added two vertical well-mixed tanks downstream for the second and third stages. Further, I agreed to not install controls on the third stage until the need was

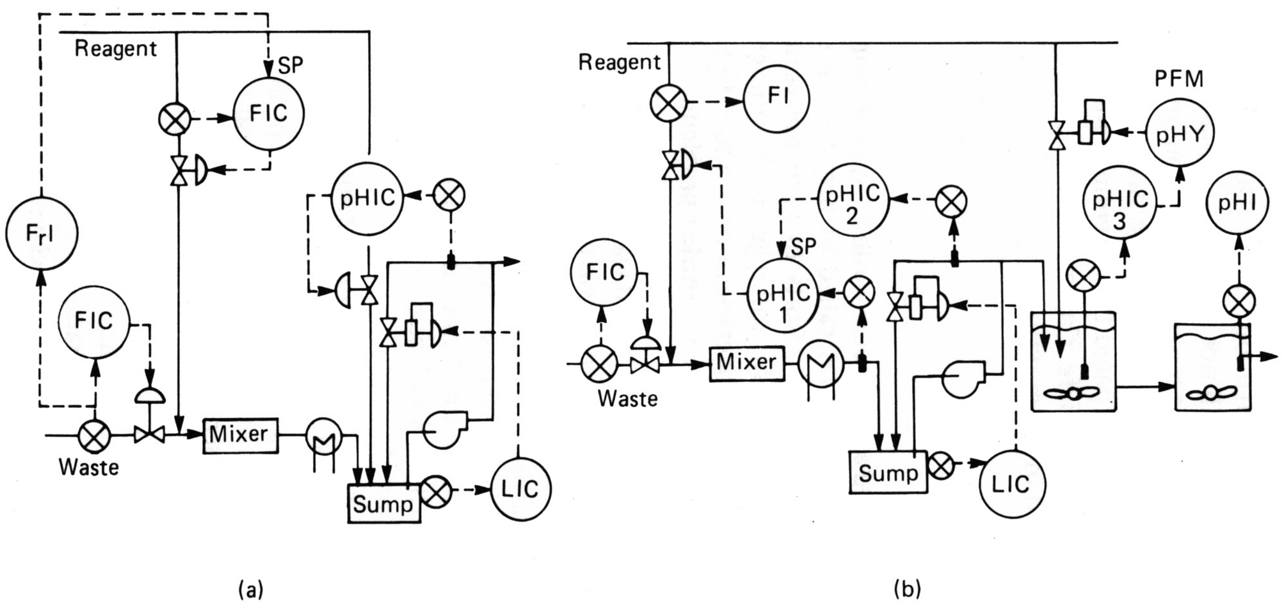

Figure 1-1. Where's the tank? (a)-unsuccessful and (b)-successful pH control systems for a continuous neutralization process initially having no mixing tank.

demonstrated. The third stage volume therefore served as a filter for the oscillation from the second stage.

For the first stage of control, we started by replacing the ratio flow system with a fast in-line pH loop. This received a remote set point from a second pH controller on the sump. The fast in-line loop would initiate the correction and depend on the sump volume to average out hydrogen ion concentration deviations. Linear control system analysis predicted that this combination would be as effective as a single well-mixed vertical tank.

It didn't work. Dynamic simulation showed that the in-line loop would oscillate between 0 and 14 pH for all controller settings. A plant test confirmed the result.

At first, I thought the sump was somehow not providing the anticipated filtering. Then I remembered *Fact of Life #2*. The filter was acting on hydrogen ion concentration, not pH. The sump was attenuating concentration oscillations by a factor of 100, but this corresponded to a decrease of only 2 pH. Attenuation was improved by reducing the distance from the mixer to the control valve and electrodes so the oscillation was faster.

The second stage had a notch-gain pH controller with an output that provided a pulse frequency proportional to an analog signal. Above 25% controller output, the valve was throttled normally; below 25%, valve rangeability was extended using pulse frequency or interval control.

Figure 1-1b shows the upgraded installation. This system could keep pH within the desired offset band at the outlet of the third stage. However, the sump controller was difficult to tune and recovery from start-up or waste flow controller set point change was slow.

If I were designing this system today, I would place a feedforward loop on the sump and would install controls on the third stage. I would also characterize the feedforward and feedback signals. The characterization would involve calculating reagent demand from the pH measurement using the titration curve, and using the result as the control command. This would reduce nonlinearity, recovery time, sensitivity, and tuning difficulty. Microprocessor-based controllers can provide the necessary calculation accuracy and ease of implementation.

As with any new system, start-up was not without bugs. Some were of the common garden variety—like transposed wires and incorrectly calibrated positioners. Others were of the magical mystery type peculiar to pH systems.

For example, at high pH levels, the measurement went downscale as the strong base reagent flow increased. As you can imagine, this drove the control system—and us—kind of crazy. The difficulty turned out to be that the measuring electrodes in the in-line loop were not specified with high-pH glass. Normally this would cause the measurement to read low by about 1 pH at the upper end of the scale. In our case, it caused a reversed response. This performance was confirmed by the vendor and was corrected by replacing the electrodes with low sodium ion error devices.

Table 1-1 The Facts of Life

1. Instrumentation is frequently the source of disturbance for pH systems, through repeatability error, measurement noise, or valve stick-slip.
2. In-line pH loops will oscillate, regardless of controller modes and tuning, if set points are on the steep parts of the titration curves.
3. pH electrode submersion assemblies with unencapsulated terminations below the liquid surface will eventually have wet terminations.
4. Reagent control valves that are not close-coupled to the injection point on in-line systems will cause reagent delivery delays large enough to describe the tools of your trade in words your sister may not even know.
5. You need either a flowmeter or a seer to diagnose reagent delivery problems.
6. Flow feedforward signals should be corrected by pH controller outputs and employed to operate linear reagent valves directly or to establish reagent flow control set points.
7. Transportation delays to pH electrodes in analyzer houses will exceed mixing deadtimes—such that increasing comfort in checking the electrodes is offset by decreasing comfort in checking trend recordings.
8. Injection electrodes should be preferred to sample holder assemblies whenever possible to reduce maintenance problems and improve response times—but not all injection electrodes are created equal.

9. Large tanks are fine if you don't have to control them; use the volume upstream to reduce reagent consumption or the volume downstream to reduce control error.
10. Install one or three but never two electrodes for a pH measurement. Use middle signal selection for three measurements for inherent protection for all types of an electrode failure and minimization of error and noise.

Another magical mystery effect was that the electrode response for the well mixed tank became erratic. We found water on the terminals inside the submersion assembly. The vendor told us if we bought an assembly that cost twice as much, the leakage would stop. We did; it didn't. The vendor then told us to buy a newly developed assembly for four times the price of the original and the leakage would surely stop. Rather than make the same mistake three times, I shopped around and found a throwaway electrode assembly completely encapsulated in plastic—at half the price of the original. It worked like a charm. A similar experience with a submersion assembly from another vendor led me to *Fact of Life #3*.

Where's the valve?

Another application required small quantities of a highly concentrated viscous reagent for continuous neutralization of a waste stream. The control system was so slow that disturbances passed through the plant long before any corrective action took effect; further, the pH trend recording had a noise band that far exceeded the allowable set point offset. When I inspected the system, I stood near the injection point at the inlet to the pipeline mixer, scanned the horizon and didn't see any reagent control valve. I quickly deduced I had a **major** problem. Figure 1-2a shows what I found.

Can you spot a control problem exclusive of the pH loop in this figure? The sump level controller sets the flow in the upper outlet branch. The mixer flow controller simultaneously manipulates the valve in the lower branch to keep a constant flow out of the sump. The system is obviously overcontrolled. We got out of this mess by cascading the level controller output to the flow controller set point.

Now for the pH loop. The reagent was being injected into the pipeline under the control of a positive displacement metering

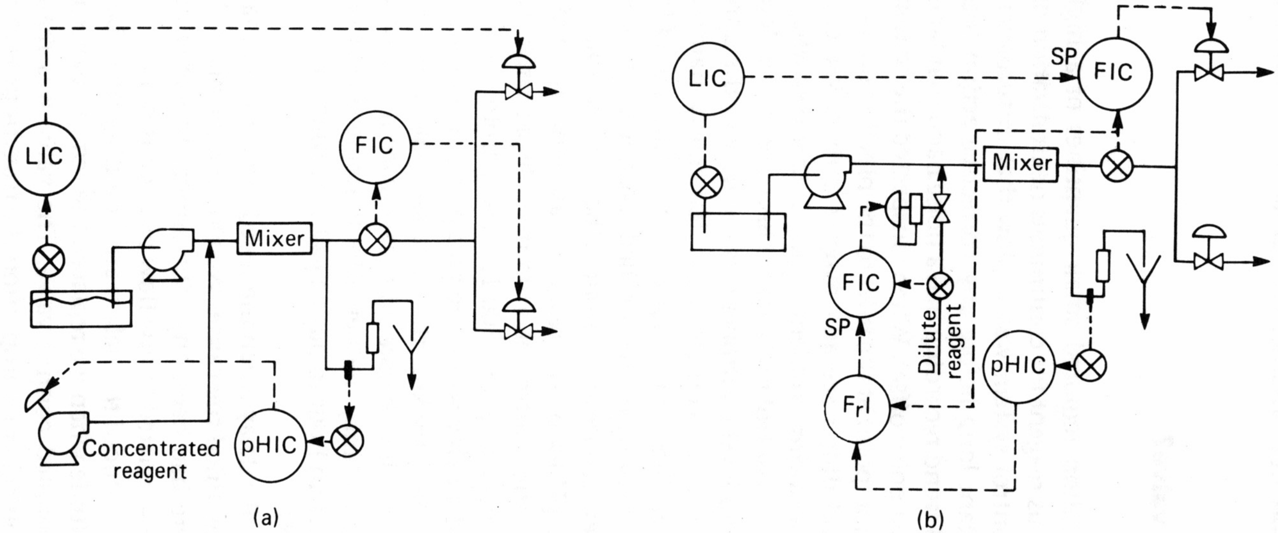

Figure 1-2. Where's the valve? (a)-unsuccessful and (b)-successful pH control systems for a process involving a highly viscous concentrated reagent.

pump. The pump was about 300 feet away from the mixer. This distance caused a delay when the pump was activated—because process fluid would backfill the injection piping and had to be pushed out of the line before any reagent could be delivered. It doesn't take much fancy mathematics to figure that at one gallon per hour, it takes an hour to push a gallon through a pipe. This led to *Fact of Life #4.*

We also found a delay when the speed of the pump changed but never really identified the cause. We would have blamed it on air pockets, if there had been any. The answer probably lies in the ketchup bottle—related to low flow of viscous fluids.

Anyway, we reduced the delays and resulting noise band by an order of magnitude when we replaced the remote metering pump with a close-coupled control valve. The valve was manipulated using a ratio controller to proportion the reagent flow to the sump discharge flow, correcting the ratio with the in-line pH loop.

Some noise still remained, due to poor distribution of the injected reagent into the pipeline. This couldn't be eliminated, because it required making the injection port smaller so the reagent velocity would be larger. Unfortunately, a hole small enough to do the job was too small to keep from plugging. The noise was more of a nuisance on the trend chart than in the system, so the record was cleaned up by passing the measurement signal through an electronic filter.

We thought our problems were over, when magical mystery reared its ugly head. As the miniature reagent valve was stroked from closed to open, the reagent flow measurement momentarily increased and then went to zero. The magnetic flowmeter was immediately suspect—but came through with a clean bill of health; we checked the wiring and found it to be correct; the vendor examined and verified the integrity of the electronics; we tested the meter on water and observed that it responded correctly. We than tried changing valve trim, but several tests yielded the same results.

I was about to throw the tiny but costly trims away, leave the engineering profession, and enter a seminary. During this period of contemplation, I suddenly noticed what looked to be a reverse taper on the trims. It was hard to tell for sure, because the parts were small, but I confirmed the observation with a micrometer. In desperation to get home from this start-up, I calculated the contour of the plug for a linear characteristic, made a sketch, and had the parts machined.

The valve worked fine with the homemade trim. The reverse taper had caused the flow to decrease as the stroke increased. The momentary surge inflow at the start of the stroke was caused by the plug lifting off the seat just enough to provide a small annular clearance. How did the reverse taper get there in the first place? I never found out for sure but did learn that the trims were too small to be standard and were specially machined by the vendor for the order. As far as I was concerned, they were too special. You can imagine how difficult it would have been to diagnose this valve problem if there was no reagent flow meter. This leads to *Fact of Life #5.*

Another instrumentation problem occurred later, when one of the design engineers decided to modify the system and recover some panel space. He installed a feedforward controller in place of the ratio station and pH-based flow controller. The device added the flow feedforward signal to the flow command from the pH controller. The vendor, anxious to sell a feedforward element, thought it was a great idea. In operation, as you should have guessed, the flow controller readjusted its output to cancel the effect of the feedforward signal and maintain flow at its set point. To work as expected, the feedforward action would have to be on the flow controller set point to provide a ratio of reagent to process flow-corrected by the pH controller output. The ratio should force the reagent flow to zero if the process fluid flow is zero Also for you control jocks, the change in controller gain with flow is negligible since the inline system's deadtime is larger than its time constant. This leads to *Fact of Life #6.*

All of these corrections are reflected in Figure 1-2b. The system, as shown, has been controlling well since start-up.

Where's the agitator?

A process used a vertical tank for neutralization. Performance was poor because response was slow and the effluent was not uniformly mixed. I looked at the drawings and noted that the vertical unit seemed a bit tall for its diameter. I asked how high it was, and the designer said, "50 feet." I gasped, "It's not nice to kid an old engineer." He responded, "Who's kidding?" I then asked, "Where's the agitator?" He replied, "You're the only agitator on this project." I instantly knew I had a **major** problem.

Figure 1-3a shows how the pH was originally being controlled. Axial agitation probably would have corrected the difficulties but could not be provided economically because the tank was too tall. A shorter tank would also have worked—again at a higher price than the plant wanted to pay.

I decided that the best way to cope with the tank would be to use its volume as a filter, estimating that it would attenuate the hydrogen ion concentration oscillations of an in-line loop by a factor of 10,000—4 pH units. A circulation pump was installed as a low-deadtime in-line mixer. Influent and reagent were added to the new suction; an injector probe was installed on the pump discharge. The new system is shown in Figure 1-3b.

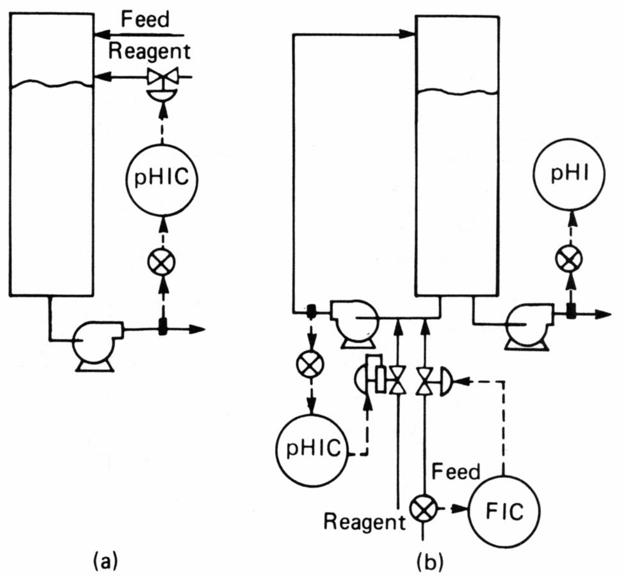

Figure 1-3. Where's the agitator? (a)-unsuccessful and (b)-successful pH control systems for a process involving an extremely tall mixing tank without an agitator.

Upsets still occurred, due mainly to the quick opening characteristic, the poor sensitivity of the spool positioner, and the stick-slip of the rotary cylindrical plug valve on the influent. However, the in-line pH loop returned rapidly to set point after a disturbance. Further, after passing through the tank volume, the pH drew the straightest line I have ever seen; for a moment, we thought someone had tied down the pointer.

Performance was so good that the plant suggested we standardize on this type of system for pH control. I warned them that the set point of this system was several pH units below the neutral zone, on a relatively flat position of the titration curve. On a steep part of the curve, *Fact of Life #2* would prevail and there would be lots of oscillations.

Where's the electrode?

I was called in to troubleshoot the pH system shown in Figure 1-4a. This simple configuration should have worked flawlessly, but was plagued by an unacceptably wide control band about the set point.

I went down to look at the exit nozzle of the vessel and couldn't find the electrodes. I rapidly surmised that I had a **major** problem.

In this case, the source of the difficulty was political. The instrument maintenance department had specified that the electrodes be located in the analyzer house to avoid the discomfort of servicing them outside during the winter. Unfortunately, this location introduced excessive deadtime in the loop. To help avoid this problem in other situations, I feel compelled to state *Fact of Life #7*.

I succeeded in getting the electrodes moved by arguing about the extreme safety hazards and product quality problems that accompanied large pH excursions. The change, indicated in Figure 1-4b, narrowed the control band to about 0.1 pH.

We used injector electrodes for this application. Experience shows that these provide better performance and require less maintenance than sample chamber electrode holders. These benefits are especially evident when the electrodes are mounted in the discharge nozzle piping where fluid velocity is high—because the flow ensures rapid response by minimizing boundary layer thickness and prevents electrode coating by impurities in the stream.

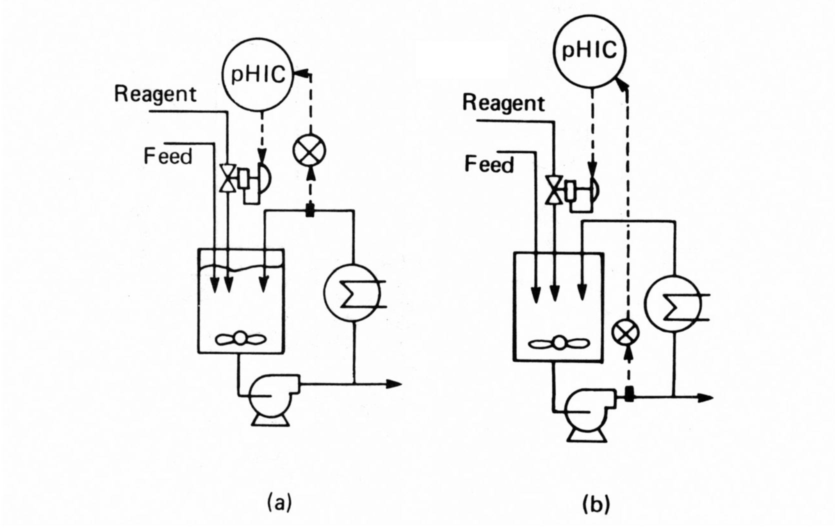

Figure 1-4. Where's the electrode? (a)-unsuccessful and (b)-successful pH control systems for a process in which electrodes have to be installed in inconvenient locations.

Injection electrodes also appear to be less prone than sample chamber elements to leakage. In checking 30 installations of injection devices from one manufacturer, I found no instances of leakage; in fairness, when we obtained products from a different source, some leakage did occur. However, every sample chamber electrode holder I have ever encountered has eventually leaked. Moreover, leakage is visible with injector assemblies but not with sample chambers. For hazardous fluids, you don't want any surprises when you open the top cover of the electrode holder. This leads me to *Fact of Life #8.*

Is bigger better?

A plant used the system of Figure 1-5a for waste neutralization. The eductor shown in the figure had been added because mixing deadtime was too long. But even with this device, the deadtime appeared to be over 40 minutes. The consequent natural period of the pH loop was 160 minutes, so the maximum reset should have been less than 0.01 repeats per minute. Since this was below the

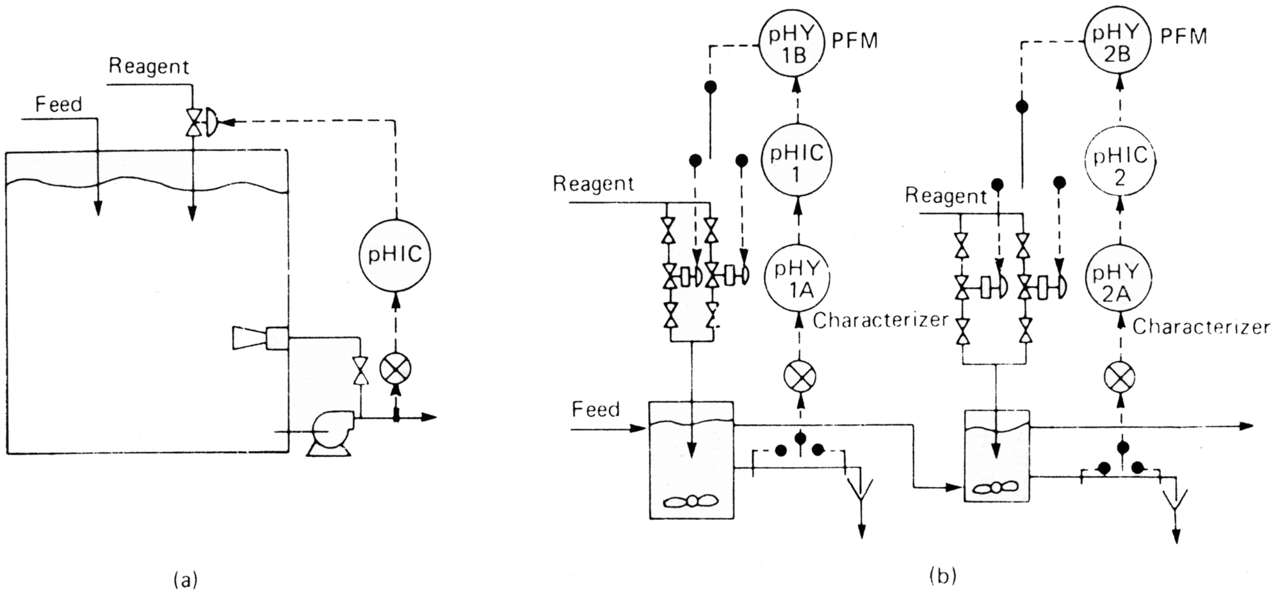

Figure 1-5. Is bigger better? (a)-unsuccessful and (b)-successful pH control systems for a process in which an extremely large tank was initially employed for mixing.

minimum setting on the controller, the loop was in a continuous reset cycle; further, the integrated error—which is proportional to the deadtime squared—was out of this world. I looked at the engineering flow diagram and spotted the largest storage tank I had ever seen. I asked the process engineer where the neutralization tank was, and he pointed to the elephant I had just thought was for storage. I immediately understood that I had a **major** problem.

The intent of the large tank was plausible. It would serve to blend acidic and basic waste streams from different sources and minimize the reagent demand. Now, as long as you don't have to put control loops on them, large tanks are useful. Upstream of a control loop, a large tank can filter out disturbances and reduce reagent requirements; downstream, it can filter out loop oscillations—which is particularly advantageous because these fluctuations are usually faster than variations in influent concentration and are therefore more effectively attenuated. This reminds me of *Fact of Life #9*.

The new control system is shown in Figure 1-5b. The large tank was replaced with two small vessels in series. A pulse frequency controller was installed to avoid valve pluggage at low reagent flows and to meet the extreme rangeability requirements imposed by the wide variations in influent flow and pH. Signal characterization was used to counteract the steep slope of the titration curve at the set point.

Start-ups are no fun without magical mystery. In this instance, we noticed that the pH measurement on the first tank was erratic. The problem could not be duplicated when we removed the electrodes and inserted them directly into the buffer solution or connected them to the measurement system of the second tank. We replaced the pH transmitter, preamplifier, cable, and electrodes individually, but the erratic measurements continued. Eventually, someone remembered that the fiberglass preamplifier enclosure supplied by the manufacturer was replaced by the field maintenance department with a metal housing—to provide more room for access. The enclosure mounting plate was grounded. This created a second ground point in the circuit, and caused a significant current flow through the circuit. The problem did not occur on the second tank because the preamplifier housing was not mounted on a conductive structure. Likewise, the erratic behavior was not observed during buffering because the bottle was plastic. The problem was solved by isolating the preamplifier enclosure from ground with a plastic mounting plate.

The control system has performed well from start-up except for periodic pluggage of the electrodes in an overflow sample line. Liquid head is too low to achieve a sample velocity sufficient to sweep the electrodes clean. A new electrode holder that provides a large flat electrode surface will be tried. If that doesn't work, we may have to shake loose enough money to install a sample pump and an injector electrode assembly.

Where's the reagent piping?

The pH in a neutralization tank was fluctuating in what appeared to be a square wave. The system was also subject to periodic glass electrode failures caused by etching and severe upsets due to a high-temperature interlock that shut off the reagent flow. Plant people were especially anxious to improve this system because reliability was critical to plant productivity. I stood at the top of the vessel wondering what to do and noticed that the reagent was being transported by a conveyor rather than a pipe. I soon perceived that I had a **major** problem.

Figure 1-6a shows the original installation. The reagent, pulverized lime, was controlled by a rotary feeder at the discharge of the hopper. Feeder speed was set by the pH controller output. Reagent delivery was subject to several minute's lag due to transportation delay on the conveyor and solids dissolution time. We made precise measurements of the pH in the tank and found that the square waves were worse than the plant thought—the process instruments recorded only the high end of the pH scale, but the fluctuations actually covered almost the whole range from 0 to 14.

Luckily, a huge tank upstream of the waste flow provided enough inventory so the pH controller could be used to throttle the waste stream. The lime feeder speed was determined by selecting the lower of a manually entered throughput set point and a command from the temperature override controller. The low signal selector therefore provided smooth transition between normal and override control. The feeder speed signal is also multiplied by the pH controller command, passed through a lag unit whose delay is set equal to the

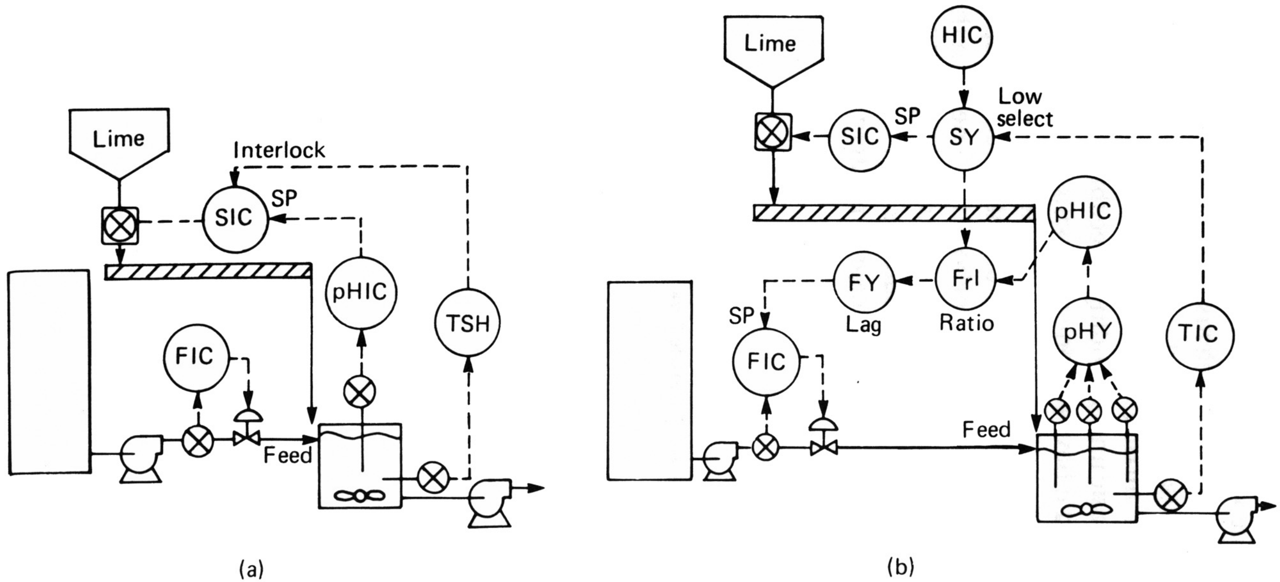

Figure 1-6. Where's the reagent piping? (a)-unsuccessful and (b)-successful pH control systems for a process in which a powdered line reagent is delivered by a conveyor.

reagent delivery time and fed forward to establish the waste flow set point.

To eliminate downtime due to electrode failures, a system was installed using three measuring elements and voting logic to establish the output signal. Use of three rather than two electrode assemblies makes it possible to determine which signal to use, if the electrode outputs disagree. This leads me to *Fact of Life #10.*

Control improved dramatically. Electrode failure due to etching that had occurred when the solution was acidic—at the unrecorded lower portion of the square wave—also stopped. And the use of voting logic to control using three electrode assemblies has virtually eliminated downtime, even when an element becomes nonfunctional.

USING YOUR SKILLS

One of the prices you pay for being an instrumentation expert in the processing industries is that occasionally someone will ask you to control pH. The job rarely proves to be easy, for instance, because you are on a flat portion of the titration curve or have wide tolerance on response and accuracy, because chances are then high that someone has done it satisfactorily without you. So the problems you get are usually **major** problems. You'll have to call on all you know about the installation and operation of electrodes, control valves, piping, and mixing equipment. You'll have to brush the cobwebs off your basic understanding of feedback and feedforward loop strategies. You'll have to hone your skills as a diplomat to get the plant to install, replace, or eliminate vessels or instruments that make life convenient for the operators or maintenance people—or represent investments for which somebody has a neck on the line—but are preventing satisfactory pH control. And you'll have to resign yourself to living out of your suitcase for a while, while the plant starts up and experiences the magical mystery of pH.

REFERENCE

1. McMillan, G K. and Cameron, R.A.; Advanced pH Measurement *and Control.* 3rd edition. ISA (Research Triangle Park NC), 2005.

Compressor Surge Control: Traveling in the Fast Lane

Let's say that instrumentation starts to become boring and you want to add a little excitement to your life. Volunteer to work on the surge controls for the next compressor project. There is nothing quite like the experience of standing next to a compressor going into surge. Noise from the collapsing voids within the unit and flexing of the intake duct during flow reversals have caused even the oldest and bravest control system designers to set records for the 60-yard dash. And operators always appreciate the exercise they get acknowledging radial vibration and axial thrust alarms; of course, if a unit stays in surge for about 20 seconds, the alarms will stop—at which time, iron will have been relocated on the plot plan.

Part of the thrill of the surge experience is that you often cannot watch it coming. Forward flow reverses direction in about 30 milliseconds. A second or so later, the reverse flow changes to forward again and the cycle repeats itself. Moving up the characteristic curve to the point of slope reversal and surge is like flooring the gas pedal of a car on the way to a precipice at the top of a mountain.

Surge dynamics carries good news and bad news. The good news is that the sudden flow transient can be used as a surge indicator. The bad news is that once the compressor reaches the point of zero slope on the characteristic curve, nothing can be done to prevent the surge. If you try to use feedback to get out of surge, you will find that the flow oscillation period is much shorter than the controller settling time. The oscillations therefore seem like uncontrollable noise to the feedback loop. Some engineers add enough damping in the transmitters so that neither the controllers nor the operators see the oscillations. This head-in-the-sand approach relies on divine intervention to prevent disaster, and suggests *Surge Control Fact of Life #1* in Table 2-1.

SOME REAL-LIFE EXPERIENCES

I have worked on many compressor control systems. They were all exciting—but three stand out in my mind no matter how hard I try to forget them. They all demonstrate the importance of good engineering practice—with the ability to keep your wits about you as events appear to roll by with the videotape on fast forward.

The phantom strikes

One of our company plants once bought an innovative compressor, designed by some aerospace wizards to run efficiently at high speeds. It was destined to be the only one of its kind in the world, for good reason.

The project team's initial delight turned to dismay when the unit proved particularly efficient at running into surge and breaking speed records that were not meant to be broken. A few people expressed doubt that a compressor could turn as fast as the gauges indicated, but excessive wear and tear on moving parts proved otherwise. More critically, the compressor would shut down through a high speed interlock so rapidly that no one could figure out what had initiated it.

The trips were soon blamed on a phantom, with irregular sleeping hours and a maudlin sense of humor, who seemed to shut the compressor down every time things were looking up. Experts from Central Engineering were dispatched to the plant to find the phantom, but their quest went unrewarded. The situation persisted for several years, at which time I was asked to join in on the adventure—with the warning that the only way to get off the project was to die, quit, or solve the problem. I was motivated.

My first goal was to get a snapshot of the operating conditions just before shutdown, to investigate possible clues that an overspeed trip was developing. I installed a fast analog derivative module to compute the acceleration and a high-speed oscillograph that could reach steady state within a few milliseconds of receiving a remote trigger signal. A voltage output from the derivative module started the recorder at the beginning of compressor acceleration. All controlled and manipulated variables, including actual valve positions on the compressor, were recorded.

The oscillograph provided tracks that led to the phantom. The little bugger was not being especially sophisticated, just operating in a shorter time frame than anyone had suspected. The records showed a classical non-self-regulating, or runaway, condition. Unloading of the impeller during surge caused the speed to increase; this acceleration caused the surge severity to increase. The slightest disturbance to the compressor, the surge control system, or any upstream or downstream equipment would cause the unit to accelerate so rapidly that it could go from normal to self-destruct speed in less than a second. We recorded accelerations as high as 2000 rpm per second, and that was during controlled tests. I quickly decided that the only time I liked positive feedback on my job is in my performance review and salary increase. The whole experience convinced me of *Surge Control Fact of Life #2.*

Understanding the problem, as usual, brought me reasonably close the solution. In this case, since the derivative switch was so good at anticipating an overspeed trip for transient recording, I tried using it for control. A second voltage switch was added to the output of the derivative module. The output operated on-off valves that removed sources of power from the compressor and relieved the surge as acceleration approached the point of no return.

The original valving system had to be modified to obtain stroking times less than 0.25 second. The conventional 3-way solenoid valve was adequate for blocking and venting the I/P signal. The problem was in exhausting the large-volume actuator on the main surge throttling valve. To achieve the required speed, we installed a double-acting piston-operated ball valve, increased tubing size, and drilled vents in the actuator on the opposite side of the diaphragm from the signal pressure; we also had to add a high capacity volume booster and air set. Figure 2-1 shows the general arrangement (Ref 1).

Table 2-1. Surge Control Facts of Life

1. Using damping in a transmitter for surge control is like walking across the Autobahn wearing blinders—suicidal.
2. You are either into fantasy or high speed recordings for surge analysis.
3. Concentrating on the speed of the controller and ignoring the response of the control valve in a loop's reaction to surge is like starting a race in a sports car but finishing in a tractor.
4. If it is important enough for you to specify a valve throttling speed or precision, you need a position transducer or smart positioner with readback so you have a way of knowing whether you are getting what you want.

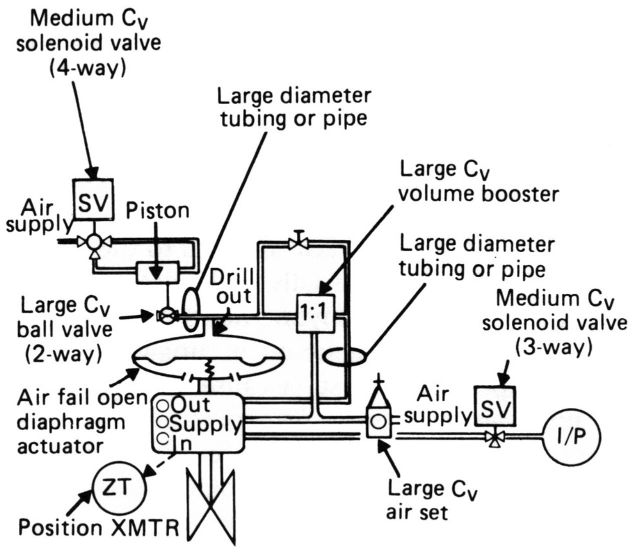

Figure 2-1. Surge control valve accessories for fast throttling and interlocking.

5. All engineers make mistakes—some of which may take years to surface; the worst mistake is not to be around to help fix the others.
6. State that you mean *throttling* rather than *stroking* time when you express the speed requirement on a surge control valve specification—unless you really want on-off rather than modulating action.
7. If an experienced technician is serious about a potential problem, you'd better listen unless you like looking like a fool.
8. A step change in the valve signal doesn't prove anything but your ignorance during a valve throttling test; the signal should be ramped at a rate similar to that expected from the controller.
9. Don't omit an open-loop backup from your surge control system unless you are planning to replace the compressor anyway.
10. Expect the unexpected in terms of disturbances to a surge control system.

One of the many things I don't understand in life is why anyone buys a surge valve without a position transmitter or smart positioner with readback. Have you ever tried to decide whether valve prestroke has deteriorated from 0.2 to 0.4 second? The stopwatch and eyeball approach just doesn't give you the tenth-second resolution that can mean the difference between having a compressor intact or scattered to the four winds. Some instrument experts are adamant about using fast digital devices to reduce the delay at the controller from 0.25 to 0.05 second but don't bother to verify the speed of the slowest component in the loop, the surge control valve. Cost consciousness is frequently the culprit. While saving money is a noble goal, false economy is something else. In my own career, the worst mistakes were the result of trying to save money the wrong way. For surge control, a $500 valve position transmitter determines whether a $20,000 valve can adequately protect a $10,000,000 compressor. If control valve dynamics were constant, the problem would not be so bad. But the speed of a surge control valve deteriorates with age. The prestroke delay and stroking time should be measured regularly as part of a good preventative maintenance program. *Surge Control Facts of Life #2* and *#4* put this issue into perspective.

This compressor project also afforded opportunity for some innovative flowmetering. The first stage of suction could not be used for conventional metering because flow was through a tapered square duct. Moreover, the usual requirement for a metering run 10 pipe diameters upstream and 5 pipe diameters downstream of the second stage primary flow element was waived because the pipe was too large and the plot plan too tight. To resolve these problems, I specified an averaging pitot tube pointed into the outlet of a silencer on top of the tapered duct. The silencer provided flow straightening for the first stage of suction. In addition, I specified straightening vanes close-coupled to a long-cone flow tube for the second stage suction pipe. All seemed well until the first stage assembly arrived. It was fabricated of thin wall tubing that would probably not survive surge-induced vibrations. All I needed was to have a mass of tubing sucked into the compressor. I designed a strong-wall duplicate of the assembly. The result was a piece of modern sculpture that I might be peddling today, if it hadn't worked.

There were also mistakes made during design and construction. The most embarrassing was the scrambling of the first out sequence alarm nameplates. Actual alarm conditions did not match the nameplates, leading to some fascinating explanations as to the causes of shutdowns.

I stayed on the job after the control problems were solved. One morning, I was roused at 5 am to spend an hour on the phone explaining the function of the derivative module—because it had been incorrectly blamed for several successive shutdowns. The technician then did such a good job with the system that he earned a promotion and I got my return ticket to St Louis.

I never completely deserted the facility, due partly to loyalty and the rest to awe about the unforgiving nature of compressor dynamics. Subsequently, I helped the plant people design a computer system to generate surge maps on-line and worked on the beginnings of an expert system to guide operators in diagnosing compressor systems.

This continuing involvement taught me how much the relationship between on-site and headquarters engineers could be improved by just staying in touch with the plant and being available for follow-up assistance. Some of the problems brought to my attention were real. Others were imagined. But there was little question about the need for reassurance that I would back up my work—that the installation was not just a passing fancy. After all, the plant had to live with the system. Most members of Central Engineering Departments need to be reminded periodically of *Surge Control Fact of Life #5*.

An upsetting experience

At one point in my career, I thought I had quite a combination of theoretical knowledge and field experience with fast loops. I was sure I could explain almost any observation and avoid almost any problem. In short, I was dangerous. Luckily, I ran into a surge control problem that brought me back to earth without causing any permanent harm.

I was asked to help design a control system that would prevent surge if any of six reactors fed by a common compressor tripped. When a trip occurred, the valve between the reactor and the compressor would stroke closed within 3 seconds. In the real world, this is practically a step change. But I thought the disturbance could be handled by an electronic analog feedback controller with anti-reset windup if the transmitter and main throttling control valve were fast enough. After all, any good textbook could tell you that a surge controller is really just a flow controller with a remote set point (Ref 2).

How do you make an 18-in. throttling butterfly valve fast and precise? It's like trying to make a giraffe do the fox-trot—even if you succeed, you are not going to win any dance contests. The actuator is too large and the linkages too sloppy for any finesse.

You must get started properly. The surge valve specification should state *throttle* instead of *stroke* for the time requirement, as suggested by *Surge Control Fact of Life #6*. Otherwise, the vendor will meet your specification with a solenoid or quick release valve. This will give you good on-off control—which, unfortunately, is not what you want. You'd also better specify volume boosters, large tubing, and high capacity air sets to be sure of fast throttling. And, you've got to keep on top of the details—since getting a vendor to make a large throttle quickly is like getting a gorilla to peel a banana in polite society—each has higher priorities than fine points of behavior.

After you get the specification correct and the valve built, you need to run some tests. I have never seen a surge control valve satisfy the specifications on the first try. There is always at least some rework of the accessories or air piping. Unless you like embarrassment, it's wise to do the testing at the valve factory. Manufacturers seem more interested in helping rectify problems if it means getting iron out the door. Also, testing at the job site puts you at the favorite hangout of such nervous types as project managers.

On this project, I didn't get a chance to do factory testing. And our plant tests showed that my fears were not unfounded. The valve was far too slow. Even after reworking some of the air piping we were off by a factor of two.

A plant technician stopped by for what we thought was a laugh at the hotshot from Central Engineering. He said the valve wouldn't work because it didn't have a positioner, although he couldn't explain why. I knew from my thorough knowledge of the literature that volume boosters were recommended instead of positioners on fast loops. In fact, one elegant theoretical study showed that even with stick-slip, performance went down hill when a valve positioner was added. I also reasoned that a positioner created a cascade system in which the inner valve position loop was slower than the outer flow loop—so the outer controller would have to be detuned for stability. All this was theoretically correct. Still, I should have paid more attention to the technician and less to theory. I recommended *Surge Control Fact of Life #7* to all instrument engineers.

We didn't have any fun the morning the surge valve was put in service. The flow transmitter showed the impending surge and the controller asked the valve to open. The valve responded by doing the worst possible thing. It slammed shut before the forward flow to any of the reactors had been established.

The technician who wanted the positioner took me to the surge valve and showed that he could move it to any desired position by tugging on its stem. Obviously, the buffeting action of the turbulent flow could cause the disc to wander and eventually close. The actuator size was checked and found to be adequate; the spring rate was increased but the results remained the same. Subsequent tests showed that the stem resisted movement somewhat better if the actuator was fed directly from an I/P transducer, and that it could not be budged at all if a positioner was installed.

Process of elimination led us to the high outlet sensitivity of the booster port as the cause of the problem. Upward force on the valve stem causes the actuator diaphragm to compress the top air enough to open the exhaust port. The exhaust causes the stem to move upward, causing the booster to relieve more, and so forth. It could happen in either direction. Positive feedback had got me again.

I subsequently discovered another incentive to add a positioner to a valve with a booster. Boosters have large deadbands, which create prestroke delays inversely proportional to the rate of change of the pneumatic signal. The deadtime doesn't show up in factory tests because manufacturers usually use step changes in the I/P signal. It does show up in service because the feedback controller usually makes gradual changes in the I/P signal. The prestroke deadtime requires the controller gain to be decreased; this makes the changes slower, increasing the prestroke deadtime. The snowballing effect causes serious deterioration of loop performance. *Surge Control Fact of Life #8* should be remembered when the throttling speed requirement is written on the valve specification.

Having discovered all this, I meekly agreed to install a positioner. Now, a positioner in series with a booster creates another instability problem. The volume seen by the positioner is much smaller than that seen by the booster. The result is a high frequency limit cycle that can only be killed by bypassing part of the air from the positioner around the booster through an adjustable restriction. This explains yet another detail in Figure 2-1.

Now that some of the air was being bypassed around the booster and the controller was detuned to accommodate the inner positioner loop, the valve was even slower. A fast transmitter was found, but the maintenance department vetoed the purchase to avoid having a new brand to stock and service. This was just as well because the fast transmitter would have amplified the flow measurement noise.

The end result was that my simple fast flow loop was not especially fast and could only prevent surge for a one-reactor shutdown. A multiple-reactor shutdown caused surge. What's more, I proved once again that after surge begins, recovery by feedback control is not likely.

To salvage the project, we added a preprogrammed open-loop backup surge control system. This calculated the surge valve opening needed to offset each reactor trip. By this time, the system was nowhere near as simple as originally intended. But it worked. This reminds me to state *Surge Control Fact of Life #9*.

I later had the opportunity to factory-test a fast throttling butterfly valve for another project.

Wanting to share my experience with mankind, I proceeded to operate the butterfly valve by manually moving the stem. Reaction in the room went from amazement at my Herculeon strength to bewilderment that anyone could do it. In the years since, I have not found anyone aware of the stability problem when boosters are used on diaphragm-actuated butterfly valves without positioners.

What—me worry?

I was called in to help with another compressor control system, on which the project team said there was no surge problem. By this stage in my career, I had been burned enough to know that I shouldn't believe them.

The reactor trip block valve was designed to close in 145 seconds, which the team thought was long enough for surge to be handled by a feedback controller. I overlayed the stroking time on a plot of the installed characteristic as shown in Figure 2-2. The flatness of the curve at the upper end of the stroke and the slowness of operation at the lower end left little time for surge control. To be precise, there was 1.7 seconds between the crossing of the feedback set point and the open-loop backup set point and the surge point. Suddenly, surge control appeared to be a problem. The open-loop backup was clearly needed, and the surge control valve would have to stroke in less than a second.

Another system like that in Figure 2-1 was installed. This worked well and prevented surge even on the several occasions when the downstream valves were accidently closed in a couple of seconds.

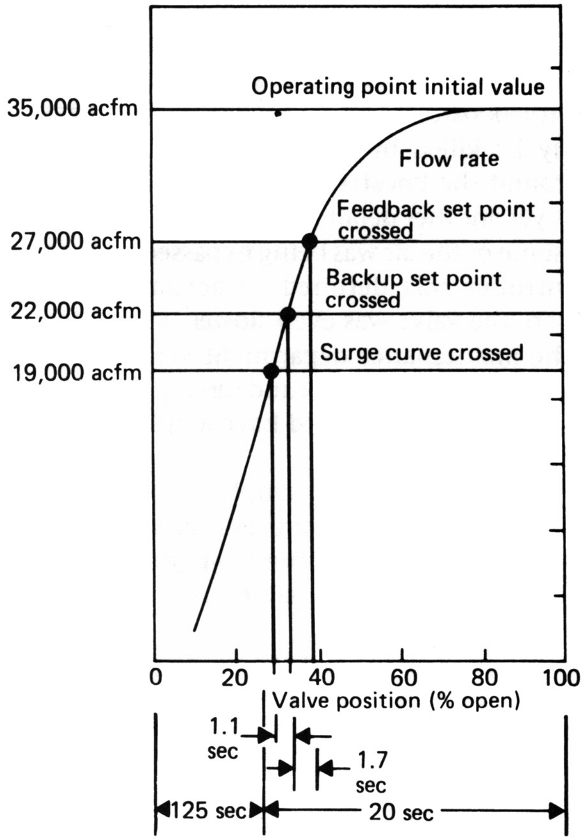

Figure 2-2. Installed flow characteristic for a block valve, indicating the short time available for operation of the surge control system.

The memory of these exciting disturbances, however, leads me to state *Surge Control Fact of Life #10.*

KEEPING YOUR HEAD

Surge control can be as nerve racking as driving the family sedan on the freeway during rush hour or as exciting as driving a Porsche at Le Mans. It all depends on your equipment and attitude. Surge control problems are solvable once you recognize the time frame, meet the speed requirements, and keep an open mind as to what might suddenly pop up on the horizon.

REFERENCES

McMillan, G K; *Centrifugal and Axial Compressor Control*, Momentum Press, 2010.

Nisenfeld, A E; *Centrifugal Compressors—Principles of Operation and Control*, Monograph Series, Instrument Society of America (Research Triangle Park NC), 1982.

Pressure Control: Without Deadtime, I Might Be Out of a Job

Have you ever watched a chart recorder connected to a process pressure transmitter(as the pen hit the peg so hard that ink splattered out the top? If so, you probably wondered what happened to that textbook world where pressure loops were so easy they weren't worth discussing. You are not alone if you have encountered tough pressure control problems. I have run into a few and am still here to talk about them. And what's a few years lopped off a life expectancy, anyway?

THE DEADTIME THAT SHOULDN'T BE

I survived my pressure control problems thanks to two words of wisdom I got when I was 4 years old and sitting on my daddy's knee. The secret passed from father to son on what turned out to be that fateful day was *deadtime*. Of course, it took me 34 years to understand what he was getting at. The intervening period has been filled with math, whose principal function seems to have been to keep me from realizing that deadtime determines loop performance.

Things aren't all bad, of course. If there were no deadtime in a loop, I might be out of a job.

Why is there error?

Think of it this way. Without deadtime, any plant engineer who didn't know about instrumentation could buy a proportional controller and set its gain all the way up. The result would be perfect control, with no error. In real life, this won't work. Loops will show peak error proportional to deadtime, and integrated error proportional to deadtime squared (Ref 1). So, somebody like you or me, who understands both the process and the instruments, has to identify the major sources of deadtime and reduce their contribution to the total. This is likely to involve changing the process equipment, the controls, the operating conditions, or all three.

I don't feel too bad about taking so long to recognize the facts of life about pressure control—and you shouldn't think your life has been wasted until now, either—because many high-powered widely published experts still don't know what it's all about. For instance, some highly paid well-regarded instrumentation specialists still claim to have designed compensators that cancel out the effect of deadtime in control loops. Well, for unmeasured disturbances, there is no fancy algorithm that will violate physical laws and shorten the deadtime interval. This leads us to *Pressure Control Fact of Life #1.*

Table 3-1 The Pressure Control Facts of Life

1. Anybody who claims to have an algorithm that eliminates the effect of deadtime for unmeasured load disturbances should seriously consider a career in science fiction.
2. Anybody who thinks all pressure control loops are easy has spent too much time in the office and not enough time in the field.
3. The instruments are the source of deadtime in pressure control loops. The control valve walks away with top honors as the most generous contributor.
4. The powerful microprocessor-based controller in your distributed control system can do more harm than good in a

pressure control loop unless it has a faster execution rate than the default rate.

5. Slow transmitters are great if you want your relief devices instead of your trend recorders to tell you that you have poor control.
6. Open-loop tuning methods provide a good test of furnace pressure ratings.
7. If plant people can't relate to you, don't expect them to relate to your control system.
8. The classic decoupler between furnace air and pressure control loops can cause inverse pressure response.
9. Distributed system consoles are great for playing starship commander as long as you can be beamed to an analog controller and recorder for an exceptionally fast pressure control loop (see April 2, 2007 entry "Analog Control Holdouts" on the website http://www.modelingandcontrol.com/
10. Electrical variable speed drives can prevent you from blowing the relief devices (and your mind) on a difficult loop.

Where does it come from?

Where does the deadtime come from in a pressure control loop? Don't look in the process. The textbooks were right about one thing. Processes have negligible pressure response deadtime, even when units include multiple volumes separated by ductwork.

The deadtime comes from the instruments—devices the textbooks always assume are ideal, and that you and I prefer to think of as solving rather than creating problems. As usual, there is good news and bad news. The good news is that an instrument practitioner should be knowledgeable enough to specify devices that minimize deadtime. The bad news is that when you try to explain this to a project manager who doesn't even think about, let alone understand, instrumentation, your planet of origin will be questioned. You can back down and claim the lead content is up in the water again and you were just hallucinating, which will please the person holding the purse strings but will leave you holding the bag. Or you can shame your boss into acquiescence with *Pressure Control Fact of Life #2*, then back it up by relating the following experiences.

FURNACE CONTROL IN THE NETHERWORLD

I was invited to a plant that had an electrical furnace with pressure control problems caused by large variations in the rate of internal gas generation. To hold the furnace pressure within safe limits, the unit had been designed with a positive displacement compressor, which pulled out gas and directed it to a flare. A pressure control loop was supposed to keep this system operating properly. Unfortunately, changes in gas generation rate created disturbances that were so large compared to the volume, it was like trying to regulate an explosion.

The unit was an ominous black structure, like what you'd imagine provides heat for the netherworld. Everybody at the plant held this unit in great awe. I understood why as soon as I saw it. I understood even better when we climbed to the top for an inspection, because the monster spewed sparks and fires broke out where we had just stepped. I noticed what appeared to be a ceremonial ring of filled bathtubs around the circumference of the furnace. The ceremony, it turned out, involved everybody jumping into the nearest tub when the water seals around the furnace blew out, releasing vapors that ignited upon contact with air. I don't mind a dip now and then. But the water was dirty, and I was never fond enough of the nursery rhyme "Rub a dub dub" to think it worth a try. So when I was asked to stay and wait for the next belch, I suggested we view it from the control room.

As soon as we reached the comparative safety of the operator's panel, I noticed the pressure trend recordings. Vertical spikes every few hours confirmed that the control system was not doing the intended job.

The control loop

The pressure control system was simple enough. An electronic pressure transmitter was connected via sensing lines to the top of the furnace. This provided signals to an electronic analog controller, which operated a control valve to bypass around the compressor. When the control valve opened, more gas recirculated from the compressor discharge to the suction so less was pulled out of the furnace and sent to the flare.

What could be done to improve performance? For a test, we placed the controller in the manual mode and found that the pressure ramped offscale within a few seconds. We considered any further experimentation with the furnace to be risky, so I decided to use a dynamic simulation program to test some ideas.

Sources of deadtime

The simulation showed that despite the multiple volumes and resistances associated with the cyclone and spray tower between the furnace and compressor, process deadtime in the open-loop response to a gas generation disturbance was negligible. The deadtime appeared when the loop was closed, however, so the fault was obviously with the instruments.

The large pressure control valve was the biggest culprit. The total deadtime of the valve, from the prestroke lag associated with filling or exhausting the pneumatic actuator, the stroke deadband or backlash associated with packing or seal friction and slop in linkages, and the remotely mounted current/pressure transducer approached one second. A second may not be much in some applications, but you've got to compare it to the process deadtime. It was apparent that a tremendous decrease in error could be achieved with a fast control valve. Your ability to solve a pressure control problem accordingly hinges on remembering *Pressure Control Fact of Life #3.*

An analog controller added essentially no deadtime. A microprocessor-based controller, with some fancy algorithms that some high-tech genius had suggested as the solution to this gritty problem, would have added a deadtime about equal to its sampling interval; this would be as appropriate in this pressure control loop as a maitre d' in a cage of gorillas at feeding time, per *Pressure Control Fact of Life #4.*

The simulation showed that the transmitter added a deadtime about equal to its time constant. The existing device was therefore replaced by a strtain gauge transmitter to obtain the fastest practical response. In replacing the transmitter, we recognized that the faster element would make the trend recording look worse for the same performance because it would provide less signal damping than the original instrument. Therefore, for evaluation, the faster transmitter would be connected to the controller, but the original transmitter would be connected to the trend recorder. Keep *Pressure Control Fact of Life #5* in mind whenever you play with the damping in your transmitter.

Runaway disturbances

The simulation also showed why disturbances caused pressure runaway. There was a process time constant, but the response curve did not bend over until the pressure was well beyond the furnace trip limit. Within the pressure control range, the pressure response therefore looked like a ramp, so the process was not self-regulating. Since furnaces have pressure control ranges in inches of water column, it doesn't take much of a disturbance to cause an excursion to race out of the control band like a runaway roller coaster.

Tuning

The high pressure gain causes the observed pseudo-integrating response with a small window of allowable proportional controller gain. This makes tuning tough (Ref 1). Most of the published tuning methods are open-loop techniques that work great for textbook curves. But, we couldn't run this process open loop for more than a few seconds. (Remember what happened when we put the controller in manual?)

Since open-loop textbook tuning methods would have provided more fun than engineers are allowed to have, per *Pressure Control Fact of Life #6*, we dusted off some old closed-loop approaches. Ziegler-Nichols ultimate oscillation tuning worked best, as long as I was ready to quickly reduce the controller gain when the oscillations started to grow too fast.

The new system worked as we had hoped. This is seen in the "before" and "after" recordings of Figure 3-1. The plant people were willing to maintain the new instruments. Part of this was because the system worked. Part was because we had developed a good working relationship with each other. This reminds me to state *Pressure Control Fact of Life #7.*

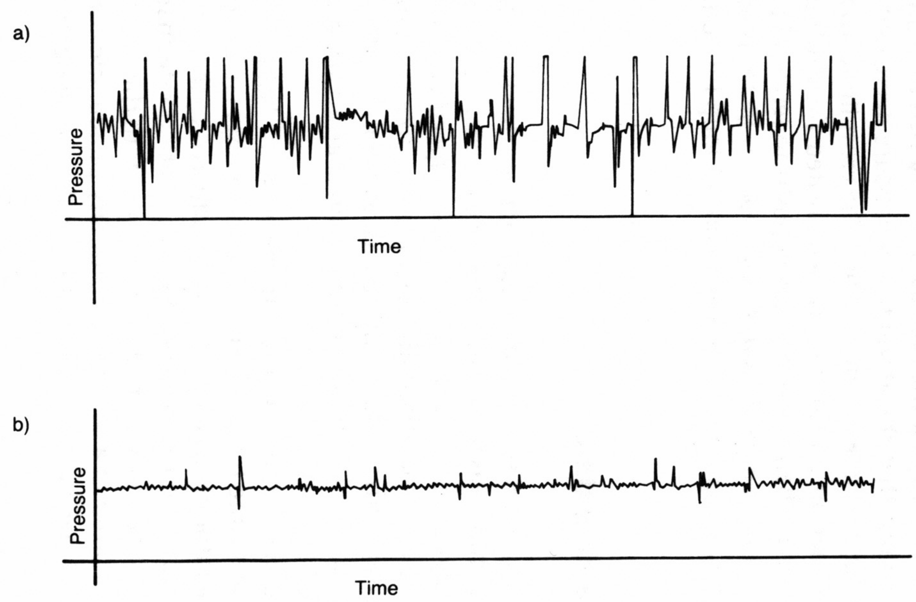

Figure 3-1. Pressure trends in the electrical furnace a) with the original control system, b) after reducing deadtime by replacing the control valve and transmitter with fast-response elements.

JOINING THE FAN CLUB

My success with the first furnace pressure control problem brought me the honor of being asked to help with another. This was a waste incinerator, which also had large pressure disturbances caused by changes in the gas generation rate. The incinerator was tripping out regularly because of wild pressure excursions—an unacceptable condition because prolonged outage forced down the production unit.

The control strategy

In this plant, the final control element was a centrifugal fan, whose speed was manipulated to regulate furnace pressure. The loop was implemented on a distributed digital control system with a fast 1/10-second sample time. Someone had taken advantage of this sophisticated system to program a decoupling strategy that broke the interaction between the combustion air flow and furnace pressure control loops. The decoupler summed the outputs of the air flow and pressure controllers so that opening the valve letting air into the furnace would increase the speed of the fan sucking gas out of the furnace.

The decoupler was certainly clever. Unfortunately, simulation and field tests showed that the increase in air caused the temperature to drop so fast that the absolute gas pressure decreased until the fuel flow caught up with the air flow. The cross limits to make sure air leads fuel on a load increase aggravated the problem—because the furnace had no soot or ash deposits to retain heat and retard temperature change. *Pressure Control Fact of Life #8* should be consulted before designing the pressure control loop for your furnace.

Looking for deadtime

After my previous pressure control experience, the first place I looked for deadtime was in the final control element. A conceptual review of the dynamic response of the variable speed drive on the fan showed it to be a velocity-limited exponential, like a pneumatic control valve, but with negligible stick-slip, wide rangeability, and no deadtime associated with air flow into or out of an actuator. Shop tests confirmed these conclusions.

One second of deadtime made a difference on the last job. Here, we had even less time to play with. Simulation tests later showed that the incinerator volume was so small compared to the changes in waste fuel flow and the associated gas generation rate that the pressure could go from its normal operating point to the trip point in less than a quarter of a second. An ordinary pneumatic control valve would not have even gotten out of its deadband before the furnace tripped.

The actual pressure response was as fast as expected, but you wouldn't have known it from the graphic displays. The distributed control system console proved useless. The minimum update was 5 seconds. It was exciting to watch numbers with three decimal places flashing without providing the slightest hint as to where the pressure had been or where it was going. The trend display showed low frequency aliases that would be disastrous if used for tuning. At this point, I searched the arm of my chair for an ejection lever. This leads me to *Pressure Control Fact of Life #9*.

The solution

Even after we eliminated the decoupler and tuned the controller, the furnace would trip out every few days. We therefore replaced the digital controller with an electronic analog device. The result was greatly improved control and virtually no tripouts, as shown in Figure 3-2.

You may now ask questions:

You: How can a tenth of a second make such a difference?

Me: In a tenth of a second, the pressure could be halfway to its trip point.

You: Isn't the variable speed drive slower than the digital controller due to inertia?

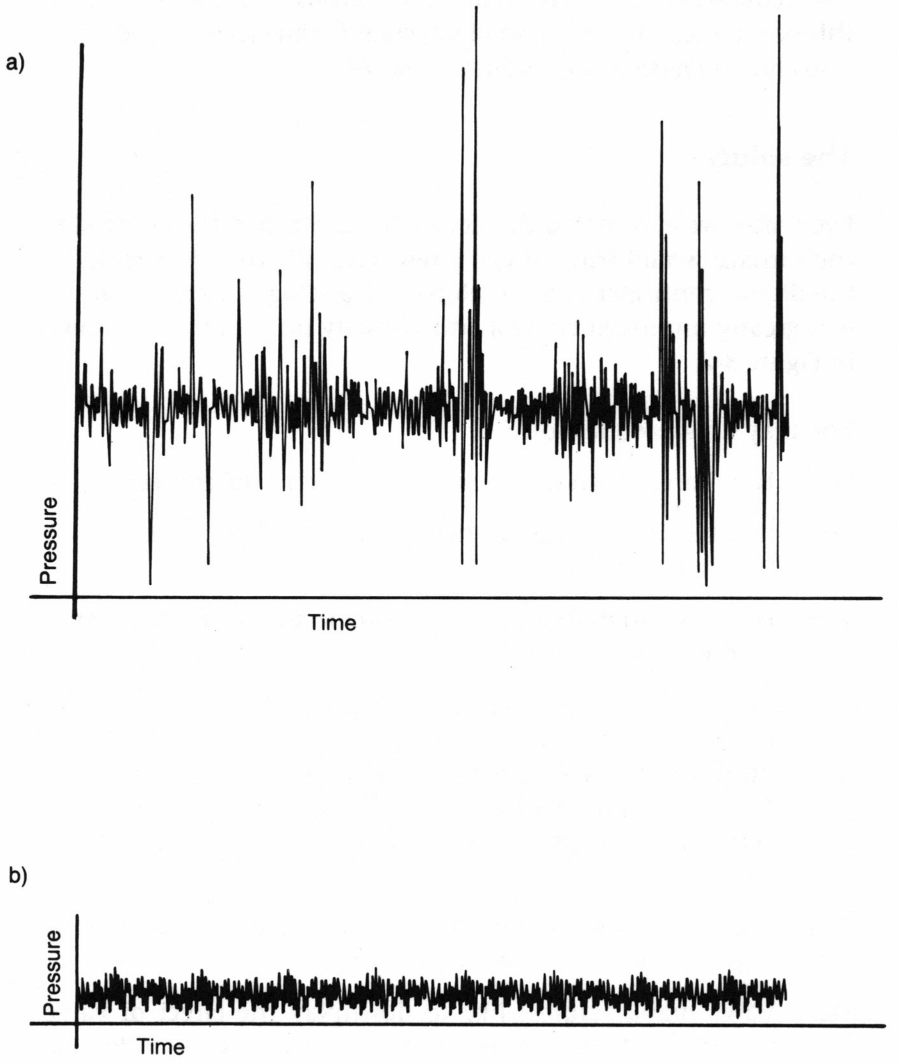

Figure 3-2. Pressure trends in the waste incinerator a) with the original control system, b) after eliminating the decoupling algorithm and reducing deadtime by replacing the digital controller with an analog element.

Me: The deadtime from inertia is negligible, and the high gain of the furnace pressure response means the fan speed only needs to change a percent or less to correct for an excursion. The time required for a very small speed change is therefore very small—confirming *Pressure Control Fact of Life #4* and adding *#10*.

You: Why aren't variable speed drives considered more often to improve the control of difficult loops?

Me: Valve manufacturers like to downplay the effect of valve stick-slip, drive manufacturers don't understand *Pressure Control Fact of Life #3*, and users tend to select among the alternatives handed to them on silver platters rather than to look for uncommon hardware.

REFERENCE

1. **McMillan, G K;** *Tuning and Control Loop Performance*, Momentum Press, 2010

Advanced Control Algorithms: Beware of False Prophecies

I have been to a lot of supposedly Advanced Process Control Symposia. In attending these supposedly high-level programs, I have seen more new predictors, observers, and compensators than anyone could possibly want, need, or evaluate.

The people who develop these algorithms must find it exciting to perform mathematical magic, prove the results on simulators predestined to do just that, then talk or write about their achievements. They certainly can't be getting much satisfaction from the widespread disinterest on the parts of practitioners who daily deal with industrial realities. It is probably unfair, but processing plants simply don't behave like simulations where everything is as the designer says it should be.

BETTER . . . THAN WHAT?

When I was wetter behind the ears than I am now, I found myself being impressed by some simulated proofs of some new algorithms. Then I began to understand how the dice were loaded.

I remember one case in which a researcher showed that his technique was better than conventional control. Unfortunately, it seems he didn't properly tune the conventional controller, neglected derivative action, omitted load disturbances, and used integrated absolute errors for the performance comparison. The advanced algorithm, which provided a nice slow overdamped response, was therefore being compared to a poorly tuned PI controller for set point changes.

In the processes I encounter, there are load disturbances, PID loops are necessary, and the controllers are tuned for minimum peak error or rapid recovery—fast underdamped responses. Numerous studies have shown that nothing can perform better than a properly tuned PID controller for unmeasured load disturbances (Refs. 1-3). This, in fact, turns out to be a secret of algorithms like Dynamic Matrix Control (DMC) and Internal Model Control (IMC): the tuning constants can be calculated rigorously in these algorithms, whereas they can only be approximated in traditional PID control.

The minimum peak error

No matter what algorithm is used, there is a lower limit to the peak error for any single loop. This limit is the excursion of the controlled variable during the loop deadtime for unmeasured load disturbances. The controller cannot do anything until it has seen the change, reacted to it, and manipulated the final element to cancel its effect. No matter where the load disturbance occurs, no correction can take place until after one loop deadtime, because the signal must make a complete pass through all the elements before any action is taken. In practice, a conventional PID controller achieves an output that effectively begins to reduce the error after 110% of the loop deadtime (Ref 4).

Attacking the real problem

If you really want to improve controller performance, locate the source of loop deadtime and reduce it. That's what instrumentation is all about. Designing yet another new algorithm to describe at a symposium is an exercise in fantasy—played in the language of algebra, calculus, and matrix arithmetic.

From the viewpoint of deadtime, some advanced control algorithms make a fatal mistake. They increase loop deadtime by inserting a significant sampling or calculation interval into the response path. Since the delays are additive, this isn't exactly the wisest thing to do. You wouldn't think the experts would make such a simple mistake, would you? Let me set you straight.

Stochastic control

Some of our plants have developed considerable expertise in stochastic control. The procedures are elegant and firmly founded on formidable mathematics.

In one case, I had heard that an optimizer was even being used to predict coefficients for the stochastic optimizer from test data. I was impressed. Then the plant ran into trouble using the system for relative viscosity control, and I looked into the details.

A trend recording of the response to unmeasured load disturbances pointed to the problem. The oscillation period shown on the charts corresponded to a deadtime about twice that expected from open-loop response plots. This was a mystery, until I found out how long it took to compute the stochastic algorithm. You guessed it . . . the calculation interval was about equal to the original loop deadtime.

The plant didn't want to replace its fancy algorithm with a plain old PI controller, but they were willing to try a form of an IMC. This did much better. As you might have surmised, it had a relatively short calculation interval. I can only imagine how well conventional PI would have done.

Advanced pH control

An even more horrible example was set by a microprocessor-based pH receiver-controller that one of our plants was trying to use on a particularly difficult loop. The controller had some nice bells and whistles, such as a *touch* calibration with a numerical keyboard and graphic loop linearization based on a segmented titration curve. It also had an advanced control algorithm that supposedly compensated for loop deadtime. This algorithm went to manual after each control action and waited one deadtime before it went back to automatic for the next control action.

The idea behind this algorithm was that the controller would not try to react before it could see the results of its actions. Sounded good. It even worked well for set point changes and feedforward action. Unfortunately, it doubled the loop deadtime for unmeasured load disturbances.

I remember watching anxiously as huge concentration disturbances in the influent caused the pH to drop off the bottom of the scale, while the algorithm calmly waited in manual to see the effect of its last control action. At times, I got so impatient watching the controller wait that I manually manipulated the output.

After several fun-filled days trying to tune the controller—whose adjustments had some creative names and explanations—I convinced the plant to switch back to conventional analog PID feedback control. The old-fashioned analog controller didn't wait when it saw the effects of a disturbance start to upset the process. Instead, it responded as quickly as its inherent deadtime would allow to keep the pH within bounds.

Some checkpoints

In spite of any doubts I may have instilled in you, you may still want to try someone's advanced algorithm to solve a tough problem. Before you do, at least make sure you have already done all you could to reduce the loop deadtime by changing instruments, equipment, or operating conditions.

Table 4-1 Advanced Control Algorithm Performance Checkpoints

1. Did the algorithm add an appreciable amount of deadtime to the loop? If so, forget it.
2. Did the algorithm perform well for unmeasured load disturbances? If not, forget it.
3. Was derivative action used in the conventional algorithm it was compared against on a slow loop? If not, your comparison was unfair; add the derivative mode and try again.

4. Was the PID controller tuned with a reputable method, such as the Ziegler Nichols closed-loop or reaction curve method with the controller gain halved to suppressed oscillations? If not, the comparison was unfair; tune the PID controller properly and try again.

After these steps, you may still have an insatiable desire to use something more sophisticated than PID control. If so, compare the performance of any proposed algorithm with a PID controller, based on the check points in Table 4-1.

Assuming that your high-tech algorithm passes the check points, please do one more thing. Make sure that the controller being used to implement the fancy strategy is self-tuning—initially and whenever loop gains or dynamics change. No matter how good the algorithm is supposed to be, it won't work unless it is well-tuned. And, surely, nobody is left outside the funny farm who would willingly volunteer to learn a new set of tuning adjustments, then use them manually every time the plant started up or was subject to a change in load.

I'd better stop here, or we shall find ourselves debating whether or not this is 1986 (which, of course, we are).

REFERENCES

Bohl, A H; McAvoy, T J; "Linear Feedback vs. Time Optimal Control II, The Regulator Problem," *Industrial and Engineering Chemistry, Process Design and Development*, Vol 15 No 1, 1976 p 30-33.

McMillan, G K; "A Performance Comparison between a Smith Predictor and a PID Controller for Control Valve Stick-slip and Process Non-Self-Regulation," ACC, 1983, pp 343-348.

Rivera, E; Morari, M; "Internal Model Control: PID Controller Design," 1985.

McMillan, G K; *Tuning and Control Loop Performance*, Momentum Press, 2010.

Direct Temperature Rate Control Improves Reactor Yield

Exothermic batch processing applications may offer opportunities to optimize yield and other performance characteristics by exercising direct control over the rate of temperature change. Why? Reaction rate generally increases with temperature. Raising the rate of temperature change allows more of the internally generated heat to go back into driving the reaction instead of being removed by coolant. A limit is necessary, however, to avoid loss of control and a runaway reaction. This could occur because above some maximum rate, the positive feedback time constant of the reaction approaches the negative feedback time constant of the heat removal system. The lower allowable proportional band on the controller then approaches the upper value—and the window of possible controller adjustments closes (Ref 1).

OBTAINING THE IMPROVEMENT

Laboratory research on one process showed that if the temperature of a fed-batch reactor were ramped upwards at 0.6 degree/minute, the yield could be optimized at a level 20% over current practice. Below 0.6 degree/minute, the yield decreased; above 0.8 degree/minute, the reaction rate increased uncontrollably—and the 0.2 degree/minute margin was considered necessary for safety.

Set point ramping

The plant first attempted to regulate temperature rise rate using a programmable logic controller (PLC) to ramp the remote set point of an analog feedback controller. It didn't work.

An obvious problem was that analog-to-digital (A/D) converter resolution error and integer arithmetic error in the PLC produced a staircase instead of a smooth ramp for the remote set point. Control action was therefore erratic.

Direct rate control

Even if a smooth function could have been generated, there was a fundamental problem with the set point ramp concept. Ramping the set point establishes a profile but does not necessarily ramp the controlled variable. Disturbances that cause a deviation from the programmed profile are corrected by raising or lowering the rate of temperature change to get back onto the original curve. To maximize reactor yield, it was essential to keep the rate of temperature change *constant* rather than enforcing a particular temperature at a given time into the batch cycle.

It was therefore decided to develop a strategy with a fixed rate of temperature change as the set point and the derivative of the temperature measurement as the controlled variable. This might have been difficult in the analog days, but microprocessor-based controllers offer all the capabilities anyone could possibly want or need to do the job. For instance, to compute the derivative of the temperature, simply subtract the value at the previous calculation interval from the present measurement and divide the difference by the length of the interval.

First, a simulation

Since the plant already had a bad experience attempting to duplicate the laboratory results in practice, I did a dynamic simulation to verify that the concept was realistic. The simulation confirmed that a strategy based on fixed rate of temperature change could deliver the expected performance. It also revealed some practical problems with noise in the digitized signals and showed that choice of calculation interval, filter size and location, and controller tuning constants would be critical to success.

The noise problem

We selected a microprocessor-based computing controller that used floating point arithmetic and a high-resolution A/D converter. As predicted by the simulator, however, there was still a small staircase effect in the temperature ramp due to the slow rate of change of the signal. This was unacceptable because internal computation of the derivative of this signal amplified the steps.

To reduce the apparent noise, the time interval between successive derivative calculations was lengthened so the A/D converter error became a smaller part of the total temperature change over the interval. This sounds intuitively obvious. However, it is opposite to the general rule in digital computation, which says the calculation time interval should be shortened to increase the accuracy of calculations from dynamic signals.

The noise could have been effectively eliminated if the calculation interval could be made long enough. Unfortunately, there is a practical limit on this approach. The calculation interval injects a pure deadtime into the loop. For the effect of this deadtime to be negligible, it must be much smaller than that introduced by the sensor, thermowell, and other elements in the measurement path.

A filter was therefore needed to further smooth the output ramp. The simulation showed that the most effective location was upstream of the derivative calculation, since the noise frequency there was higher than after the computation. This meant that a shorter filter time constant could be selected for the same overall attenuation—with less effect on the measurement lag.

Figures 5-1 through 5-3 show how lengthening the calculation time interval from 0.02 to 0.05 to 0.10 minutes decreases the noise in the rate signal. The figures also indicate how lengthening the calculation time increases the deadtime.

The optimal filter time constant was found to be 0.05 minute for the measurement signal increasing at 0.6 degree/minute. Figures 5-1 through 5-3 show the rate signal, with such a filter located upstream of the calculation. Figure 5-4 shows the rate signal for a 0.10-minute calculation interval, with the same filter downstream of the rate calculation. Based on data of this type, the optimum calculation interval size was predicted to be 0.10 minute; similarly, the best filter location was shown to be upstream of the derivative calculation.

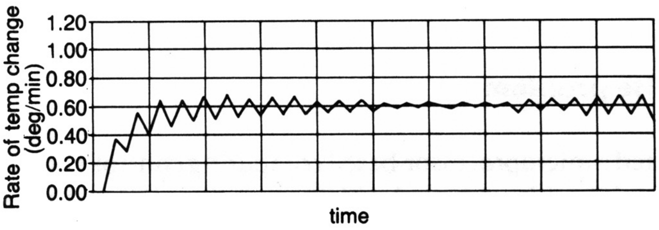

Figure 5-1. Derivative signal with a calculation interval of 0.02 minute and a filter having a 0.05-minute time constant located upstream of the derivative calculation.

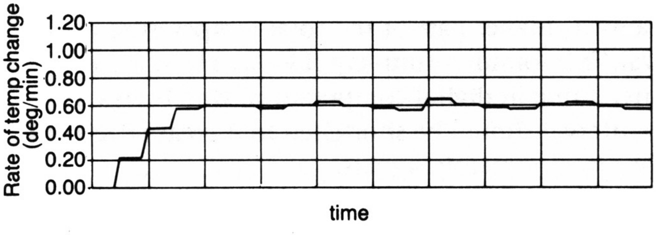

Figure 5-2. Derivative signal with a calculation interval of 0.05 minute and a filter having a 0.05-minute time constant located upstream of the derivative calculation.

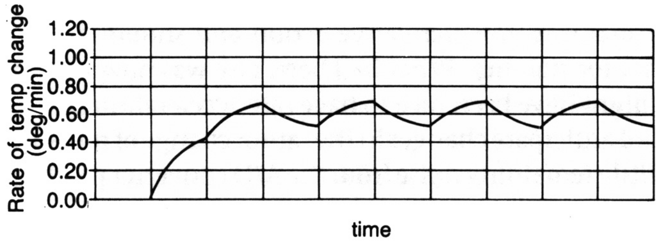

Figure 5-3. Derivative signal with a calculation interval of 0.10 minute, and a filter having a 0.05-minute time constant located upstream of the derivative calculation.

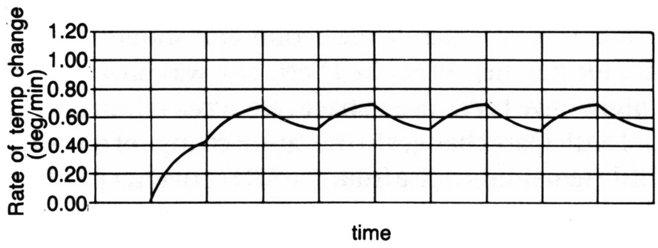

Figure 5-4. Derivative signal with a calculation interval of 0.10 minute and a filter having a 0.05-minute time constant located downstream of the derivative calculation.

Tuning adjustments

Normally, reset action would be small, proportional action large, and rate action somewhere in between for exothermic processes, due to the non-self-regulating response (Ref 1). The simulation program showed otherwise for this system. It indicated that the reset setting should be increased, the proportional setting decreased, and the rate setting turned off. After discussing these unexpected results with my colleagues, I realized the reason. Using the derivative rather than the value of temperature as the controlled variable caused the tuning modes to shift—such that:

Reset Action ➡ Proportional Action
Proportional Action ➡ Rate Action
Rate Action ➡ Rate Acceleration Action

Controller mode settings therefore required a translation from the traditional values. In this light, there was no longer any reset action with respect to temperature. This was just as well, because it was not desirable to enforce a temperature profile with respect to time. What had been reset now acted like proportional action and should be increased to provide the major control effect. What had been proportional was now derivative action and should be decreased somewhat for stability. What had been rate was now rate acceleration. It might have been nice to have rate acceleration action, since this would anticipate changes in the rate of change of measurement, but the little remaining noise from the A/D converter precluded this possibility—so this term was set to zero.

The digital controller also afforded the opportunity to use adaptive tuning to compensate for the change in process gain with batch time. The heat release per pound of reactants decreased and the amount of heat transfer area covered by the material increased as the batch proceeded. Consequently, a given change in chilled water flow caused a greater decrease in temperature later in the batch. For stability, the overall loop gain should be kept constant—by decreasing the controller gain in proportion to the increase in process gain. This was achieved with an adaptive scheme that made the controller gain an exponentially decreasing function of the total reactants charged.

PERFORMANCE

The first batch with the new system was better than the best previous result with the original analog controls. Not only was the temperature ramp smoother, but the chilled water flow was less erratic and total water usage decreased. Figure 5-5 compares temperature ramps, and Figure 5-6 shows chilled water flow rates before and after the microprocessor-based system was installed.

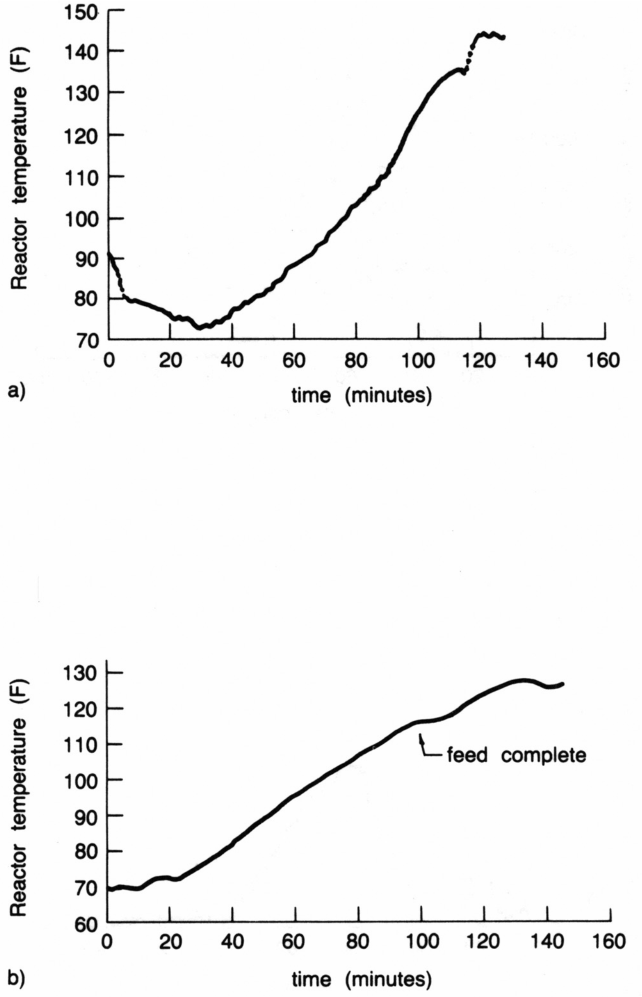

Figure 5-5. Reactor temperature as function of time a) with the original analog controls, b) with the microprocessor-based system and the direct rate control scheme.

Incidentally, the initial calculation interval length, filter size, and tuning parameters were taken directly from the simulation. The values were twiddled after the system was installed but no better settings were found.

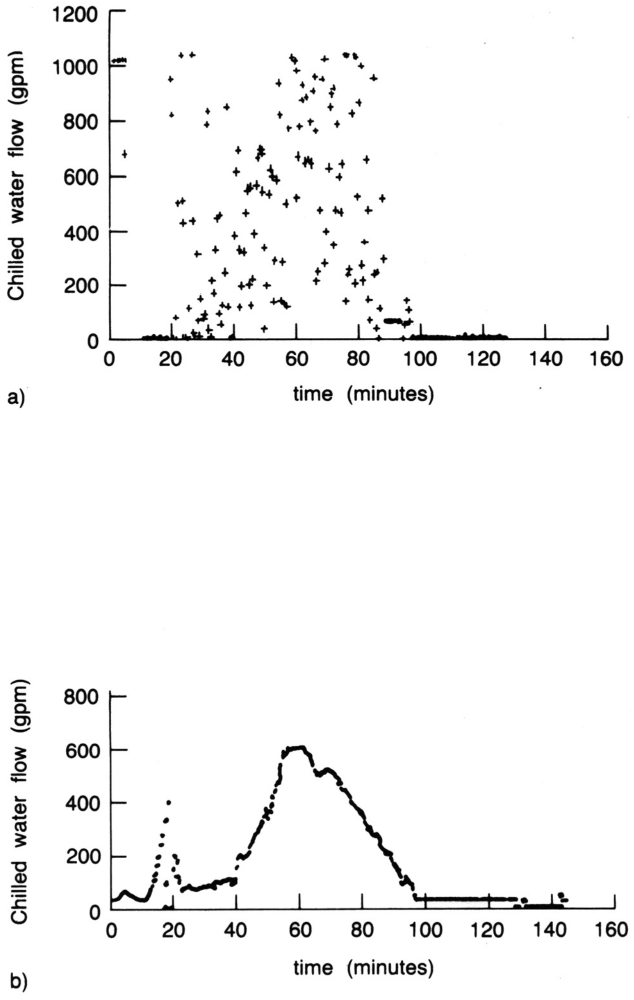

Figure 5-6. Chilled water flow as function of time a) with the original analog controls, b) with the microprocessor-based system and the direct rate control scheme.

REFERENCE

1. **McMillan, G K;** *Tuning and Control Loop Performance*, Momentum Press, 2010

pH Titration Curves—Trick or Treat?

WHO'S AT MY DOOR?

In some ways the chemist standing at your door when you need a pH titration curve is like an overgrown kid standing at your door on Halloween. Not only are both wearing funny costumes that you can't laugh at, both are liable to perform some trickery if you don't give them exactly what they want. I don't mean you should give your chemist candy (although it might not hurt) but you should be prepared to give him exact requirements on how a titration curve is to be constructed.

Titration curves provide ample opportunity for the chemist to unintentionally become a master in the art of deception. An improperly done curve is like a Rorschach inkblot test; your interpretation tells you more about your personality than the process. Such curves become a virtual warehouse of potential blunders. I don't know about you, but I am perfectly capable of making more than enough blunders without any help.

WHY BOTHER?

If pH titration curves are so misleading, why bother with them? Well, if pH loops were easy, you probably could forget about them. Unfortunately, most pH loops are so incredibly difficult that you need the performance edge gained from a correctly done curve.

In order to get a titration curve suitable for use as a design tool, the abscissa should be converted to a ratio of reagent to influent flow. Since most samples are dilute solution titrated with small reagent volumes, changes in density and ionic strength can be ignored. If you ask the chemist to provide the initial sample volume, the reagent concentration, and the reagent volume added at each point on the titration curve, you can now convert his abscissa to an abscissa that has the desired flow ratio by use of the equation:

$$Ratio = \frac{Q_r}{Q_i} = \frac{V_r}{V_i} \cdot \frac{C_r'}{C_r} \tag{1}$$

where:

C_r = concentration of reagent used in control loop (in normality units)
C_r' = concentration of reagent used in titration (in normality units)
Q_r = reagent flow in control loop (gpm)
Q_i = influent flow in control loop (gpm)
V_r = reagent volume added for titration (ml)
V_i = influent volume used for titration (ml)

It would be nice if the chemist would use the same reagent concentration for titration that you use for control, but this is usually not the case. A more dilute reagent is used for titration to avoid working with hazardous solutions and very small incremental volumes. A small drop of a concentrated reagent can cause too large a pH change for steep titration curves and small sample volumes (Ref 1).

A pH titration curve that uses the above flow ratio for its abscissa can provide important information for the specification of a reagent control valve (Fig. 6-1). If a distance 'A' is marked

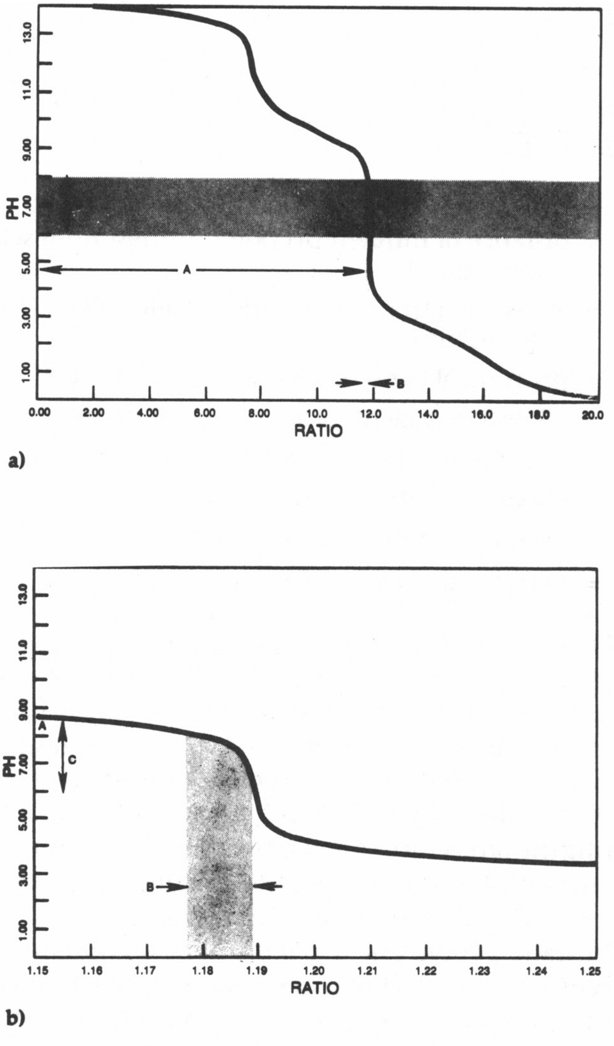

Figure 6-1. Part a) shows a general, wide-range weak-acid/weak-base pH curve as a function of reagent/influent flow ratios. Distance A indicates maximum reagent flow required with respect to an allowable pH error (pH control band) shown as band C. Band B indicates the allowable flow ratio as determined by the allowable pH error. Part b) shows how the curve reduces to a single 's' shape with a seemingly longer allowable flow ratio (band B) by expanding the ratio scale by a factor of 20.

along the abscissa from the influent pH to the center of the pH control band and the pH control band 'C' is translated to a band 'B' of allowable error in the flow ratio as shown in Fig. 6-1a, then the reagent valve size, allowable valve stick-slip, and required valve rangeability can be estimated by equations (2), (3), and (4). You can appreciate the importance of this information if you ever had a reagent valve that was the wrong size or a pH loop that oscillated even when there were no influent disturbances. If you don't like to use equations, no matter how useful, program them in your personal computer (Ref 2).

$$Q_{rmax} = 1.25 A Q_{imax} \tag{2}$$

$$E_r = \frac{80B}{A} \tag{3}$$

$$R_r = \frac{1.25A}{B} \cdot \frac{Q_{imax}}{Q_{imin}} \tag{4}$$

where:

A = distance of influent pH point furthest from set point (gpm/gpm)
B = allowable error in flow ratio for allowable pH error (gpm/gpm)
C = pH control band or allowable pH error (pH)
E_r = allowable reagent error or valve stick-slip (% of stroke)
Q_{rmax} = maximum reagent flow for 25% excess capacity (gpm)
Q_{imax} = maximum influent flow (gpm)
Q_{imin} = minimum influent flow (gpm)
R_r = maximum reagent valve rangeability

If the allowable reagent error is too small or the required valve rangeability too large, more than one stage of pH control may be needed. Such a solution is symptomatic of a difficult pH system typically caused by a very steep titration curve near the set point (B << A). Thus, titration curves can indicate how much mixing equipment and what instruments are required and should be obtained during the project scope and estimation stage. I usually don't see them until it's too late and project managers seem to balk at adding another neutralization tank during start-up.

The pH titration curve can also tell you how accurate your pH measurement has to be. pH measurement accuracy is affected by high and low pH excursions from sodium ion error and acid error. Steep curves can cause such excursions. On the other hand, flat curves mean that a given pH measurement error (0.1 pH or more for industrial applications) translates to a larger flow ratio error. Thus, for flat curves the measurement error can approach and even exceed the control error.

The pH titration curve can also help you tune the controller and even linearize the system. The steady-state open-loop gain can be estimated from the capacity of the control valve, slope of the titration curve at the set point, and the pH measurement span. If the titration curve is programmed into a microprocessor-based controller, the use of the flow ratio instead of pH as the controlled variable will linearize the system (Refs 1,2).

If a pH titration curve is generated by titrating the worst-case spill from your plant into a sample from the city or state sewer, a very simple program can be written to show the dilution and attenuation effects of sewer-system volumes on a pH upset from your plant. In many cases, the large volume dilutes and attenuates pH pulses so much that the actual pH change at the final stage of city or state treatment is negligible. The results of the program can be used to relax city and state restrictions on pH upsets to more reasonable values (Ref 2).

EXPERIENCES TO BE AVOIDED

All of these benefits can be realized, if you carefully detail your requirements for the chemist. Table 6-1 serves as a good checklist (Ref 1). It summarizes the many problems associated with the trick-or-treat syndrome. Also, remembering the following four case histories might help avoid some potentially bizarre situations that can develop.

Case Study One. Process engineers and chemists at a Canadian plant diligently collected samples for several weeks and generated twenty titration curves. It must have been the McKenzie brothers, because neither the ordinate nor the abscissa was labeled. I didn't have the heart to tell them the curves were worthless. I stared at them for several hours to come up with some use for them. No matter how hard I tried, I couldn't think of anything except a pop art exhibit. The dilemma was finally resolved when my company started a housekeeping program and requested that everyone pitch useless material. This allowed me to hurry the curves on to their rightful destination with a clear conscience that I was doing what was best for the company.

Case Study Two. On the next job, I got diverse opinions from various engineers as to how much trouble we could expect from a pH system. It turned out the engineers who thought the system was easy had the partial curve shown in Fig. 6-1b, whereas the others had the full curve in Fig. 6-1a. The actual curves had the abscissa labeled so that you couldn't tell one was a blowup of the other. Imagine the interesting control valves that would be selected based on Fig. 6-1b for a control band between 6 and 8 pH. Also, since I was going to linearize the system, some important wiggles would have been missing from the curve used for signal characterization.

Table 6-1 Common Problems with Laboratory Titration Carves

1. An insufficient number of data points generated near the equivalence point (steepest portion of curve).
2. The starting pH (influent pH) data point was not plotted.
3. The sample or reagent solids-dissolution time effect on the abscissa was omitted.
4. The gaseous-reagent-dissolution time and escape effect on the abscissa was omitted.
5. The sample volume was not specified.
6. Reagent concentration was not specified.
7. The sample temperature during titration is different from the process temperature.
8. The influent sample was contaminated by absorption of carbon dioxide from the air (nearly pure water sample).
9. The influent sample was contaminated by absorption of ions from the glass beaker (nearly pure water sample).
10. The laboratory measurement electrode type was different from the field electrode type.

11. A composite sample instead of individual samples was titrated.
12. The reagent used for the titration curve was not the same type as that used by the control system.

Nothing can match the graphical deception of a titration curve for a strong acid and base system. Have you ever heard someone say the set point is on the linear part of the curve? There are no linear parts.

This misconception is the basis for the nonlinear or notch-gain controller that approximates a titration curve with three linear and symmetrical segments. The slope of a titration curve constantly changes with pH. The steep part of the titration curve shown in Fig 6-2a, when magnified 10,000 times, reveals another S-shaped curve shown in Fig 6-2b. Similarly, a blowup of the vertical portion in Fig 6-2b would yield another S-shaped curve. For a strong acid and base, the slope changes by a factor of 10 for every pH-unit deviation from 7 pH.

Case Study Three. I once had the misfortune of being given a curve very much like that in Fig 6-2b for a pilot plant. One day, in a meeting to discuss the control system, it dawned on me that the 8.5 in.-wide paper should be extended about 7000 ft to show how far the influent pH really was from set point.

Case Study Four. When I was given a titration curve that involved the titration of a strong acid into a lime solution, I asked if the curve included the lime dissolution effect. The chemist said there was no time delay. When he added the strong acid, the pH dropped immediately. I was puzzled until I realized he was not waiting for the rebound effect. Sure, the pH plunges instantly, but then gradually rises as the lime dissolves and reacts. I asked for a plot of this, and got the time response shown in Fig 6-3a. It took nearly 20 minutes for the pH to reach its equilibrium value. You can imagine how this overshoot and subsequent delay deteriorates pH control loop performance. Whenever I see a pH loop with lime, I try to find a different application to work on. If the titration curve is run slow enough, the lime dissolution effect will be included (Fig 6-3b). A fast titration of an acid or base into a lime solution will yield a shifted curve (Fig 6-3c). A fast titration of lime solution

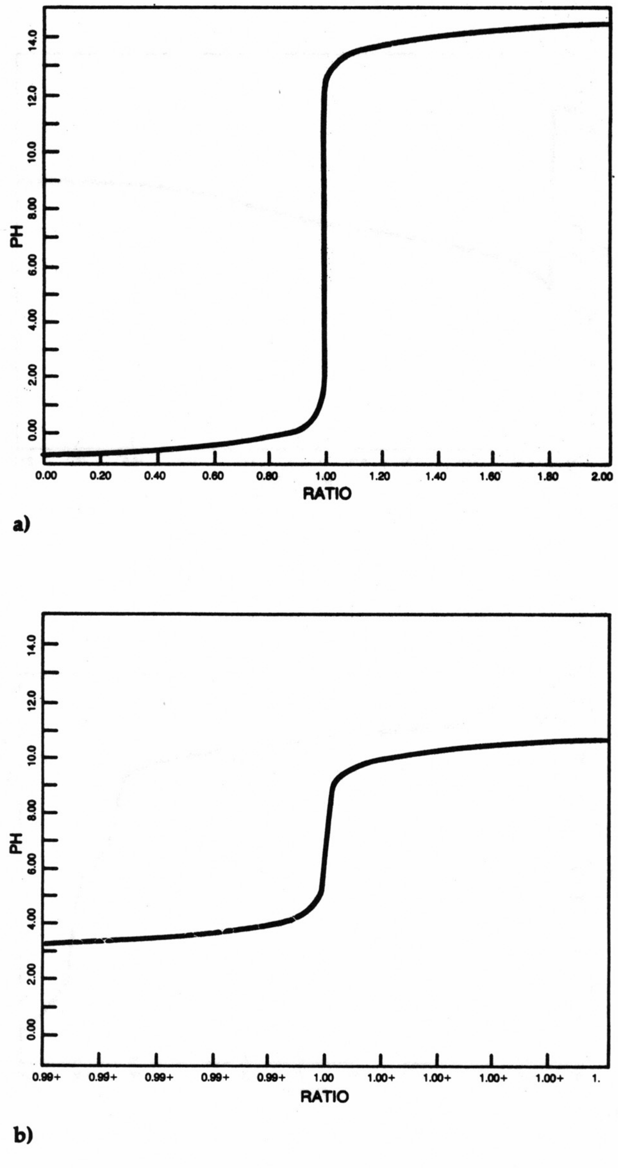

Figure 6-2. Part a shows a general, wide-range strong-acid/strong-base pH curve; Part b is a 10,000 expansion of the Ratio ordinate and graphically indicates the inherent nonlinearity of the "steep" portion of the curve brought about by pH numbers actually being exponents.

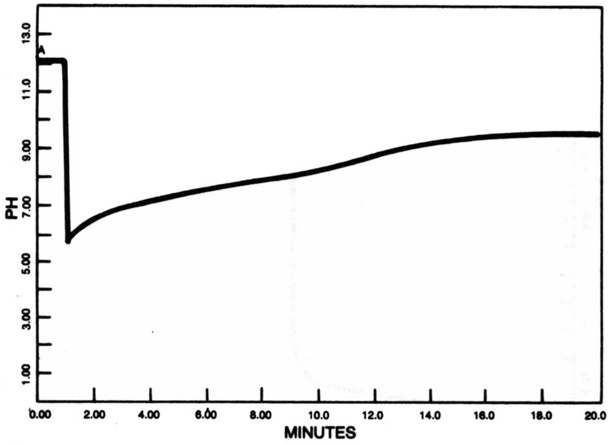

a)

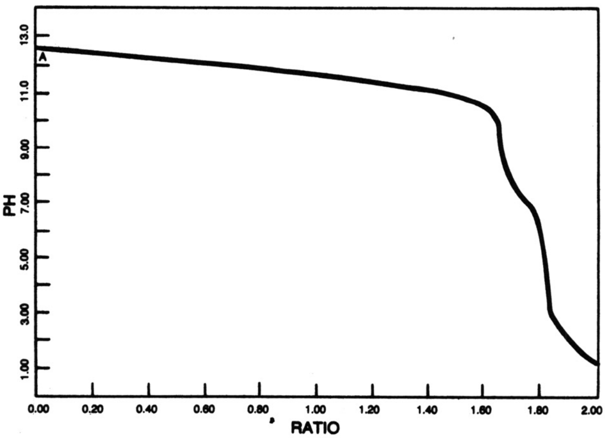

b)

Figure 6-3. Part a shows the lime dissolution time effect on pH. Part b shows a ratio curve developed by a fast titration of a lime solution.

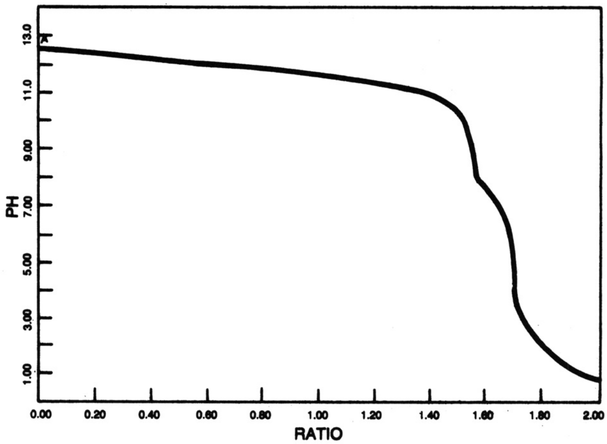

Figure 6-4. Figure 6-4 shows the ratio shift that occurs by performing a slow titration.

into an acidic or basic sample may yield no curve at all or, at best, a distorted one.

Based on these four case histories, you might think it is best to avoid the whole experience and pretend that titration curves don't exist. This is like turning out the lights and pretending you are not home on Halloween. Whether you answer the door or not, you are still vulnerable to some nasty tricks. Your process design engineer may have some curves hidden in the desk that were the basis of the strange reagent flows shown on the process flow diagram. Your best bet is to try and get the titration curves to work for you rather than against you. It requires attention to details, but isn't that what instrument engineering is all about anyway?

REFERENCE

McMillan, G K, and Cameron, R.A.,*Advanced ph Measurement and Control*, 3rd Edition, ISA (Research Triangle Park NC), 2005.

Wally and the Beave Automate Reactor Start-ups

The start-up of gas feeds to catalytic reactors is often the most labor intensive and hazardous part of the reactor operation. When these reactors are part of a large continuous chemical processing unit, delays on start-up result in costly downtime. Microprocessor-based controllers of distributed control systems can duplicate the best of the operator actions required for fast and safe start-ups.

The trials and tribulations of four engineers (Eddie—the project manager, Lumpy—the mechanical engineer, Wally—the experienced process control engineer, and Beave—the new engineer) show that successful implementation depends not only upon the engineer's technical expertise but also on the quality of the start-up support As problems develop, the engineer must have an open ear and mind to opinions from the operators and be flexible enough to make changes on the run. While the facts are true, any similarity between the characters portrayed and real people is purely coincidental.

Somewhere in the skies over North America, four people are in route to the first of two reactor start-ups. The people are immediately recognizable as engineers by the other business travelers. Maybe it is because they are wearing khakies and hardhats, their clothes are *wrinkled and their shirtails are hanging out, and they are waving calculators in the air, but more probably, it is because they are wont to quantify everything. If the plane would suddenly go into a nose dive, they would be estimating its terminal velocity instead of trying to make peace with their creator.*

Beave: Gosh, Wally, why do we have to go on this start-up? The simulations you did showed the stuff would work.

Wally: Gee, Beave, have you been taking stupid pills again? Don't you know that if anything can go wrong on a start-up, it will? There is plenty that can go wrong with this system. The new control system must automatically pressurize and establish normal gas feed to the reactor within 30 sec with a 30-in. control valve. The reactor operating pressure is within 3 psi of the fatigue point of the new 30-in. rupture discs. If the reactor upstream valve is too far open or the downstream valve is too far closed, the pressure can rise at the rate of several psi per second. The leakage flow through the new upstream valve used for pressure control is enough to cause the pressure to rise at the rate of a half a psi per second. Presently, the operators use nitrogen hoses and regulators to slowly pressurize the system.

Beave: How come you couldn't use the existing eccentric disc that had tight shutoff for pressure control?

Wally: Don't be a goof, Beave. The high breakaway torque of that valve results in a 15 to 20% overshoot upon breakaway from the seat. An overshoot of just a few percent is enough to overpressure the reactors and blow a rupture disc. It takes a crane half a day to replace a disc. If the system starts blowing discs, the plant people may not yell at you, but they will look at you kinda funny.

Eddie: Hey, kid, if the system doesn't work right the first time, I am going to get you transferred to the Alaska plant.

Beave: Wally, I don't want to live in Alaska. All they have to talk to there is elk.

Wally: Hey, look, Beave, Eddie was just giving you the business. I told the plant it would take five or six start-ups to get the controllers tuned and all the parameters for the set point ramps and decoupler to be adjusted. The plant won't install the new rupture discs with the lower burst pressures until the system is proven.

Beave: How come they want new rupture discs?

Wally: They want to operate the reactor at concentrations closer to the explosive limit for greater yield. The new rupture discs are designed to blow before the reactor does.

Beave: Would you go over the control diagram with me so I don't act like a goof tomorrow? (See Figure 7-1.)

Wally: Sure, Beave. We don't want the plant to think you are crummy or something. The strategy uses several different ramp speeds for the set points of the pressure and gas feed controllers. A fast ramp is used at the beginning to reduce the time required for start-up. A slow ramp is used at the end to reduce overshoot caused by too large an inlet valve opening (manipulated by the pressure controller) and undershoot caused by too large an outlet valve opening (manipulated by the flow controller). A portion of the output of the flow controller is added to the pressure controller output to help decouple the two loops. If the pressure is too high and is increasing too fast, the pressure control valve is closed in proportion to the rate of pressure rise. If the pressure still doesn't decrease, a small vent valve will open at the inlet to the reactor. Because the pressure loop has such a high gain and fast response and the flow loop has a low flow shutdown point within 5% of the normal operating point, fast and accurate instruments are needed to pull this whole thing off. The digital controller sample time, transmitter response time, and control valve stroking time must all be smaller than normal. The pressure control valve must throttle

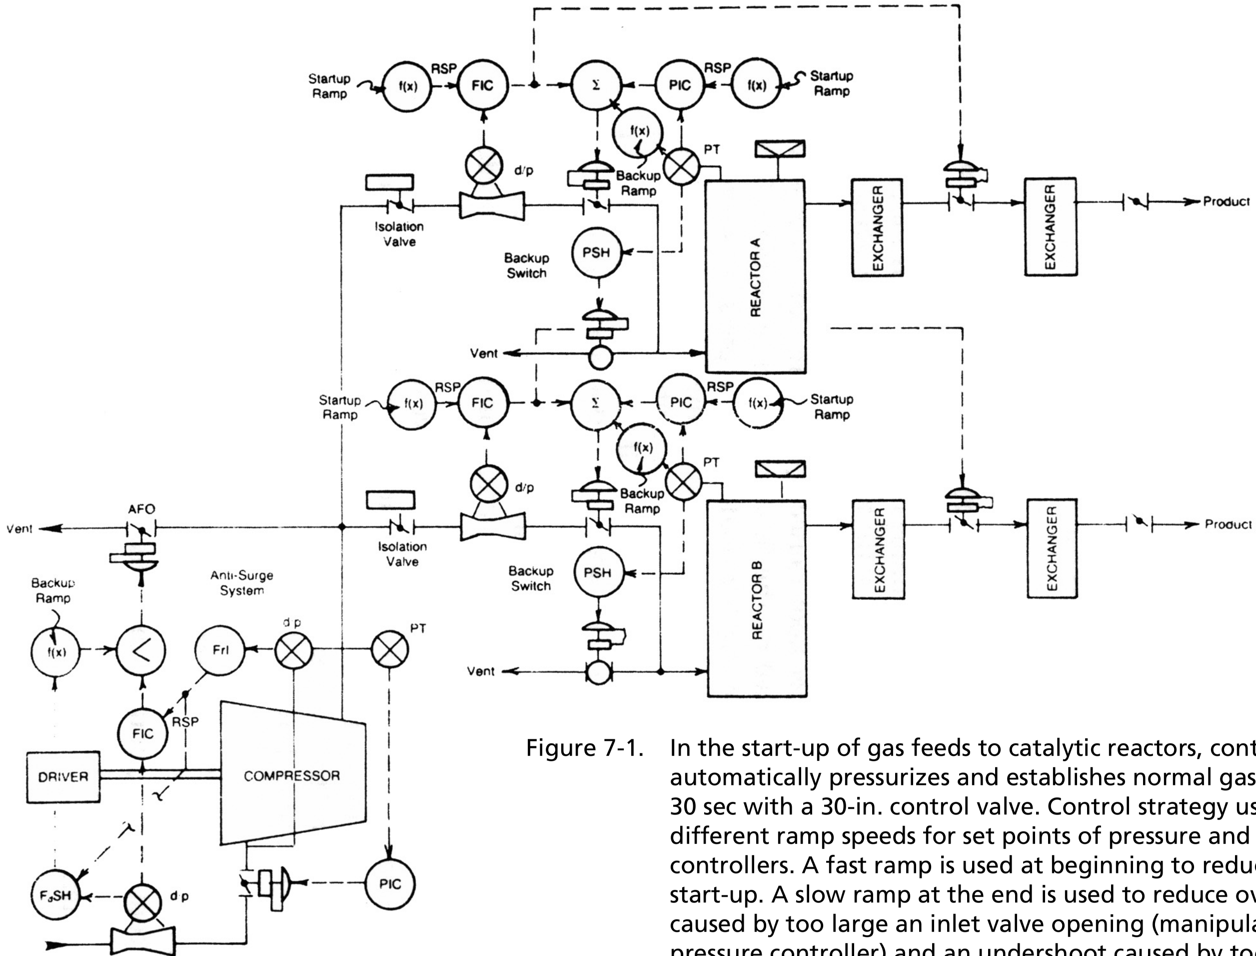

Figure 7-1. In the start-up of gas feeds to catalytic reactors, control system automatically pressurizes and establishes normal gas feed within 30 sec with a 30-in. control valve. Control strategy uses several different ramp speeds for set points of pressure and gas-feed controllers. A fast ramp is used at beginning to reduce time for start-up. A slow ramp at the end is used to reduce overshoot caused by too large an inlet valve opening (manipulated by the pressure controller) and an undershoot caused by too large an outlet opening (manipulated by the flow controller).

full scale in less than three seconds with less than one percent overshoot. This is no easy trick for a 30-in, butterfly. It is a good thing we throttle tested the valve at the factory.

Beave: It was real neat how the replacement of the large spring and diaphragm actuator with a volume piston actuator for the pressure control valve got rid of those disc oscillations caused by the quick release valve the vendor added to make the diaphragm actuator faster.

Wally: Heck yea, Beave. Those oscillations would have been enough to blow the rupture discs. We might both have ended up in Alaska.

The next morning Wally and the Beave enter a ballroom-sized control room full of consoles with colorful displays and flashing alarms.

Beave: Golly, Wally. What a setup. Those consoles must be just about the smartest machines there are. Can we get them to make something really neat like a Snoopy picture?

Wally: What is impressive is the computing power and flexibility of controller cards mounted in the panel upstairs. I am counting on them for my ticket home.

Lumpy and Eddie Arrive.

Lumpy: Great, I always wanted a ride in a starship. When do we go into warp drive?

Eddie: Wally, are you guys finished yet?

Walley: Aw, cut it out, Eddie. Give us a break. Ever since the project started you have been asking me that.

Beave: Wally, how come Eddie never asks us how the system works?

Wally: Beave, he has too much worrying junk on his mind, like schedules and budgets.

Beave: Hey, Wally, the field operator says the existing flow control valve sticks or something.

Wally: Gee whiz. That could mess up our start-up ramp. Let's go watch it stroke.

Wally and the Beave go outside to the flow controls valve.

Beave: How come none of the engineers are out here?

Wally: It is not air conditioned out here. Keep quiet and watch the valve, Beave.

Beave: What's going on Wally? My books said valves start to open at zero percent signal, but this valve didn't do a thing until the signal increased to 15%, then it jumped.

Wally: Beave, what did I tell you before about eccentric disc valves?

Beave: Oh yeah, I forgot. How come engineers use them?

Wally: It is pretty goofy for throttling service. I would only use them in on-off applications or if I never needed to throttle the valve below 20%. Let's go inside and modify the program to give a 15% initial signal to this valve for start-up.

Wally and Beave are in the room with the engineer's console on which they have adjusted parameters after each test start-up. Except for the first test, the overshoot is small enough to avoid the actuation of the rate-of-rise system. However, the undershoot after the overshoot lasts too long despite the addition of more and more reset action. It takes too long for the pressure to recover from the undershoot and reach set *point. Also, the operator can close the flow control valve so fast that it causes the pressure to reach the disc setting within two seconds. A nice finishing touch is the yo-yo improvisation by the rate-of-rise system.*

Wally: Beave, we better go in the control room and tell the supervisor we need to do some thinking, program modification, and more tests.

Beave: Wally, if we go in there now, we'll be killed. If you are going to get killed, it is better to get killed later. Let's go hide.

Wally: They must not be too mad, they aren't hollering or anything yet. Besides, I have got some ideas. I think the decoupler is counteracting the reset action. When the pressure control valve opens more, the pressure increases, which pushes more stuff through the flow control valve. The flow controller cuts back its output, which in turn decreases the pressure control valve position through the decoupler, negating the increase from reset action. You could increase the reset setting till you are blue in the face and still have an offset.

Beave: Actually my face is a little green.

Wally: If we reduce the decoupler gain, put the flow controller in override and ramp its output instead of its set point, and slow down the flow valve ramp during the overshoot, we can get to set point faster. If we add a velocity limiter in the program for a decrease in flow valve position, we can slow down the pressure rise when the operator accidently closes the valve. If we replace the rate of rise system with an override system that decrements the pressure valve position until the pressure drops below the system actuation point, we can stop the yo-yo.

Beave: Wally, how come you know stuff like that?

Wally: By studying variable speed recordings of valve positions, measurements, and set points, discussing the problem with the operators, and using some simple cause and effect analysis, you can learn enough to solve most problems.

Beave: Boy, it sure was exciting watching the numbers flash on the console. With its minimum update time of five seconds, I couldn't tell where the pressure was or where it was going. It was more excitement than a new engineer should be allowed to have.

After a long night, and many program changes, the start-up is fast and smooth and the system can stop the pressure from getting too high after the closure of the flow control valve. A bunch of weary engineers head back to the motel.

Beave: Gee, Wally, are we going to have to go through this sort of thing again? I almost had a mental meltdown (see Figure 7-2).

Wally: Don't worry Beave, I will help you through all the bad stuff. All the trouble you are getting in, I have already gotten in. You can't be a good process control engineer unless you go on start-ups, find out what your mistakes are, and correct them. Our next start-up is going to be another reactor system, but this time we are going to automate and coordinate the opening of the suction valve and closing of the surge valve for a centrifugal compressor on the feed gas. Since the override functional block showed how it enabled the controllers to make such a smooth transition between feedback and preprogrammed open-loop control, I am going to use it on the next start-up for backup protection against compressor surge and overload. Also, we are going to reduce the energy consumption of the compressor by slowly adjusting the suction valve to keep the reactor feed valves as far open as possible. Presently, the operators surge the compressor every time they start up, and they run with the reactor feed valves near the closed position so that they can handle large increases in the reactor pressure without running out of valve. While this system response is not quite so fast as the last one, it is continually being perturbed by the operation of blowback

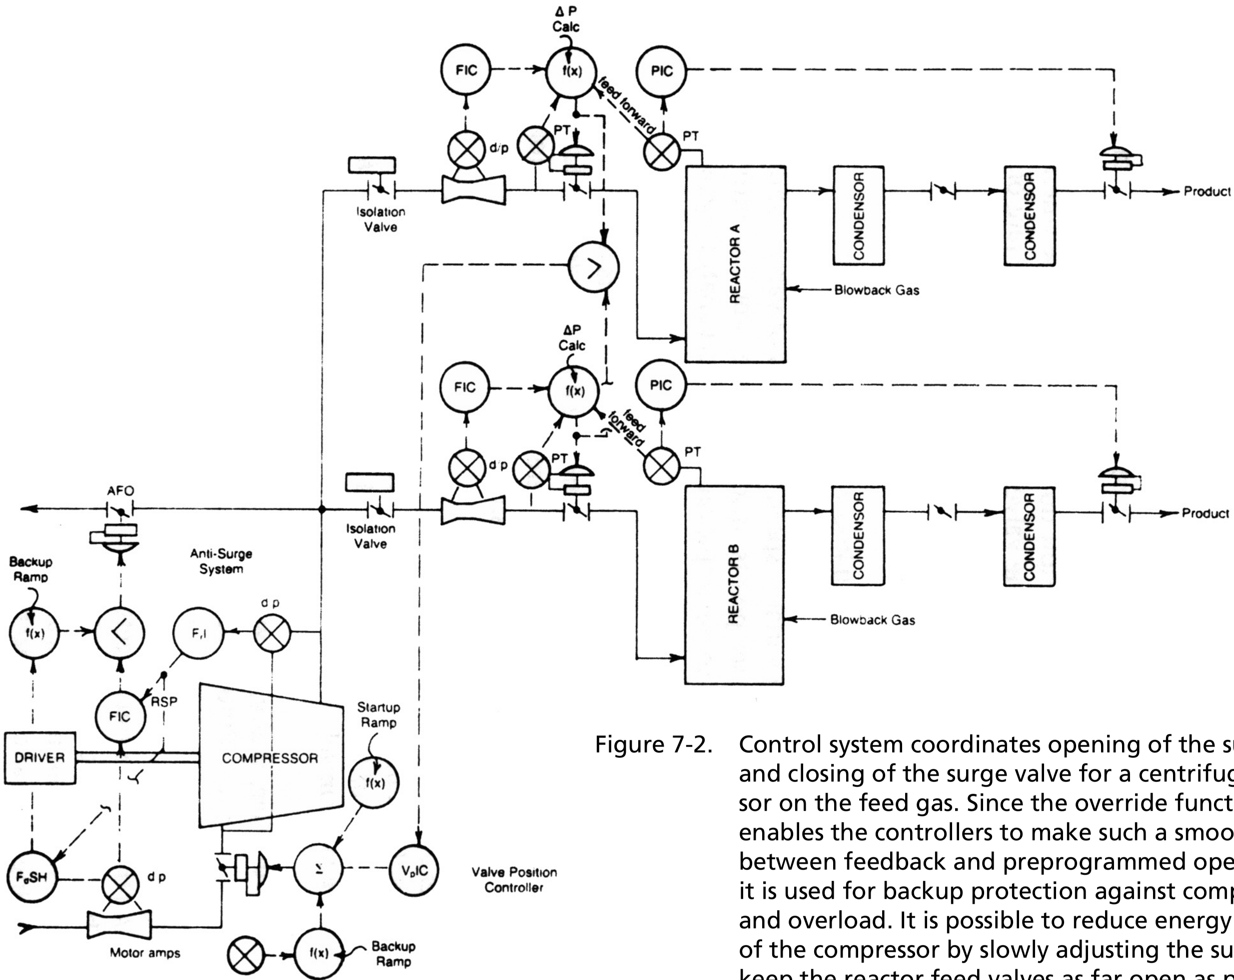

Figure 7-2. Control system coordinates opening of the suction valve and closing of the surge valve for a centrifugal compressor on the feed gas. Since the override functional block enables the controllers to make such a smooth transition between feedback and preprogrammed open-loop control, it is used for backup protection against compressor surge and overload. It is possible to reduce energy consumption of the compressor by slowly adjusting the suction valve to keep the reactor feed valves as far open as possible.

filters for the reactor. The system setup is similar except the flow controller throttles an inlet valve and the pressure controller throttles an outlet valve of the reactor.

Beave: Will they have a neat control room too?

Wally: No, they have a small control room crowded with a mixture of analog controls of various vintages. Instrumentation, unlike fine wine, does not improve with age. Asking an instrument engineer to interface with that maze of wires and junk borders on mental cruelty. Fortunately, some other engineers are going to take care of getting us the signals, so we can concentrate on the control strategy. We will be commissioning a stand-alone microprocessor-based controller similar to the ones we have just been working with. It just won't be connected to any fancy displays until they get a distributed control system.

Beave: How come they don't have one yet?

Wally: It takes a lot of paperwork to get one approved. One last thing, Beave. If you say 'how come' once more, I am going to be sick.

A few weeks later, Eddie, Lumpy, Wally, and Beave squeeze into the old control room of their next start-up. Lumpy scavenges for doughnuts while Eddie introduces the crew to the operating supervisor.

Eddie: Sir, what a well-maintained system you have.

Sir: Why, thank you, Eddie, but we would like a distributed control system.

Lumpy: (aside to Wally) What a mess. There is nothing here a few sticks of dynamite might not fix. That will give them their distributed control system.

Wally: Cut it out, Lumpy.

Beave: Hey, Wally, the operators say that they have to wait about 10 sec for the compressor to accelerate before they load up the compressor, otherwise the high current trip will shut them down. Then they open the suction valve as fast as they can but the compressor always surges. While running, they adjust the suction valve based on amps. Too high an amp reading indicates an impending overload and too low an amp reading indicates an impending surge. Also, they say the buildup of pressure in the reactors from plugged filters and blowback plays havoc with their feed-gas flow control.

Wally: Gee whiz, Beave. We had better make some program changes right away to duplicate their actions but faster and more accurately. We can use ramps and overrides to keep the compressor out of trouble. For flow control, we can compute the changes in the pressure drop at the feed valves caused by the cycling reactor pressure and automatically correct the valve position before they cause a big deviation of flow from set point. This feedforward action will only be enabled after the reactor is up to set point so that it won't overcorrect during pressurization.

On the very first attempt, the compressor starts up smoothly without surge on automatic. After adjustment of the gain applied to the feedforward signal, the flow fluctuations from pressure cycling are reduced by an order of magnitude. The only remaining problem is the eccentric-disk surge valve. Once it closes after start-up, it takes two minutes for it to break loose. Since a surge cycle period is about one second, the compressor would surge over a hundred times before it opened. To prevent Ms, a low output limit of 5% is used to keep the valve from closing until the valve can be replaced with a low leakage conventional disk butterfly valve. Weary but happy, Wally and Beave head home.

Beave: Wally, my mind is as soggy as a bowl of Farina.

Wally: Yeah, Beave. But look at what we learned about automating reactor start-ups.

Beave: Wally, do you think that what we ended up with is an expert system?

Wally: No, Beave. The program language wasn't fancy enough, and engineers don't think of operators as experts.

Beave: They are, kind of, if you listen carefully.

Wally: You aren't such a goof after all, Beave.

Wally and the Beave Return to Automate Another Reactor Startup

Years ago, we sneaked along on a trip with Wally (the seasoned instrument engineer), Lumpy (the enduring mechanical engineer) and Beave (the new engineer) to automate the start-up of large continuous catalytic reactors. Through some strange twist of metaphysical fate, the team is reunited to tackle the start-up of a waste treatment reactor to destroy some very nasty stuff. We rejoin the entourage as they board the plane that will whisk them away to another exotic place.

Beave: Wow, Wally, this is really neat getting these frequent flyer bonuses.

Wally: Don't be such a goof, Beave. What do you want with another trip when you travel too much?

Beave: But Wally, we can bring our friends along. You could bring that girl that you met.

Wally: I think the company has a rule against having that much fun.

Lumpy: It would be great to have a friend along to watch the Beave get clobbered. Maybe they could take pictures or something?

Wally: Cut it out, Lumpy. How come you are always giving Beave the business?

Lumpy: Well, I don't want him to go out in the world and get slaughtered.

Wally: I don't know how to tell you this, Beave, but Lumpy could be right. The plant has a waste pond that keeps blowing up. We are going to put the stuff in a large tank with a removable lid. If we control the pH real tight, dilute it enough, and keep the reactant-to-waste flow ratio just right, the nasty stuff will be destroyed.

Beave: Golly, Wally, what if we don't?

Lumpy: We will launch the biggest flying saucer since War of the Worlds.

Wally: Pipe down, will ya, Lumpy. Beave, we have added a hundred interlocks and alarms to keep the operation safe.

Beave: What a relief.

Wally: Not really, Beave. The plant already has so many interlocks and alarms that it is difficult to get the plant to run right, and bad batches and trips cause the biggest upsets to waste treatment. If we screw up, the distractions and problems we create will make it difficult for the operator to focus on the process. It will be like standing at the bottom of a ski slope and watching a snowball at the top rolling towards you change into a boulder and then into an avalanche.

Lumpy: You ought to see how the operators tap out the Morse code for help on the alarm acknowledgement button. They know it so well they don't even have to look at the screen to do it.

Beave: Figuring out interlocks makes my head hurt. Will I get a standing around part?

Wally: No, Beave, you have to make sure the interlocks work right.

Beave: If I yell and holler and everything, will they let me go home?

Wally: Listen, Beave, the real-time simulation shows that we can automatically restart smoothly and safely after an interlock trip without disturbing the operator. But we have to stick around and make sure it can handle all the strange things caused by real instruments and operating conditions. Otherwise, we could do more harm than good. A distributed control system (DCS) has so much power it is like a Corvette idling. You can stomp on the throttle, but you better know where you are going and be prepared for the unexpected. We need to take a lot of test rides before we turn our control system loose. Besides, Beave, you are not a real engineer until you survive a pH start-up.

Beave: Gee, Wally, what is the control system?

Wally: I am glad you asked, Beave, otherwise the readers wouldn't know what we are talking about. The heart of the control system is a lead-lag or cross-limit strategy similar to what is used for air-to-fuel ratio control for boilers but extended to include a third flow controller. For a load increase or start-up condition, the water flow leads the reactant flow, which leads the waste flow. For a load decrease or shut-down condition, the water flow lags the reactant flow, which lags the waste flow. This ensures safe operation for upsets, start-up, and shutdown by over-dilution and excess reactant for the feed of the waste from the first tank to the next tank in the reaction. The demand for waste flow is set by a level controller on the first tank. Interlocks bottle up the first tank but an automatic restart of feeds occurs when conditions clear. The DCS controllers are coordinated with the programmable logic controller (PLC) interlocks by use of the output track function. The remote set point of the level controller ramps down from the level during shutdown to the minimum set point limit to minimize inventory to facilitate maximum upset capability. The pH control system is pretty much standard stuff except for a middle select of triple redundant pH measurements and separate controller for the small and large reagent valve for improved operator interface

. . . A week of water batching and low load operation has passed. We catch up with Wally, the Beave, and Lumpy in the field . . .

Lumpy: A guy could get tired of watching water run.

Wally: I guess you are not going to be a plumber.

Lumpy: But a plumber gets good money to watch water run.

Wally: We have learned a lot of neat junk and made a lot of adjustments critical to the success of the system. We have reduced trim sizes to improve turndown, raised the level controller low output limit to prevent excessive turndown, lowered the level controller high output limit to match maximum waste valve capacity, raised trigger points for the flow measurement dropout to insure cleared conditions for an automatic restart, changed the velocity limit location and valve for analysis signals to hasten reestablishment of correct signals after a download or calibration, added a level measurement filter so that level noise doesn't upset the ratio control system, speeded up the level set point ramp rate for low loads, and reduced the magmeter calibration range for the reactant. Oh, I almost forgot the time delay we added for restart. Beave, remember when you tested the interlock to trip the feeds?

Beave: Sure, Wally, the valves kept closing and reopening faster than the update time of the DCS screen. It was really weird. I couldn't figure out what was going on.

Wally: Without a time delay, the system was attempting a restart every few seconds but failing due to the forced trip condition. You must have tested the trip and restart about a hundred times. A forty second time delay before restart solved the problem.

. . . The pH system is commissioned for low loads. The pH control is so tight, the middle measurement stays within a hundredth of a pH of set point until the Beave gets a hold on it . . .

Wally: Beave, what did you do to my pH control system?

Beave: Gee, Wally, the chemist thought the measurements looked frozen, so I recalibrated them all and forgot to put the pH controller in manual. They sure weren't stuck. They all zoomed offscale high when the big reagent valve popped open. Are you sore at me, Wally?

Wally: No, but I hope you have learned not to mess with something that works. Chemists and process engineers always think instruments are wrong, even when things are going great.

Beave: Wally, what do we do now?

Wally: We wait for ten days to see what happens.

Beave: I wish I could go to sleep for ten days.

Wally: That is not what we are supposed to do.

Beave: I wish we were two men from Mars so we wouldn't know what we are supposed to do and could do what we want to do.

. . . The whole start-up goes better than anyone expected. Wally and the Beave are at the DCS console staring in amazement at the trend recordings . . .

Beave: The system runs and restarts so smoothly, the operators don't have to mess with it at all and can concentrate on making good product. How come it works so well?

Wally: Well, Beave, the pH control system has a really low dead-time-to-time constant ratio due to the small turnover time from good mixing and the small reagent delivery delay from side injection and a recirculation line. The lead-lag strategy prevents the ratio system from hitting the trip points and the automated restart is foolproof.

Beave: But, Wally, we had to tweak a lot of parameters and change a bunch of trim sizes, and the process engineer looked like he was going to holler at me.

Wally: When you get older, you tend to remember the good and forget the bad.

Beave: Does that mean I will think those days were not so bad?

Wally: Sure, Beave, you might even write a story about them.

Can You Say "Process Interlocks?"

Mr. Rogers, a public relations specialist, walks into the offices of the engineering department of a chemical company singing a sprightly tune. Mr. Rogers is going to visit Mr. Fellow, a process control engineer, to learn all about process interlocks.

Mr. Rogers: It's a beautiful day in this engineering department, a beautiful day for an engineer. Would you be mine, could you be mine? It's an engineering day in this beautiful department, an engineering day for a beauty. Would you be mine, could you be mine? I have always wanted an engineer just like you, so let's make the most of this beautiful day, since we are together we might as well say, would you be mine, could you be mine, won't you please, won't you please, please be my engineer?

Mr. Fellow: You must be from public relations.

Mr. Rogers: Yes, I bet you know all sorts of special things about process interlocks.

Mr. Fellow: Not much more than other control engineers, but I have had to organize my thoughts because in a weak moment I volunteered to be the chairman of a committee to develop a corporate guideline for process interlocks. I should have become suspicious when I found out the other chairmen have quit the company. Getting approval of a guideline in a large company is like taking several thousand kids to an ice cream store and getting them to agree on vanilla.

Mr. Rogers: I like tutti frutti. Everyone in the world is different from everyone else. We look different, we smell different, we sound different, and we have different human thoughts inside us.

Mr. Fellow: Some people thought the guideline was too long, some thought it was too short, and some thought it was boring. I made a video tape presentation of it. Now the people who think it is too long can play it on fast forward, the people who think it is too short can play it in slow motion, and the people who think it is boring can play it backwards. Through some freak break in the space-time continuum of the standards world, agreement was reached and remnants of my sanity were salvaged.

Mr. Rogers: This is so much fun I think I could stay here a week.

Mr. Fellow: Don't do that—you might have to fill out a time sheet or even a goals document. Why do you want to know this stuff?

Mr. Rogers: Because of public concerns and company image. Can you say "company image"? Sure, you can.

Mr. Fellow: Well, the guideline has four classes. It is 33% better than the previous draft because it had only three classes. Also, we are now as good as our largest competitors since they have four classes. There are some real benefits. The additional classes allow you to separate the economic issues from the safety issues and focus on those interlocks that require special attention. Prior to this, the protection of human life was grouped with major property protection and formed a very large class for which it was difficult to enforce stringent design and test requirements. Also, the cost of production losses, which can exceed property losses, after a failure was largely ignored or relegated to a class of so-called "operational" interlocks. Now everything that affects the bottom line of company profits is considered in the economic classes. The classes are as follows in order of decreasing severity:

- Class I - Community Protection (safety)
- Class II - Employee Protection (safety)
- Class III - Major Property and Production Protection (economic)
- Class IV - Minor Property and Production Protection (economic)

Mr. Rogers: Does that mean people outside the plant gate are more important than the people inside?

Mr. Fellow: No. It means there are more of them outside and they are less prepared. Also, there are children out there and people who act like children.

Mr. Rogers: Without children, I would be out of a job.

Mr. Fellow: Since you probably don't want to know the nitty gritty details of the design and test requirements for each class, I thought I would explain some of the ideas on which the guideline is based.

Mr. Rogers: It is such a good feeling to know you have ideas to talk about.

Mr. Fellow: To date, a shotgun approach has been used for the design of interlock systems.

Mr. Rogers: Does that mean you were shot if the interlock didn't work?

Mr. Fellow: No. It means that interlocks were scattered about without much focus or attention to the most direct and distinct causes of a hazardous release. Interlocks tend to proliferate like alarms to the point where they detract from the importance of the critical ones. I could issue a blank piece of paper with the title "Process Interlock Guideline" and a thousand interlocks would be added from people thinking about them in a random fashion. To avoid this, a logical cause and effect approach to interlock classification is needed. Fig 9-1 shows a generic sequence of events that lead up to a hazardous release. Most people think from left to right in this figure whereas they should work backwards from the release. The most immediate causes or, in other words, the direct causes of the hazardous release are events A, B, and C. However, events B and C are concurrent or simultaneous. Thus, the most direct and distinct causes are A and B or C. These direct and distinct causes should have the most severe classification. Events D, E, F, G, and H are

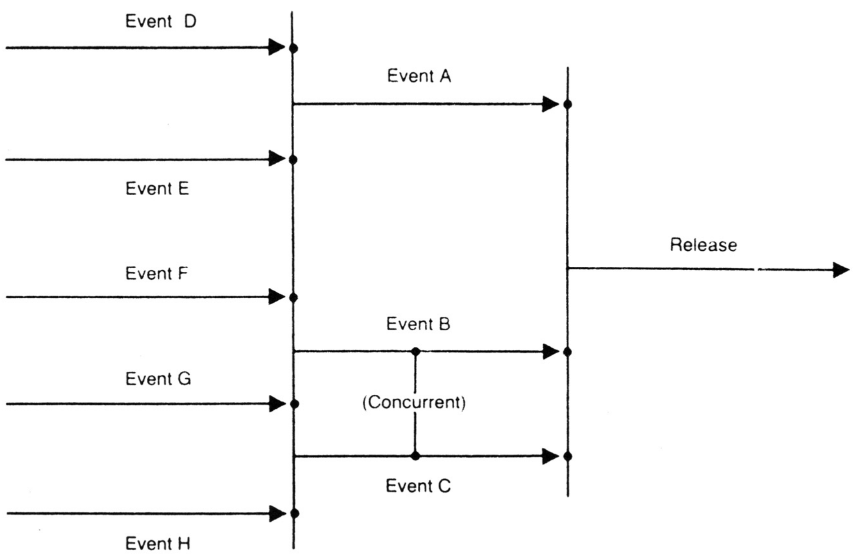

Figure 9-1. Generic sequence of events leading to a hazardous release.

indirect causes and have an interlock classification that is based on the undesirability (usually economic) of the intermediate operating condition. If events D, E, F, and G were interlocked but event H was neglected, a release could occur if the direct and distinct causes were not identified and interlocked. In fact, it is difficult to cover all the possible indirect causes. Also, it diverts attention and dilutes the effort. Hence, it is better to identify and classify the direct and distinct causes as class I and II interlocks and the indirect causes as class III and IV interlocks.

Mr. Rogers: I love generic things.

Mr. Fellow: But engineers don't, so I worked out some examples. Fig 9-2 shows a chlorine vaporizer. In the past, interlocks would have been placed on many but not all of the indirect causes such as the wide open nitrogen regulator, the closed downstream block valve, and wide open steam valve and regulator. The design and test requirements for the direct cause were the same as that for interlocks in general. Due to the large number of interlocks, insufficient time and money was available to assure the integrity of the pressure

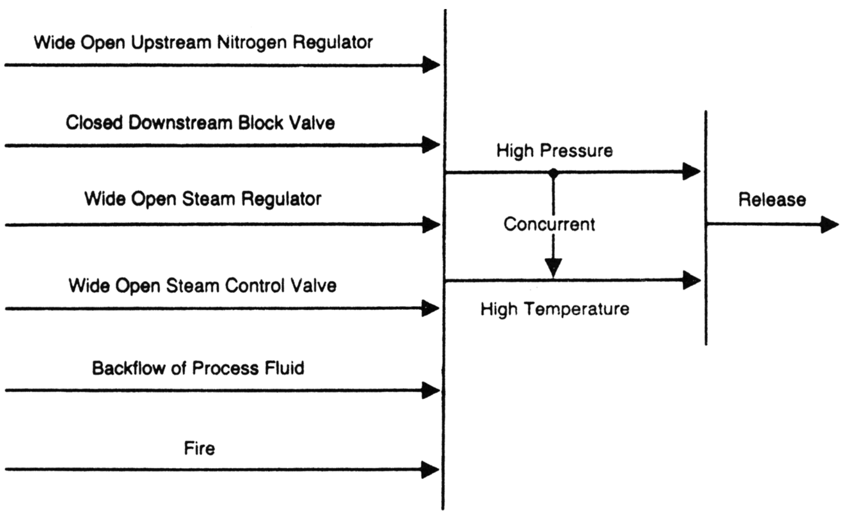

Figure 9-2. Chlorine vaporizer example.

interlock. The most direct and distinct cause of a potential release is high pressure. It is distinct because high temperature is always accompanied by high pressure but not vice versa for the boiling mixture. Thus, the high pressure interlock should receive the most attention and the most severe interlock classification. Fig 9-3 shows a chlorinator. Here a reaction is going on so that high temperature is not always accompanied by high pressure. Accelerated corrosion occurs with high temperature. Thus, the direct and distinct causes in this example are high pressure and temperature and their interlocks should be designated as Class I or II. As you might have guessed, in terms of the most frequent direct and distinct causes, pressure rates first, temperature second, and explosive mixtures third. Explosive mixtures pose an additional problem because an

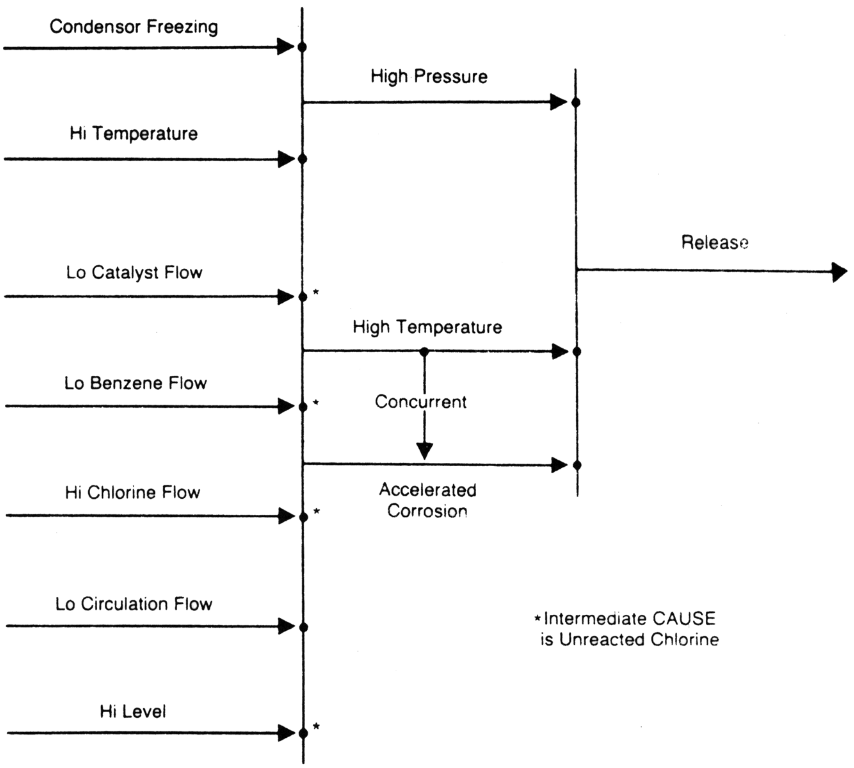

Figure 9-3. Sequence of events leading to a hazardous release (chlorinator).

analytical measurement of the direct and distinct cause is not used due to perceived reliability problems. If redundancy cannot alleviate these concerns, a material balance computation from flow measurements may adequately predict an out of range mixture.

Mr. Rogers: Sure. What about instruments and special things like that?

Mr. Fellow: Before you can evaluate the merits of various instrument systems, you need to realize that 80% of interlock failures are due to failures of field devices. Thus, only 20% of interlock failures are due to failures of control devices.

Mr. Rogers: Oh, it must not be very nice in the field.

Mr. Fellow: That's right. The field devices are exposed to harsh environmental and process contitions. They are not as thoroughly tested as control room devices because they are not as accessible and the condiditions are not as comfortable. Also, field devices have mechanical components. Statistics show that 45% of the interlock failures are caused by measurement device failures and 35% of interlock failures are caused by valve or valve actuator or accessory failures. If you consider that 75% of the control room software device failures are due to programming mistakes, that means only 5% of the interlock failures are caused by failures of software device electronics. The availability of electronic simulators has in some ways aggravated the differences in reliability. The substitution of a simulated signal for the real signal gives a false sense of security and only checks out the devices least likely to fail. I have seen many written descriptions of test procedures for interlock systems that detail a thorough test of the control room devices via simulated signals but omit any functional test of the field devices, particularly the sensors.

Mr. Rogers: All those statistics make my head dizzy. Let's pretend we are measurement devices. It helps to play about things. It helps you find out how it really feels. Oh, it must be a special feeling to touch the process.

Mr. Fellow: A lot of our chemicals would make it a very special feeling. The weakest link in an interlock system is the measurement device. The second-weakest link is the final element. My start-up experiences over the last few years have shown that the distributed control systems after the initial burn in are sitting there waiting for me to solve all the field device problems even though the complexity of the microprocessor-based devices is an order of magnitude or more greater than anything in the field. The new digital controllers have proven to be more reliable than even analog controllers. When you consider this plus the fact that these controllers have the computational capability to detect a field device failure and take corrective actions, you realize the proper use of distributed control systems can improve the system's reliability.

Mr. Rogers: That's a special thought.

Mr. Fellow: It is now common practice within my company for critical loops to monitor for upscale or downscale failure of a measurement. If such a failure is detected, the controller tracks the last valid output and the operator is then notified. He or she then has the option of leaving it as is with a frozen output or adjusting the output in manual until the instrument is repaired. On really critical loops, triple measurements and middle selection is used. By selecting the middle measurement, you can ride out any type of failure. Even if the measurement sticks at an intermediate value, the other two measurements will move away and one of these remaining valid signals will be selected. The middle selection also increases measurement accuracy by discarding signals that start to drift and decreases measurement noise without the introduction of an additional lag. Many signals start to become noisy before they fail. The detection of a final element failure is more difficult and not widely practiced. It requires the addition of a valve position transmitter to verify the valve stem actually moved or a flow measurement to confirm the flow actually changed when the controller output changed.

Mr. Rogers: You certainly are a man interested in different things.

Mr. Fellow: Now we are ready to discuss the relative amounts of various functional characteristics for the four major types of instrument methods used for protection against a hazardous release.

Mr. Rogers: Sure. I knew we would.

Mr. Fellow: Table 9-1 shows that the field switch that was traditionally thought to be the best for a critical interlock because of its simplicity has a low amount of ruggedness. It is inexpensive and you basically get what you pay for. It has low amounts of

Table 9-1
How Interlocking Systems Compare

	Relative Ability to Satisfy Functional Characteristic			
Functional Characteristic	**Field Switch**	**Relief Device**	**Analog Device**	**Software Device**
Ruggedness	low	high	medium	low
Complexity	low	low	medium	high
Accuracy	low	low	medium	high
Response speed	high	high	medium	medium
Predictability of failure mode	medium	medium	medium	low
Verifiability of integrity	low	low	medium	high
Self-test capability	none	none	low	high
Diagnostic capability	none	none	low	high
Security from jumpers & tools	low	medium	low	high
Security from changes	medium	high	medium	low

complexity and its response speed is high, which is good. But there is little verifiability of integrity, self-test capability, and diagnostic capability. The relief device has about the same characteristics but is generally more rugged. The analog device now provides a signal to monitor but lacks the computational capability to do much with it. The software or digital device can do most of your troubleshooting for you. The main limitations in a programmable device are the availability of signals and your imagination.

Mr. Rogers: Why do some people refuse to use software in interlock systems?

Mr. Fellow: They think it is unsafe because it has an unpredictable failure mode and is not secure from unauthorized changes. What these people don't realize is that they really don't know how their hardware devices will fail either. I have seen transmitters fail upscale, downscale, and midscale. If you consider that the cross checking and self-testing of software devices is a piece of cake and that these devices are in the protected environment of the control room, you realize the problem has been blown out of proportion.

Mr. Rogers: I like chocolate cake but it blows me out of proportion.

Mr. Fellow: There are a lot of misconceptions about the security problem. I have trouble figuring out how to make a program change on a new system even after a week or more of schooling. You have to wade your way through many menus and have to know the name of addresses of several devices before you can do some damage. Even if you have several programmers, they tend to specialize in different systems or portions of the same system. Compare this with the number of people in a chemical plant who know how to use a jumper or screwdriver. Many of the hardware devices have terminals and adjustments readily available to them. Also, a software device is smart enough to detect when someone is messing with it. Thus, in my mind, the field switch and analog device is actually less secure than the software device.

Mr. Rogers: Wow. Should software devices be used everywhere?

Mr. Fellow: We are not ready for that yet. The proper application of a software system requires considerably more expertise. There is presently a shortage of engineers proficient in the configuration of such systems. The most vulnerable point is the programming. Presently, there is no way of deterministically proving that a complex program will work for all conditions. There can be bugs that show up only when a strange signal or parameter is used. While this sounds somewhat hopeless, most of the risk can be alleviated by developing methods for software testing and validation before start-up and after every change. There are also some tricks that can be used like the use of diverse and separate programs. Just like diverse (different principles of measurement or methods of actuation) and separate (separate process and electrical connections) hardware can be used to reduce common mode failures in an instrument system, diverse and separate operating systems and calculations can be used to reduce common mode failures in software. It is important to remember that even if you have triple redundant software devices, if you have the same program mistake in all three, you are still in trouble.

Mr. Rogers: Thank you, Mr. Fellow. Do you make every day a special day?

Mr. Fellow: No. My company does that for me.

Mr. Rogers: Would you like to go to the land of make believe?

Mr. Fellow: No thanks, I have been to enough management meetings.


...OH YES, THOSE SWITCHES, TEE-HEE-HEE... I'M AFRAID WE GOT A LITTLE CARRIED AWAY WITH INTERLOCKS. ONE HAS TO ACTUALLY FLUSH THE TOILET TWICE BEFORE ONE CAN START UP ONE'S PLANT.
HIGH TECH TOILET
Williams

The Survival of an Instrument Engineer

We would like to say that a person becomes an instrument engineer after years of diligent study and academic preparation, but this is excerpted from a nonfiction book. No one plans to become an instrument engineer. The prospective engineer doesn't even know the profession exists until the job offer comes. There are few curricula in universities with even a remotely similar title, no mention of the profession by counselors, and no recognition of the profession in such technical societies as IEEE and AIChe. Also on top of this, public understanding of the profession is nonexistent. Have you ever tried to explain to a spouse or a friend what an instrument engineer does? Tell them you specify D/Ps, RTDs, PLCs, and PCVs. It leaves them speechless.

One usually becomes an instrument engineer by default or, in other words, because one has found nothing better to do. The lure of gainful employment has attracted chemical, mechanical, and electrical engineers to this technology. (There was once a civil engineer in instrumentation, but he kept trying to bury his installation mistakes in concrete. He now works for the nuclear industry).

If you are one of the chosen few who have a job offer in front of you right now and want to know if you've got what it takes to become an instrument engineer, you need to truthfully answer the following questions:

- Do you like to thumb through glossy pages in hundreds of catalogs?
- Do you like to keep track of intricate details of thousands of gadgets that change yearly?
- Are you capable of forgetting all the math you learned in college? (The few cases in which algebra might be needed, such as control valve and orifice sizing, were done first with special-purpose slide rules and then with personal computer programs).

There is some debate as to whether an instrument engineer even needs to be able to count! This does not imply that the village idiot can become an instrument engineer. You must, however, be able to learn on the job how to specify, install, check out, and start up complicated systems.

Instrument engineering is basically a lot of fun. It is like being paid to go shopping at your nearest electronics or computer store.

Americans love to go shopping. Just look at how crowded your favorite shopping mall is on weekends. As an instrument engineer, you get to buy thousands of gadgets with somebody else's money. You also get to watch these gadgets in action. The visual and audible feedback in a control room is impressive.

The Learning Experience

It is an undeniable truth that instrument engineers learn by making mistakes. In fact, to be a really good instrument engineer, you should leave a trail of disasters in your past as long and incomprehensible as the course of the Amazon. Your goal is to go from being a lowlife engineer on the early shift of winter start-ups to being a consultant after the golden handshake in a plush office complete with stereo and a wet bar. If you can't wait that long, you can acquire the same comforts by becoming either a manager or an instrument sales representative, or both.

Some chemical companies have recognized the need for on-the-job training (usually after a new engineer is found wandering aimlessly on a job site with his calculus book) and have tried to establish their own instrument engineering school or have brought in an outside consultant. The schools until recently either have taught process control from the professor's viewpoint (the same one who thinks the universe is described by partial differential equations) or have taught instrumentation from the illegible notes of comatose engineers who mumbled in monotone. What is unfortunate is that experienced engineers typically don't like to speak in public or write. Also, they are usually too busy to extrapolate beyond their experiences or to develop principles to guide their own future actions or those of others. The result is a shortage of good mentors and instructors. Professional societies, such as the Instrument

Society of America, have done an excellent job of seeking out and finding engineers able to communicate their experiences.

To be really effective, a school should give the student engineer a trial run at a career in instrumentation by condensing years of mistakes into one very intense experience, which should be as close to real life as possible. First, the student should be bombarded with thousands of facts on how instruments work until he or she suffers a mental meltdown. To insure plenty of mistakes will be made, no information should be given on how to determine what instrument is best for a given application. The student should be given a project where the instruments have already been purchased and the design drawings issued for a lump-sum bid. Verbal harassment by a real-life project manager about schedules and budgets should be used to add a greater touch of realism. The student should receive a set of process and instrument diagrams drawn with disappearing ink so that process revisions can be made with a clear conscience. The process should be revised hourly. Chemicals and/or concentrations should be chosen that do not appear in handbooks on physical properties or materials of construction. Electrical, mechanical, and chemical engineers should be given solids-handling projects, computer projects, and electrical distribution monitoring projects, respectively. For start-up, each engineer should be reassigned to a completely different project.

A tour of duty with a large engineering department is equivalent to attendance at the school described above. Such organizations create an environment for the maximization of mistakes. They like to play musical chairs with the assignment of engineers to projects so that there is no continuity or foundation of knowledge to build on. You will never be asked to work on a project that is technically related to anything you have worked on previously.

Of course, you must know when you have made mistakes and have a chance to correct them in order to learn from them. Therefore, it is not wise to go to work for a contract engineering firm until you have been exposed to a lot of mistakes or, in other words, until you are experienced. In such firms, you are either in a design or field section with no overlap in function. In the design section, you move from one project to another without having to try to make your designs work. In the field section you spend all your time fixing other people's mistakes, but don't get to make any design mistakes of your own.

A stint as a plant engineer is necessary to round out your perspective. Only by facing the day-to-day repair problems can you appreciate the impact of instrument selection, documentation, location, and installation on its serviceability. Many of the operational problems don't surface until after start-up. If all your field experience ends after commissioning, your viewpoint is distorted. You may be happy in your ignorance, but the plant engineer who has to live with your design will not be.

As a rule, plant engineers do not think highly of either corporate or contract engineers. These organizations tend to have engineers with average or above average technical capability but below average maintenance experience. Corporate engineers have more start-up experience than contract engineers, but they are less likely to treat the plant engineer as a customer. Nothing infuriates a plant engineer more than a corporate engineer refusing to alter a design for better reliability because it threatens the budget or the schedule. Even though the project money is allocated to and paid by the plant, corporate engineering acts as though it has total responsibility for the money. Therefore, plant engineers feel they have no control over corporate engineers, and so, prefer contract engineers. The lesson here is to treat whoever uses the installation as a customer and seek to identify the short and long-term needs of that customer. While all this seems like common sense, it isn't widely practiced. As a result, corporations have begun to make their employees sit through several days of sheer excitement called "Total Quality" programs to explain this fundamental principle. (This comes too late for most, based on the present binge of decentralization.)

Playing The Game

In most large companies, you don't have to worry about your mistakes adversely affecting your potential for raises and promotion. Your performance review is based on a set of goals concerning whether the instruments were purchased on time and within the budget but not on whether or not they function properly. You can include lots of extra instruments in your estimate, buy the ones with the shortest delivery, have instrument systems fail right and left, and still meet your goals. Even if you mess up the estimate and fall short of your goals, don't worry. Your promotion and raise are based on how well you play the corporate game, not on your performance review.

There are many rules to the game, but the foremost is "Don't rock the boat." The primary goal in the life of a manager is to not upset the next level up manager. The sole purpose of each level of management is to soothe the next level through numerous reports.

Disastrous results are OK as long as everything is done in compliance with corporate procedure and the data is properly reported. Innovative thinking is encouraged figuratively, but not literally; free thinking and radical methods make managers nervous. If you want to maximize your monetary rewards, smile and say "yes" a lot, and attend all management training programs. You might even become a manager, in which case your raises will automatically be larger and you never have to do anything more technical than to analyze the performance of your investments.

You have then reached the ultimate goal of being a non-practicing engineer who gets to set the salary of all the practicing engineers. Even when a technical ladder of promotion equivalent to the management ladder exists, managers at the same level earn more money because they set the raises. Imagine what the raises of the managers would be if they were set by the practicing engineers!

These words are not intended to make you a manager. After all, this is a technical magazine. However, if you learn everything presented here, you will be well on your way to becoming a prima donna who can strike fear into the hearts of at least first level management. Any direct influence on upper levels is too high an aspiration and nearly impossible to achieve because these lofty personages do not want to hear and be confused by mundane facts.

Success?

Somewhere along the line, someone may ask you if you are a successful instrument engineer. First, you must define "success." Some people like to use the words money, prestige (including fancy titles), company or national reputation, job security, performance, or even knowledge.

For an instrument engineer, all of the above may be wrong. The real source of "success" probably depends on how you are perceived by your management or clients. It doesn't matter if you work for a large corporation or a small engineering company. A little luck, work on high visibility projects, and being fortunate enough to have agreeable, competent people on your jobs are the real keys to success.

Learn and understand the "Good-Good" and "Bad-Good" theories. Good-good engineers try to do the best job possible, while helping other design engineers, and even project managers. Bad-good people are rotten to anyone when it advances their own reputations. If they are very intelligent, they are also very dangerous. Everything that goes wrong is someone else's fault. Identify these villains quickly and stay away from them. Most nice guys can't play their game and win.

INSTRUMENTATION RULES OF THUMB

1. Being experienced in instrument engineering means you have made and corrected a lot of mistakes.
2. You must check out and start up your own designs; otherwise you will make the same mistakes again.
3. You must treat the user like a customer if you and your organization are to survive.
4. Confrontations with management indicate suicidal tendencies.
5. If you want a sense of accomplishment, stay technical. If you want to make money, become a manager, sales rep, or outside consultant.
6. Attend as few meetings as possible. If you like to attend meetings, you will never be a good instrument engineer and probably will end up in management.
7. If you work for a corporate engineering department, join every airline's frequent flyer club.
8. Buy and use luggage that doesn't have to be checked. Checked luggage gets lost or destroyed. Having dirty socks and underwear fall out of your briefcase at a meeting is tacky.
9. Create the perception that you do wonderful things that produce amazing results.

Woody's Performance Review—What's Inline Next for pH Control?

Through some existential twist of fate, Woody Allen finds himself somewhere in the seven hidden dimensions of the eleven dimensional universe where he is assigned the task of commissioning pH systems. We join Woody at the end of the year during his performance review with his supervisor, Mr. Bossman.

Woody: What is all this stuff on your desk? Is this a swap meet or something?

Bossman: Sit down, Woody, and tell me what you think of pH loops.

Woody: I feel pH loops are divided up into the horrible and the miserable. The horrible ones are the ones with the steep titration curves and all the others are the miserable ones. If they are installed on a sump they are worse than death. Anyone who has spent an evening with an insurance salesperson knows what death is like.

Bossman: I heard you refused to start up a sump pH loop.

Woody: Suicide is not a middle-class alternative.

Bossman: Some of our most educated environmental specialists specify sumps for waste treatment pH control.

Woody: Ph.D.s show you that you can be smart and still be stupid.

Bossman: Do you have something against Ph.D.s?

Woody: Are you kidding? Some of my best friends are Ph.D.s.

Bossman: I have a Ph.D.

Woody: There are exceptions.

Bossman: Let's get serious. What pH systems did you commission?

Woody: I survived 4 pH start-ups. They all had relatively steep titration curves near the set point. The first had a pipeline static mixer for the first stage and a well-mixed vertical tank for the second stage. The traditional method was two well-mixed tanks in series because inline systems develop nearly full scale oscillations for steep curves. The plant had one inline system on a flat curve that worked great so they wanted an inline system. I had a good pH book with me but I figured I needed something more—a Rabbi, analyst, or interplanetary genius. When the inline system got going and the field pH meters pegged upscale and downscale every two seconds, my whole life flashed before me culminated by a performance review. Fortunately, something was different about this system. It had a microprocessor-based controller with an adjustable measurement filter. In desperation, I increased the filter's time constant to about 12 seconds, which cut the oscillation amplitude displayed in the control room to about two pH units. This looked better, but appearances can be deceptive. To my amazement, with tuning settings normally for a flow loop, I could change the set point and control the average of the oscillations. The pH in the downstream tank confirmed the average. To make a long story short, the last tank's pH system did not work, even though it had a turnover time about ten seconds, because of excessive reagent delivery delays from reagent dribbling down a large and long dip tube. So we did the entire neutralization with the first stage. The well-mixed tank was used as a wide spot in the line to average the process pH oscillations in a manner similar to how the digital filter averaged pH measurement oscillations. I told the operators not to look at the field transmitters if they had a nervous stomach (they still oscillated wildly). (See Fig. 11-1.)

Bossman: Didn't the measurement filter add lag to the loop?

Woody: Yes, but the rest of the inline loop was very fast due to close coupling of instruments and high velocities so that the loop period was still small. More importantly, my deadtime on the job site was shortened.

Bossman: OK, let's move along. I have important things to do like letters and, uh, letters.

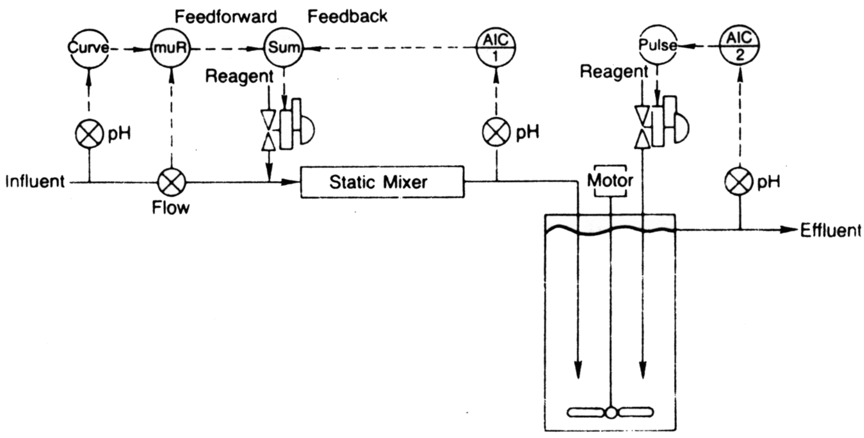

Figure 11-1. Start-up #1.

Woody: Is it true they call you "memo man"?

Bossman: Yes, and I am thinking of one right now about your performance review.

Woody: I better move along. Since the last inline system worked so well, I decided to use it on the next system, where the flow rates were so large that the cost of installing and agitating a well-mixed tank with sufficient residence time would be too expensive. However, the plant was violently opposed to inline systems since they had tried them for steep titration curves. I explained to them that with a small signal filter, close coupling of instruments, precise control valves, median selection of triplicated measurements, pressure control of the reagent, and the use of the titration curve to translate the controlled variable from pH to reagent demand, the oscillations would be controllable. They thought I had a degenerate's mind. I took this as a compliment and proceeded with the installation. During start-up, the field transmitters showed severe oscillations, but the controlled variable displayed in the control room drew a straight line. Using flow loop tuning settings, the loop was on set point within a couple of minutes. The pond that the inline system discharged into lined out within a few tenths of a pH unit from set point. It eliminated the need for further treatment downstream (which they never were able to do successfully) and enabled the plant to declassify the surface impoundments. It saved the plant the ten million dollars by not having to reline these impoundments. Can I have a raise as large as yours? (See Fig. 11-2.)

Bossman: No, but I might overlook your attitude problem.

Woody: The next start-up had the classical three well-mixed tanks in series. The tank sizes were minimized (one to two minutes residence time) and electronically set metering pumps, median selection of triplicated measurements, and titration curve translation of the controlled variables were used to insure the final tank was within one tenth of a pH from set point for extreme variations in the waste stream. When I got to the job site I found metering pumps

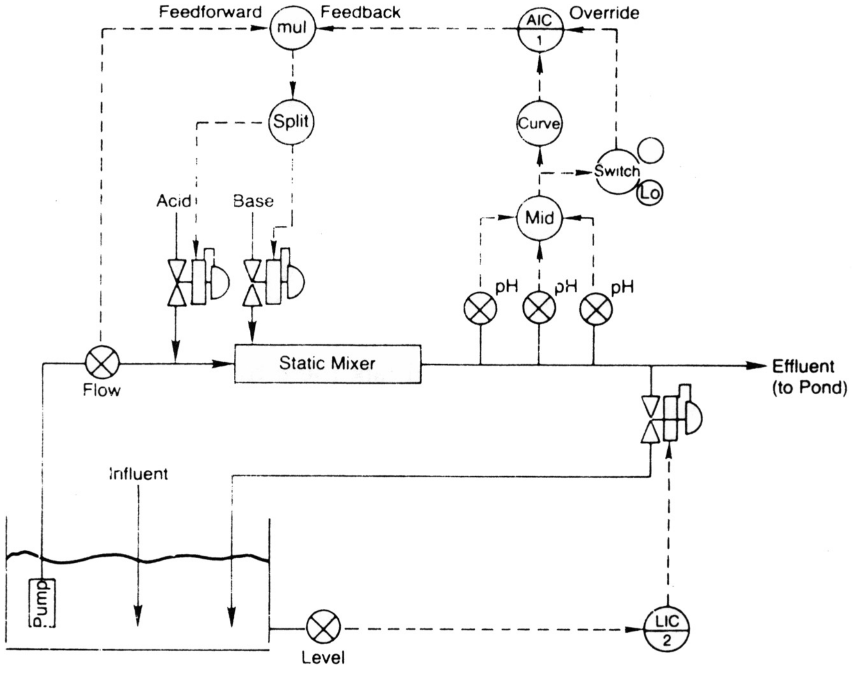

Figure 11-2. Start-up #2.

as big as battleships on the first and second tanks despite the small reagent flow requirements. The discharge piping and dip tubes were also large to be consistent with the impressive size of the pumps. The last tank had a metering pump that was small enough to fit inside a breadbox. The mechanical engineers didn't like this one because it was too tiny to kick. Burned by excessive reagent transportation delays so many times that I was scarred for life, I insisted on a reagent delivery test prior to start-up. Forty minutes after the acid pumps were turned on, the indicated pH still had not decreased from neutrality. The chemist got suspicious, took a sample, and proudly announced that his lab meter read two pH and that my pH measurements must be screwed up. Since I had triplicated measurements that all agreed, I insisted on triplicated samples. The chemist came back later and sheepishly announced the next samples were neutral. He had used a breaker with dried up residue of a two pH buffer for the first sample. Between one and two hours after the start of reagent flow, the pH in the second and then the first tank finally dropped. The results were impressive enough to convince them to install a restricting orifice and small diameter tubing for the dip-leg so that the lines would run full. The start-up went extremely smoothly with controller gains of four, forty, and over one hundred for the first, second, and third stages, respectively, due to successively smaller reagent pump capacities. The titration curve translation was not needed for the last tank due to the high controller gain and tight control band. As I stood on the grating near the top of the tanks and watched the digital pH displays flash different numbers, I realized that even in well-mixed tanks the measurements never agreed within a tenth of a pH and that median selection helped average out this randomness without the introduction of an additional lag. For one brief moment, everything seemed to fit together perfectly in my mind and I was moved in a very profound way. (See Fig. 11-3)

Bossman: What was the most important thing that you learned?

Woody: Everything that our parents said was good for us has turned out to be bad like sun, eggs, red meat, and big dip tubes.

Bossman: What about your last start-up?

Woody: I did it by phone. I had to be in two places at the same time. Normally, start-ups take priority, but if you had the choice of a plant near the most beautiful beach in Florida and a plant in the middle of the corn fields of Iowa, which would you choose? I was somewhat confident that I could talk them through the start-up because it was an inline system and everything was in working order when I checked it out. I knew that an autocatalytic reaction could develop in the scrubber that could cause the pH to plummet to zero, but the use of inline control of recirculation pH set by sump pH instead of direct control of sump pH would probably work. Then, another plant said they tried pH control of a similar scrubber and had failed and were now putting in several very expensive analyzers. I got a little worried. To everyone's surprise,

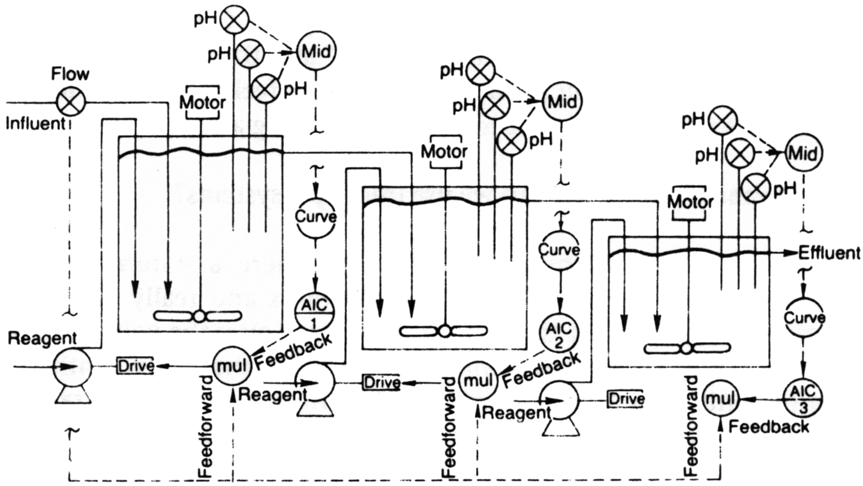

Figure 11-3. Start-up #3.

the start-up went well after commissioning the cascade of sump to inline pH and tuning the sump pH for slow action (little reset) and tuning the inline system like a flow loop (lots of reset). I talked the plant engineer through the adjustments. I was lucky that he picked up on key ideas quickly and had an open mind (he hadn't been out of school long enough to develop a bias against corporate engineers). (See Fig. 11-4.)

Bossman: Have our reorganization and cutbacks improved the plant's opinion of us?

Woody: Their attitude reminds me of the joke about two people at a restaurant they have been going to for years. One of them says, "The food at this place is really terrible," and the other says, "And they have such small portions." The plants feel that we are lousy but that we are not onsite long enough during a start-up.

Bossman: Where else can we use inline pH systems?

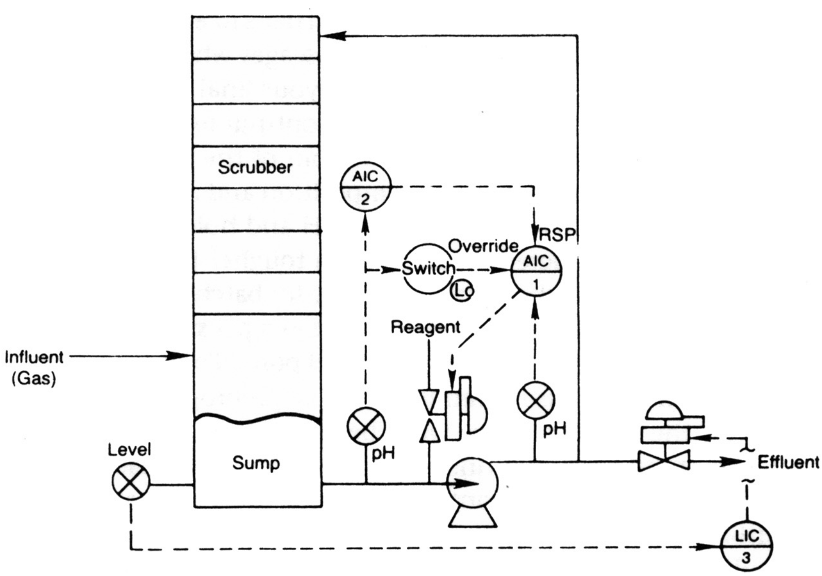

Figure 11-4. Start-up #4.

Woody: For continuous pH control where a volume exists downstream to average out the oscillations and really tight pH control at a set point is not required. The pH after the volume may draw a straight line, but it will usually be more than a few tenths of a pH unit away from set point. The nonlinearity of the titration curve guarantees the average of the pH measurement and the hydrogen ion concentration oscillations will not match exactly.

Bossman: What about batch pH control?

Woody: Batch pH control is much easier. Remember back to the titrations you did in your chemistry lab? That's like a batch system. As long as you are careful to use tiny doseages while on the steep portion of the curve or as you approach your final pH, you can do a good job of reaching it exactly. A continuous system is like someone had drilled a hole in the bottom of the beaker and was adding more sample of different composition and at a variable rate behind your back. Reaching the same pH and holding it there by just looking at the meter is now a much tougher job. Almost any strategy that avoids overdosing is suitable for batch pH control. For flat titration curves, this can be as simple as a pH switch set to trip off the reagents just before reaching the set point. For steeper curves, proportional-only pulse duration control or preprogrammed pulses based on the titration curve work better. It is important here to realize that reset action cannot be used for batch feedback control of pH in a vessel and that control action must wait at least one loop deadtime after the finishing doses to allow the pH to coast toward set point. Another alternative exists if there is a recirculation line and microprocessor-based controller available. An inline pH control system can be used (reset action is OK here) with a vessel pH cutoff as a backup to provide a simple and fast method of batch pH control. Here the reagent flow is automatically maximized at the start and minimized near the end. The operators find it easier to understand and interface with than some complex batch strategy for reagent doses. It is critical to remember that for steep curves an adjustable measurement filter of ten to twenty seconds, valve positioners, linear trim, reagent pressure control, and close coupling of components to the inline mixer are absolutely necessary and that median selection of the measurements and titration curve translation of the controlled variable are highly desirable.

Bossman: I think you deserve a good rating, but you do some things that are not consistent with company policy, like wearing a tee shirt with a tie and and a tuxedo jacket with jeans and tennis shoes.

Woody: Despite the fact that I do a lot of silly things, surely you realize that down deep I am not evil or anything and that I am just reaching or searching for something.

Bossman: Don't call me Shirley. How about you reaching over and signing this appraisal form?

Woody: Good grief! It says I was a good engineer this year. Didn't I have to do something really important like buying savings bonds or something?

What If the Starship Extraprize Had a DCS?

As I reached deep down into my Cracker Jacks box while sitting at the control room console on the graveyard shift, I found an extra prize. It looked just like a starship. My mind started to wander (as it usually does), and in a brilliant flash of bizarre creativity I thought to myself what if the starship Extraprize had a distributed control system?

Captain's log star date 0625 1988 2204.32: The starship Extraprize ventures out on its first voyage after a major restructuring of its engineering department and installation of a distributed control system (DCS). Our mission is to boldly go where no control engineer has ever gone before.

Captain: Bridge to engineering. Stouty, we are going backwards. Stouty, where are you?

Solo: Captain, Stouty took early retirement offered in the incentive program to disappear during the reorganization.

Captain: Who is left in engineering?

Solo: Captain, no one. It was decided that engineering isn't needed anymore.

Captain: Cut the "Captain" formalities. Who installed and configured the DCS?

Solo: The vendor who sold us the DCS.

Captain: Where is the vendor?

Solo: About a million light years away.

Captain: Where is Spook?

Solo: He is taking a stress management course.

Captain: Do you mean it's just you and me on this starship?

Solo: No. We have three thousand lawyers on board to deal with possible litigations from Klingons.

Captain: I thought the Klingons used phasers against us.

Solo: They found out that lawsuits are more effective.

Captain: We are still going backwards. The DCS console says there is a configuration error. It can't be too difficult to figure out how to modify the software. Let's go find the instruction manual.

Solo and the Captain enter a large room whose walls are completely covered by manuals, all of which are the same color and labeled with just the manufacturer's name.

Captain: Have we entered another world?

Solo: This is the world of DCS documentation where everything looks the same but is different.

Captain: Which manual is the one we want?

Solo: Well, it depends upon the type of console and controller used and its revision status and what level in the hierarchy of software the mistake resides.

Captain: Let's just grab a manual. The software can't be that much different from one device and level to another.

Solo: There are seven different configuration languages and five different keyboards. You need to go to DCS school for ten weeks to see the whole picture, but it is physically impossible to become proficient in all aspects of it Our best bet is to seek help elsewhere. Maybe we can flag down that spaceship we passed before we got stuck in reverse.

When the unknown spaceship is found, Solo manages to put the DCS in track mode so that the starship is latched into the same trajectory as the spaceship. Solo and the Captain beam onboard the spaceship.

Captain: Greetings. Can you help us with our DCS?

Being: Sorry, we don't have the time or the money.

Captain: Who are you?

Being: We are project managers (PMs) from the planet Mirth.

Captain: Why are you in a spaceship?

Being: Just before our engineering department was dissolved, the instrument engineers arranged a special mission for us that only we could accomplish.

Captain: Where are you going?

Being: Our destination is the planet Omega. We don't know why we are going there, but the trajectory is fixed as defined in the project definition. We are on target—within budget and on schedule.

Captain: Do you realize you don't have any landing gear and that Omega is a barren planet?

Being: The correction of that will have to be a post-project expense. You must detach your starship from us. You are causing extra fuel consumption, and we only have enough to get to our destination.

Captain: Sounds like you need to be reprogrammed but we can't even figure out how to reconfigure our own DCS. Goodbye and good luck.

Back on board the starship, Solo finds a console key that says "opposite." After pressing it, a warning message flashes on the screen that says you are now anti-matter and are proceeding to the nearest anti-matter planet for reeducation. After landing on the planet, Solo and the Captain are introduced to the aliens. The aura of well-being is awe-inspiring. Did these aliens have super intelligence or did they have a better insight and perspective of what is right and wrong?

Captain: Can someone here tell me how to reconfigure my DCS?

Alien: No problem. Since you have been transformed into anti-matter where everything is the opposite, the software is now easy to understand and consistent. Your DCS uses the same language as ours. Its rules are documented on this card in my wallet. You can learn them in an hour.

Captain: Do you have many software revisions after commissioning a starship?

Alien: No. Our PMs prod us to spend enough time and money to get the best hardware and software. The DCS manufacturers insist that all systems be designed to self-document and be easy to select, apply and use. The vendors understand our applications and ensure we get the least expensive DCS to meet our needs. Our managers and directors are truthful and straightforward and seek ways to be adaptive and flexible even if it reduces their own power. The corporate engineers are dedicated to making the DCS operable and maintainable and don't infringe on the freedom of choice of the end user. Everyone has learned to sacrifice personal interests and short-term gains to achieve real success.

Captain: We better leave soon so the people on our planet of origin can benefit from this knowledge.

Alien: That is not such a good idea. If anti-matter and matter meet, there is an annihilation and a tremendous release of energy.

Captain: Do you have time sheets?

Alien: What are time sheets?

Captain: We'll stay.

Alien: You're welcome, but the lawyers have got to go.

What If the Starship Rent-a-Ride Had an Expert System?

We left Solo and the Captain in the anti-matter universe (see the previous article). While it seemed like paradise, the fact of the matter was that everything was too perfect. Solo and the Captain got bored; there was nothing to laugh about. They decide to make the transformation to ordinary matter and journey back to their headquarters to see what's new or in this case, what's old.

When the starship arrives home, they are informed the Extraprize has been sold. They are given two tickets for a rented starship, the theory being that a zero divisor will make the ratio of return to capital look great.

Captain's log star date 0520 1989 0852.30: The starship Rent-a-Ride ventures out on its first journey. Their mission is to boldly go where no control engineer has ever gone before and where no management would ever want to go.

Captain: Where is Spook?

Solo: After his stress management course, he completed a breakthrough session. He is scheduled for courses on total quality, managing a diverse workforce, and vision verbage. He no longer says stuff is illogical, says he sees the big picture, and smiles a lot. He has been labeled a rising star in the Federation.

Captain: I guess his journeys into outer space are over. Why do we get an alarm every five seconds?

Solo: This starship seems to be composed of packaged equipment with instrumentation left over from the twentieth century. There are a lot of devices called "d/p cells." They have primitive analog circuits. While the catalog accuracy specs are OK, they are for laboratory situations. I suspect the "d/p cells" are having trouble with our ambient conditions. The good news is that we have a real-time expert system. The bad news is that we have to program it in our spare time.

Captain: What spare time? With Stouty, Spook, and the rest of the crew gone, we're already overworked, but, we have no choice. These alarms are driving me crazy. Can we rent a knowledge engineer to help us get started?

Solo: Yes, I have already made arrangements for one. However, he wants us to beam his "Beamer" aboard.

Captain: What's a "Beamer"?

Solo: I believe it is a BMW. It's quite common for these Yuppy creatures to refuse to travel without them.

A creature, immaculately dressed in a navy suit with a power tie climbs out of a black BMW in the transporter room.

Bawhstonian: It is a pleashaw to meet yuh. I would have gotten huh soonah but it took an ouwah to pock my caw. Yuh cuhtainly have a lodge numbah of alahms.

Captain: Solo, what language is he speaking?

Solo: I believe it is a form of English spoken in the planet called Boston.

Bawhstonlan: Yuh will nevah know what's wrong unless yah expuht system diagnoses instrument foilahs.

Captain: We have thousands of different instruments that we know nothing about, how can we program it working part time?

Bawhstonian: Yuh need to make the faumulahs and rules as genuic as possible and develop a knowledge gathahing strategy to provide the detail poametahs.

Captain: Can you help us?

Bawhstonian: No, I'm too expensive. Besides, I have to leave for an ahlee powah lunch.

Unlike their distributed control system, the expert system is schematic-based, learns relationships from connections, uses menus and a mouse, provides prompts based on information already entered, and has no need for a mammoth documentation system. Since it is almost fun to work with the system, the Captain and Solo travel at warp speed to dilate time enough to allow them to complete the program, which is fortunate, since they are about to face the Klingons with packaged equipment—an instrument engineer's worst nightmare come true.

Solo: The Klingons have fired their phasors. Our shields withstood the first salvo, but they are starting to oscillate. The expert system says a thousand "d/p cells" have developed erroneous readings.

Captain: Use inferred measurements from models based on the installed characteristics of the final elements to functionally replace the lost signals, tell the Klingons they are missing a staff meeting to buy time, and fire our phasers to show them we mean business.

Solo: The Klingon shields are oscillating. The amplitude keeps getting larger. Some of our phasers are making it through, and severely damaging their ship, its like a video game and we're racking up a big score.

Klingon: Stop firing!

Captain: Why are your shields oscillating?

Klingon: Why did yours stop? We're using the same special instrumentation supplied with packaged equipment.

Captain: We'll let you go if you tell us how you found out confidential information about our instrumentation.

Klingon: A Klingon disguised as a computer repairman took the personal computer with your instrument list on its hard disk. It indicated the pervasive use of packaged equipment and a secret new device called the "d/p cell." The whole Klingon fleet is now outfitted with them.

Captain: There's a critical mass before the secret powers of the "d/p cell" are unleashed. Obviously, you don't have enough of them. (Privately to Solo: Let the Klingons go, I don't think we have to worry about them for awhile. Besides, we'll be barraged with a staggering number of liability lawsuits.)

Postscript: The expert system provides documented evidence of the starship's severe performance problems with packaged equipment instrumentation. This provides enough justification for the massive replacement of such instrumentaiton instead of the ageless tradition of gradual substitution. The "d/p cells" are placed on the open market where they are eagerly grabbed up by packaged equipment vendors and sold to Klingons. The Klingons become the joke of the universe. Years later, when some astute technicians figure out what's wrong, the Klingons scrap their systems and sue the Federation for improper software security, releasing false information, and flooding the market with inferior devices.

Funny You Should Ask a Process Control Engineer

Have you ever been asked a question on process control that leaves you at a loss for words? Here is your chance to amaze friends and relatives (or at least baffle them) by learning some good concepts and analogies. If you understand the ideas and their significance conveyed in the answers to the following devastating questions, you might just become famous by Friday.

What is loop deadtime? Deadtime is that period of time from the start of a disturbance until the controller makes a correction that arrives at the same point in the loop that the disturbance entered. The controller needs to see the upset, react to it, and send the correction to the right place. To appreciate deadtime, consider when you go to a party and start drinking. The period of time between when you first take a drink and when you first recognize the effect and bypass the next round for coffee is deadtime. As the deadtime approaches zero, the portion of the open-loop error that appears as the closed-loop error approaches zero, and nonlinearities and high process gain (high proof and small body) become unimportant.

What is the open-loop error? Well, if you continued drinking, your state of intoxication would exponentially increase and reach a steady state with you passed out on the floor. The open-loop error is the alcohol concentration at this point. It is the result of having the controller (your mind) in manual or disconnecting it from the loop (your body). A smart party animal knows when to say "when," according to a famous seer.

What is the closed-loop error? The peak excursion of the alcohol concentration until the coffee first enters the bloodstream is the closed-loop error. For the well-tuned PI and PID controllers, the best you can do is limit the closed-loop error to the maximum deviation from set point that occurs at 150% and 110%, respectively, of the loop deadtime. Often, the party mind does not behave like a well-tuned controller, and the closed-loop error reaches an unacceptable level for driving, especially if there is a non-self-regulating process.

What is a non-self-regulating process? If your liver breaks down and removes alcohol at a fixed rate, the alcohol concentration ramps, which is characteristic of an integrating process. If an increase in intoxication causes a decrease in your restraint, you might start drinking faster. This is positive feedback, and the acceleration of the concentration corresponds to a runaway process. People with such non-self-regulating processes are easy to spot at parties; they are the ones with the lampshades on their heads.

How can you reduce loop deadtime? When I first left home, my father said to me, "Be as honest as the day is long, don't talk when you should listen, and don't be fooled into thinking that deadtime compensators can eliminate deadtime from the loop." I thought this was extremely strange in that I only planned to walk around the block. Some professors must not have gotten the same words of wisdom, because they are convinced they can mathematically cancel out the deadtime term in a loop by an advanced control algorithm. Deadtime cannot be reduced without violating the basic principles of classical physics. While the theory of relativity concludes that you could contract distances and dilate time and hence shorten transport delays as you approach the speed of light, unless you have Scotty and warp drive on your loop you are stuck with deadtime caused by the vessels, piping, and instruments within your loop. If the loop deadtime is well known and considerably larger than the largest time constant in the loop, a deadtime compensator, such as the Smith Predictor, can make the loop perform better. However, the use of a Smith Predictor on a non-self-regulating loop is downright dangerous. In general, your best bet is to look for the biggest sources of deadtime and work on a change in the plant or instrument design to reduce their contribution to the loop deadtime. Figure 14-1 is valuable in the search for deadtime because it reminds you of all the places to look. Also, the figure helps you to remember that if the disturbance is zero, the open-loop error is zero, and if the disturbance is much slower than the loop deadtime, the closed loop error is small. However, you are again faced with a change in equipment design. Additionally, the figure shows that output and input signal characterization should be used for nonlinearities on the input and output, respectively, of the process. The figure can be used in a lot of ways. Have you ever had guests stay too long? Just flash a slide of Figure 14-1 on the wall and start talking about loop performance. Some guests have hurt themselves in the rush for the door.

What if you can't change the equipment design? Then your best bet is to use feedforward control for disturbances that can be measured when the deadtime in the disturbance path is smaller than the deadtime in the correction path to a common point in the process; use cascade control for disturbances that can be isolated by an inner loop when the inner loop deadtime is smaller than the outer loop deadtime. The advanced strategy should be checked out via real-time simulation.

What is real-time simulation? Decades ago, we used analog computers with gobs of potentiometers and a maze of wires to simulate the plant. Fortunately, we survived the analog computer scare of the sixties. The principle is the same, but now one uses spare or virtual digital controllers or a host computer connected to accept the control system outputs and return the control system measurements

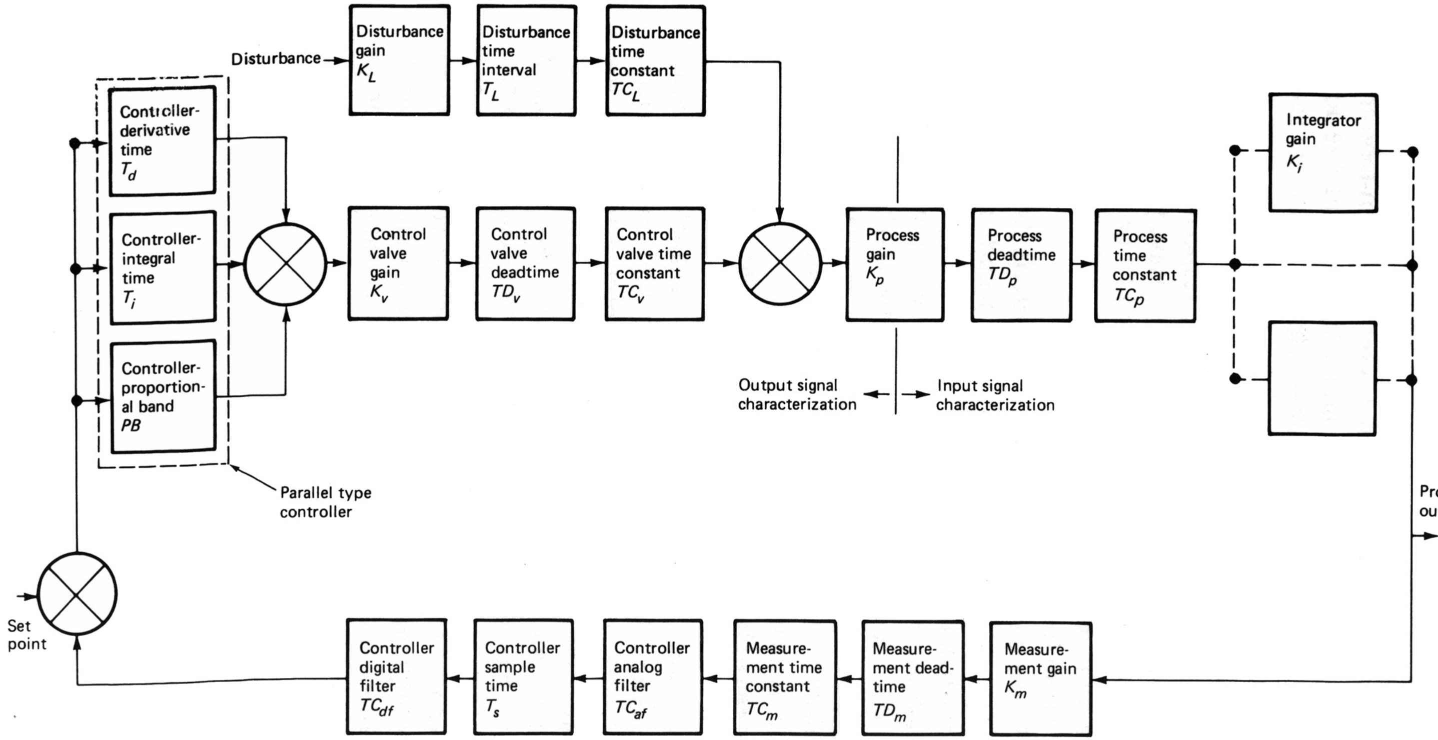

Figure 14-1. Output and input signal characterization should be used for nonlinear components between the controller output and the process and between the process and controller input, respectively. *From the ISA ILM Continuous Control Techniques for Distributed Control Systems.*

in the same time frame as the real process. It is needed whenever you have a batch operation or a complex control scheme. While real-time simulation is never as good as the real thing, the real thing can be too exciting. It's a lot easier on your nerves to confront configuration problems during simulation than during plant operation. Real-time simulation is not needed if you plan to retire before start-up.