Future of Supervisory Systems in Process Industries: Lessons for Discrete Manufacturing

Ian Nimmo
Honeywell Industrial Automation and Control

Edward L. Cochran
Honeywell Technology Center

This paper provides a brief review of the history of process control and automated systems and an overview of what we can expect in the future, together with some lessons that may be applicable to the development and operation of highly automated discrete manufacturing systems.

Supervisory systems in the past

During the early 1970's—the beginning of the computer control era—memory and disk space were very limited, and the interface to the computer was primitive. The initial control capability provided what we describe today as Supervisory control. The computer had limited output capability, but could influence instrumentation systems by changing the setpoint to PID electronic controllers.

Early applications

Starting up a continuous polymer process is challenging: The characteristics of the product change as operating conditions change, and they also change as a function of plant throughput . These changes require different process control setpoints. In those days the operator used to manually tweak the controller setpoints but would often overcompensate and introduce new disturbances to the process. The early computers helped by automatically changing the setpoints to the controllers. Within a couple of years technology development enabled the computer to carry out full PID control using a mathematical model of a PID controller. This allowed not only automatic changes to the setpoint but changes to the tuning constants of the control loops as well. This was especially useful for non-linear processes.

However, the user interface was still very crude: Plant supervisors used to enter data by identifying the address of a variable and entering data by first entering the address and then the

---

Acknowledgments. The authors gratefully acknowledge the contributions made by Peter Bullemer to the ideas presented in this paper.

Preparation of this paper was in part supported by a grant from the National Institute of Standards and Technology's Advanced Technology Program to the Abnormal Situation Management Joint Research and Development Consortium.

data, using a set of 16 piano-style keys. The process was very dangerous, because the opportunity for error was so high, especially when we consider that the data was not entered in normal base 10 numbering but in octal (base 8) or binary formats.

Soon more powerful monolithic control computers were introduced but were limited by small amounts of memory and little or no disk storage capability. The Supervisory Control Computer System held the main program and instructions for the "target" control systems. The Supervisory system downloaded the software to the control system and the supervisor initiated a booting or startup sequence. During plant operations, the target system ran its PID control algorithm and fed information about the process back to the supervisory computer system. Other information, such as batch records, historical plant data, and any abnormal condition reports was recorded manually.

The control system user interface improved slightly, allowing operators to change the control computer setpoints and alarm values directly at the control computer system. However, this required regular updates back to the Supervisory system to allow easier synchronization should the computer shut down and require reloading of the currently running program and data. This meant that the Supervisory Control Computers required several versions of the controller software, including an original or "virgin" image, a copy of the original image which had been modified during operations and which held the latest setpoints and other data parameters, and any development or test versions of the software.

During the late 1970's the Supervisory system became a very powerful system, especially on batch processes, and as Digital Corporation's VAX computers and commercial VMS operating system became available and cost effective, the use of off-the-shelf database products improved the historical data capture and replay capabilities.

As Supervisory systems continued to evolve, Honeywell introduced a more powerful Distributed Process Control Computer System (DCS) that could work independently of Supervisory systems. The DCS had better human interface, historical data capture and replay capabilities, and powerful alarm management systems. Supervisory Systems then became a place to store large amounts of historical data and the optimization or mathematical modeling environment.

Today, the DCS works independently of any supervisory system, although some responsibility for advanced control remains delegated to more powerful supervisory systems. As new open system technologies are developed for control, it is becoming difficult to tell were the control system starts and ends and if a separate supervisory system as we know it is required anymore.

Adopting (and adapting to) new technology

The evolution of any technology has consequences for those who the technology is intended to help. The development path is rarely straight, and unexpected consequences must often be dealt with. Control systems are no exception: During the evolution of control systems, plant management and operations personnel have been required to continually reassess their abilities, progress, and needs, and adopt the technology as it becomes available. There are human and institutional consequences to this, and often the technology must be adapted to by the users as much as it is adapted for the application. The lessons learned during this process transcend the particular process and control technology involved, and provide the only hints we have available to us about what the issues of the future will be. In the sections that follow, we describe a number of examples of the sometimes unexpected consequences of the evolution of automation technology.

The ICI shrink-wrapping robot

The development of more capable computers allowed mechanical systems such as robots to be used for many highly labor-intensive and perhaps dangerous or unpleasant activities that required little in the way of human cognitive skills. Initially the robot was used for repetitive tasks that people added no value to, and for safety reasons people and robots would never work together. The robots were put into large safety cages and isolated from all contact with people. The safety concerns, although based upon legitimate concerns, often had deeper origins, as the following experience by the senior author illustrates.

> I remember my first robotics application within Imperial Chemical Industries (ICI) in the UK. The company was keen to identify and use time-saving and cost-effective automation. A group of engineers with ideas for applications of this technology was formed and several trial applications were investigated. As the technology was tested it became obvious that there were many reasons why people should not trust this technology. In particular, the British Government had identified high risks to people from the robots themselves as they were very powerful and could seriously hurt or even kill a person them should they get in the way of the robot's movements. The initial solution was simple: Isolate the robots from people by placing them in high security prisons. Should a person need to enter the room were the robot was working elaborate isolation and safety interlock equipment was employed to secure the robots and make them safe. Unfortunately, this solution made robots expensive and no longer cost effective for a variety of applications.
>
> My first robot was no exception to this rule. I had the great idea to use my robot as a shrink wrapping system for 300 Kg bales of staple fiber. These bales looked like very large cotton pillows, and were wrapped in polyethylene which needed to be shrunk by heating the polyethylene at very high temperatures. The existing system consisted of a very large open electric oven with a conveyor through its middle. The conveyor carried the bale from the wrapper to the oven and then to labeling and the warehouse for transportation to customer sites. The electric oven was very inefficient (large amounts of heat escaped out the open ends) and was costing millions of dollars in energy usage; further, the wrapping was poor due to the large side ears that resulted in the wrapping and shrinking process.
>
> My idea was to train a robot to hold a hot air gun and effectively shrink the bale by blowing hot air in a motion similar to spraying a car. I had witnessed robots performing this activity very effectively in the automobile industry. The first step was to develop a gun suitable for this activity and a local company soon came up with an effective prototype which we could use with the robot. By chance, another part of the company had been trying to use robots in the manufacturing of explosives and after an unfortunate accident the explosive was re-classified and no electrical equipment was allowed during the manufacture of the product. This allowed me to acquire the perfect robot for my application, and one that had

already been paid for by my company. With minimal negotiating the robot was transferred from Scotland to England and was on its way to save the company a lot of money and produce a better package for our customers. What could go wrong?

The robot arrived in bits but with the aid of some untrained personnel and in true engineering tradition we figured out how to re-assemble the robot. Soon the basic computer was connected and, after a little experimentation, the robot was programmed in assembly language and soon responded to my every request. The sight of such new and exciting technology set the political wheels in motion. The craftsmen who would have to maintain it saw it as a large threat; the operators who were to interface with it could not make their minds up, and senior management needed to be convinced. So my boss decided that we should gather all the interested parties together and demonstrate of the capabilities of this new technology.

Given 48 hours, without all of the equipment for the final application and no time to do a correct installation in any event, we needed a demonstration task. We came up with the idea to put a box in front of the robot and write the name of a the senior executive who would be present to demonstrate the flexibility and precision of the equipment. Programming was easy and soon we had a working demo and we had confidence we could win everyone over to this new idea and effectively use this technology. At the end of the demo I was to give a short presentation on the real application and the benefits and savings.

On the day of the demo one of the mechanical craftsmen noticed a small hydraulic oil leak and without asking anyone he decided to just go ahead and fix it as he would any other piece of plant equipment. However, the very small leak was not fixed by a couple of turns of his wrench, so he took the pipe off and went to find new seals for the connection. He was not aware that the piping was designed for very high oil pressure and required that specific maintenance procedures be followed. A simple oversight?

After the audience assembled, and after a few brief words, the demo was started and to my horror the robot went crazy: Instead of writing the executive’s name on the carton it used the felt pen as a weapon and destroyed the carton with several lethal blows. Then, with one effortless blow the robot knocked the three-foot-square carton towards the amazed audience and then repeatedly drove the pen into the ground as if it were a bayonet into a dummy until all semblance of the pen was totally gone.

By which time I was able to make a safe path to the emergency button on the console.

Without any opportunity for discussion, I was told to get rid of the machine, and to never use the word “robot” in any discussion with anyone ever again.

Lessons

There are lessons to be learned from this story on multiple levels.

Configuration management as root cause

Traditionally speaking, the root cause of this incident was inadequate configuration management. In fact, there were a host of organizational culture issues that needed to be resolved, involving work permitting, skill certification, project planning, and safety.

This was a large set back for technology at ICI, and from the experience the senior author learned to never do anything by half measures: Every piece of equipment introduced should have the protection of competent people working on it, professional installation, clear operating procedures, and management of change control. If we had had the system identified as a plant asset, it would have been put on the maintenance system, and the craftsman would not have worked on the robot without a work-order and a permit. The work would have been reviewed and the correct replacement parts would have been identified from the manufacturers recommendations. As we assessed the incident we discovered that the craftsman was also in a dangerous situation during his maintenance activity — he was not aware of the high pressures involved; the system had not been isolated correctly, and he could have been seriously injured. The only isolation was that the computer program was not running and traditional electrical and hydraulic isolation which was not done would probably not have been sufficient based on the capability of this beast.

Why did this one event temporarily stop progress in the application of automation?

This one incident brought the introduction of automation into this company to a screeching halt. It had a disproportionate impact across the whole workforce because it reinforced preexisting doubts, fears, and insecurities, and it confirmed negative beliefs not only about robots but about technology in general. Every robot became a bad thing despite success in other parts of the company. People wanted to believe that the technology was unnecessary and dangerous and the old way of doing things was the best. After all it had been that way for hundreds of years why change? This incident confirmed their suspicions.

The impact of technology on cultural values

The demonstration was the wrong thing to do. The engineers involved wanted to demonstrate impressive technology, but the audience wanted a demonstration of safety and job security. The engineers thought that the goals being demonstrated addressed important underlying values, in this case efficiency and cost savings. However, the audience was interested in a different set of values: The impact on jobs, the change in responsibilities and training required to be competent on the new technology , and the impact of the new technology on the value of the existing workforce. The engineers involved still truly believe that this was the right solution, and that it could have saved large amounts of money and been very efficient and produced a better product. The organization is still trying to achieve other objectives in addition to efficiency.

ICI Robot Automated Guided Vehicles

Some years after having survived the debacle of the first demonstration of robot technology, the senior author discovered another problem that could be solved with automation, and has hope for the redemption of robot technology:

> Well after the demonstration of the first robot , another opportunity arose. The same plant was having difficulties transferring the same bales of fiber from the baling machines to the

wrapper and to two different storage locations. The existing transportation was via computerized hydraulic hoist. The hydraulics were worn out and the leaking oil contaminated the bales. More importantly, the hoist periodically dropped the bales on the floor and an operator would have to drive the hoist under manual control and pick the bales of the floor and transport them to the wrapper. In the process, the unique identification of the bale would be lost and the automatic labeling system would get out of step and put wrong labels on the bales. This often would cause a major customer relations problem because large batches of clothes would be ruined, resulting in hundreds of thousands of dollars in losses.

Conveyors didn't work very well with this product because fibers would get into the bearings and seize the conveyor in a short time. The hoist system's inherent problems with hydraulic systems and reliability made it a poor choice for replacement.

The best solution from an engineering point of view was a robot truck, called an Automatic Guided Vehicle (AGV). AGVs are used in many warehouse applications and can be designed especially for carrying problem loads. An additional advantage of the AGV in this application was that it could maintain the orientation of the bale out of the box of the baler guaranteeing perfect wrapping and perfect orientation for autoamtic labeling every time.

All we had to do was write a capital expenditure request to get some money and do it, except for the one minor problem that I was not allowed to use the "R" word wiithin 20 miles of my management team!

Past lessons applied

The senior author, presented with an uncommon opportunity to push a favorite but disfavored technology using learnings from his past attempts to overcome the nontechnical barriers to the effort, takes up the quest:

Well this time before I mentioned "robots" I meet with the Unions and the operations folks and identified and confirmed the existing problems and issues. We all agreed that we could not accept the current way of working and we needed a better solution. And, in an almost fully automated plant, the most unpleasant place remaining requiring the operators to interact with the process was the bale handling and labeling area.

During the discussions we reviewed all the options and identified conveyor transportation issues but agreed we would investigate and get quotations for the best available conveyors. After the initial investigation we arranged for sample conveyors to be tested and, as in the past the loose fiber found its way into the works and the conveyors stopped working. The union and the operators, concerned that we were not keeping up to date with progress, demanded that we find other alternatives—surely technology has progressed so that we could move a 300 kg bale of staple fiber a couple of hundred yards and maintain orientation for wrapping and labeling? So my management demanded we come up with other alternatives to the equipment we had so effectively ruled out.

Well this was the opportunity I had been waiting for to introduce the "R" word, but there were cultural issues still outstanding from the first failure. So this time I worked with the manufacturers to identify all the issues and to demonstrate what other industries were doing about them. With the aid of films the team was able to review what other folks do. We showed a frozen food warehouse with extremely cold areas where people did not want to work and how AGVs were being used in that environment. We showed how AGVs interface with and behave around people. We showed how full size automatic forklifts were handling difficult products and how reliable and accurate they were. We demonstrated how they could maintain bale orientation because they followed guided wires buried in the floor. We also presented an estimate of the proposed system which was significantly less than the proposed conveyor system. Finally, we described how we could maintain the trucks by working with the existing maintenance group.

It was the Union that recommended that management make a case to justify this work. The union would support the successful implementation of this system as it was the only sound alternative to a major problem that was now a bottleneck to the whole business. The management decision was not difficult; the whole work force was united behind the AGV solution.

The fully computer controlled system was installed and worked from day one without any problems. It was interesting that the operators took pride in the AGVs from the early days and demonstrated their acceptance by naming them and creating faces on the glass shields of the fork. Each truck had its own personality which the workers identified by watching the way they worked. The final acceptance was expressed in the form of a cartoon where the AGVs were depicted as a team members taking a break and playing dominos, and the plant manager pulling his hair out as he had recently had a campaign to keep people on the job and out of the mess room.

Initial success enables more

The initial success of the AGV system led to extensions. The robot truck system was later fully controlled by a supervisory system which tracked bales and handled grade changes and handover to the warehouse for putting into campaigns for customers delivery. The supervisory control computer was an important part of the system but most people probably did even realize that it existed. The only time people interfaced to the system was so request information about the product and where it was in the plant or to call a truck to do something unusual such as store and retrieve bales with a unique identification during wrapper maintenance or other outages.

What culture lessons were learned

It is not very often that case studies on technology implementation can be found with so much in common: The automation attempts were made in the same company, the same plant, and even in the same area. The cultural issues to be solved were the same. The engineering group was the same. What was different was the approach: Technology was introduced as a solution to a problem that people were having. The solution to the problem was consistent with the values of the people involved. Concerns based upon misimpressions were allayed through early and thorough exposure to the problem and the possible solutions.

Unlike engineers, most people are not interested in technology per se—only how it impacts their way of life. They are more interested in what they will see and how they will deal with problems and failures. Modern industrial plants are introducing more and more technology in every operation and people are often stretched and usually not considered partners or co-workers. The

AGV system demonstrated how people and equipment could work in harmony and be one team. Preparation for change is important and winning sponsors is extremely important.

People do not like to face the unknown they like to see well managed and professionally implemented projects. Having a supervisory system was not important—it was the functionality and how people identified and worked with it that mattered.

Often in automation projects, the need to communicate with the affected users is identified, but as an end in itself instead of as a means to an end. Involving people in automation changes is the only way to understand all of the obstacles to the success of the technology, most of which are not technical in nature.

Formentor

Formentor is a research program conducted in Europe with funding by ESPRIT. Participants include by Cap Gemini Sogeti, the Joint Research Centre of the Commission of the European Communities and Aerospatiale Protection Systems. The goal of this team was to develop a risk management solution for complex systems, in the form of an "intelligent watchdog". Formentor is a decision support system that provides the process operator in charge of process supervision with:

- monitoring,
- situation assessment,
- diagnosis support,
- reactive planning, and
- prediction capabilities

Formentor was intended as an added-value system to the existing control systems. It was to provide a global and permanent overview of the process state, and help react to deteriorated or disturbed operation. It was to allow the anticipation of severe process states by detection of precursor signs and malfunctions, and measure the impact of actions on the process behavior, thus limiting plant shutdowns and improving quality of the products.

The Formentor team have completed two industrial prototypes: one at the BP Grangemouth chemical plant in Scotland, and the other at the Total Refinery in LeHavre, France. Both of the systems have met many of their technical goals.

The impact of the technology life cycle

Both of the prototype applications have been abandoned by their industrial sponsors due to the high cost and required knowledge skill set to maintain the system, and the lack of benefit achieved during actual use.

Formentor systems need to reflect the actual plant in order to provide technically accurate information. The system is initially accurate, but plants are always being maintained, upgraded, or expanded, and small changes in the plant equipment or the process operation that occur over time require specialist computer experts to make corresponding changes to Formentor. This turned out to require resources beyond the ability of the plants to provide.

The model-based approach used by Formentor (and some other efforts with similar objectives) is technically very promising. However, the same monolithic approach that enables good diagnoses of problems limits the accuracy if the models are not accurate. Thus the approach may require more maintenance effort in practice than can be expected.

Multiple applications for each problem

Another issue with monolithic approaches is that not all problems are best diagnosed using the same approaches. For subtle imbalances in processes, the performance of good models is hard to beat. However, other kinds of problems may be better diagnosed with empirical (e.g., statistical or neural network), knowledge-based, qualitative, or fuzzy techniques. Aside from diagnostic accuracy, an advantage to using multiple techniques is that the installed system can be developed incrementally—the user gets benefits from each new diagnostic capability that is added—which helps get over barriers raised by the large initial cost to install a comprehensive system.

The value of current practice

Had Formentor provided significant benefit in actual operations, maintenance costs might not have been a significant issue, but it turns out that Formentor was not used during process upsets as much as it could have been. This was due to the fact that the Formentor display was not incorporated within the system used by operators, but on a separate, independent system. In the event of an upset, this required operators to leave their primary view of the process—and their only means of process control. History has repeatedly demonstrated that operators will not do this regardless of the extra information that may be available. The robot examples showed the need to take into consideration the users' practice of values, but we also need to take into consideration the value of practice.

For a diagnostic or decision support tool to work in the real world, it must seamlessly fit into the way people work with the existing control system. More generally, for any improvement to an existing system to succeed, it must accommodate the ways users work with the existing system.

The Advanced Automation System for air traffic control

Beginning in 1984, the FAA attempted to replace large parts of its obsolescent air traffic control system with better technology. After a rigorous design competition, a $2.5 Billion contract was issued in 1988 calling for initial deliveries in 1992. The schedule quickly slipped and costs escalated. At the peak of the program in 1992, more than 1000 software engineers were working on the system, and the estimated cost to complete the system had risen to $7.6 Billion. With no end in sight, most of the program was canceled in 1994.

Lessons revisited

Every one of the issues raised in our examples so far played a role here as well. There was inadequate attention to configuration management. There was insufficient attention paid to the culture of air traffic controllers and to their operations practices [In fact, at least some informed observers lay the ultimate failure of the program to intractable problems associated with automating the racks of flight progress strips that controllers physically manipulate to help them stay oriented to the progress of flights in their airspace.] Inadequate attention was paid to lifecycle costs, and the solution being built was too monolithic [Calling for the consolidation of 200 air traffic control centers into 20, for example].

The system also serves as an example of two new issues that frequently arise in the evolution of complex systems, which will be encountered more and more frequently, and which may represent the largest challenge in the future: The program attempted to apply the wrong kind of technology, and it did not address the right problem, in the first place.

New technology for new solutions?

In 1984, the Advanced Automation System attempted to modernize air traffic control through the development of a system using newly developed parallel-processing UNIX workstations running hundreds of thousands of lines of Ada code and with a seven nines availability requirement (allowable downtime: 3 seconds per year). To the best of our knowledge, none of those elements have yet been commercially demonstrated, five years after the system was supposed to have been fielded. This was not a system depending on well tested off the shelf components!

Conversely, petrochemical companies, manufacturing industries, and even NASA often have the opposite problem: It is not uncommon to find key control software running on the equivalent of IBM XT technology—if it is computerized at all. What is lost in capability is more than made up for in reliability: A desktop PC might crash once or twice a week (or even once or twice a day); a control PC is not expected to crash during years of continuous operation. In these industries, even upgraded technology that seems old by desktop standards.

The challenge to those introducing automation is to select technology that is new enough to be useful, but old enough for the risks be understood.

What is the problem here?

The FAA attempted to better automate the existing air traffic control system, and in particular to make it better at doing the things it already did. Thus, designers went after throughputs, and planes per controller, and other measurements of the system currently in use.

Fundamental goals—the reason the system exists in the first place—were given short shrift. Thus, the problem left unaddressed was something like, "How can we have more planes in the same airspace with more safety and less cost?".

It would be unfair to hold up this particular project to ridicule as a particularly bad example, because the issue is endemic, and very difficult to overcome. Most of engineering is incremental—we attempt to do the same things, only better. The engineers at ICI did not think much if at all about whether there was a better way to move fiber to customers than in bundled and labeled bales; they instead tried to bundle, label, and move bales better.

The real advances—the breakthroughs in the application of technology—come not from solving the same problems in better ways, but from solving new problems in new ways. All of us, even the most inventive, have to watch out for this one. [To take just one example, the Wright Brothers spent years trying to improve the aerodynamics of their front-mounted horizontal stabilizer, even after their competitors had put it in back.]

At this point in the history of control systems, it may be time to look away from efforts to gain another .5% of efficiency, and instead focus on something that can give us 10%.