[Human] Supervisory Control and Decision Support: State of the Art

Ted Cochran
Honeywell Technology Center
3660 Technology Drive, Minneapolis, MN 55409

Duncan Rowan
DuPont Engineering - Process Control Initiative
DuPont Co - N6545A, 1007 Market St., Wilmington, DE 19898

Abstract

We review the state of the art in technology available for decision support, as well as the extent to which that technology has been applied in process control. We conclude that user interface hardware is less than state of the art due to the requirements of the domain, and the introduction of some user interface software technologies has been delayed as a result. On the other hand, decision support technology has been given a thorough trial, and process control applications have sometimes advanced the state of the art. Current research into the development of environments to support collaborative interaction of multiple operators with multiple tools may be the key to more cost-effective use of available technology.

Keywords

Process control, Decision support, Expert systems, On-line aiding, Supervisory control, User interface.

Introduction

This year?1995?may well be the year selected by future historians to mark the beginning of the information age. Last year, for the first time in history, more was spent by consumers on new computers than on television sets, and it is only a matter of months before computer sales exceed television sales in raw numbers, as well. The world-wide web is growing at 10-20% per month, and press coverage is ubiquitous. Email addresses and URLs are making appearances in popular publications and in television and print advertising?they are no longer confined to Wired. The rate of improvement to desktop computer power is accelerating: Performance is no longer doubling every three to four years, but every two to three. Networks are being installed which are capable of delivering more information, more rapidly, than can most hard disks. CD-ROMS are now standard accessories, pre-installed on the majority of new personal computers.

So, what heights have these developments enabled supervisory control of process operations to achieve?

In this paper we'll review the state of the art in supervisory control of complex processes. We'll begin by describing the roles academic and industry researchers are playing in this area. We'll discuss the state of the art for some of the key generic technologies in terms of recent research achievements, what's available commercially, and what's typical in current process control systems.

We'll then describe the current state of process-specific supervisory control and decision support technologies, offer an opinion about what's still needed, and assess the prospects for getting there in the near future.

Technology Development

Academic research, and research directed at technology development by private enterprise, tends to define the leading edge of the state of the art in most technology-driven fields. Petrochemical companies have historically been conservative about connecting the latest bells and whistles to their processes at low levels (a reliability track record is crucial, and takes time to attain), but have adopted relevant new technologies relatively quickly at the supervisory control level when cost-effective applications are evident.

Applied Research

The objectives of industrial research organizations are now more focused on core technologies essential to the business. These groups are now focused on extending process understanding and driving the development of new technologies by vendors. Special attention is being devoted to the development of process models (fundamental and empirical, steady state and dynamic), on new control technologies (e.g., non-linear model predictive control), and on control strategy improvements. Models are often developed using commercially available tools such as Aspen Technology's SPEEDUP product, or Honeywell SACDA's TRAINER. The application of mature technologies such as decision support are no longer receiving extensive attention from corporate research organizations; instead, these are used by corporate and plant engineering groups and by some groups in Information Systems organizations.

User interfaces

The state of the art in interface hardware includes the 9000 x 3000 pixel (25 ft. x 9 ft.) composite display system assembled by the Electric Power Research Institute (EPRI) to simulate mimic boards (Fray, 1995), and by virtual scenes generated by head mounted displays (stereo color 1280 x 1024 displays throwing 60K polygons/sec). Technologies such as gesture recognition ("put that there"), unconstrained speech input and output, handwriting, and tactile feedback (gloves, manipulators with force feedback) are at varying stages of maturity, but all have been developed sufficiently to enable practical application in at least some domains. User interface "software" technologies include virtual reality, synthetic vision (e.g., realistic displays created by combining terrain models with position and attitude information and projected for pilots of windowless aircraft), augmented displays (e.g., infrared data overlaid on the users' existing view of the world), and new navigation, data visualization, intent recognition, and user modeling approaches.

Commercially available hardware includes displays with resolutions up to 2048 x 2048 pixels at 24 bits per pixel, integrated multimedia (for example, up to four simultaneous real-time 640 x 480 video windows, stereo CD-quality audio), constrained speech I/O, rudimentary handwriting recognition, and 100 mb/s networking (28 kb/s wireless). Commercially-available user interface software technology includes simple intent recognition (e.g., Microsoft Excel's assistant), low resolution virtual reality, first generation groupware (Notes), third-generation user interface building tools (Visual Basic, TAE), and second-generation graphical user interface technologies (Quickdraw 3D and VR, Scripting, Agents). The performance of such systems is best measured from the user's perspective (we assert that CPU speed is irrelevant). Users typically demand response times of less than one second for frequent operations and will tolerate 60 second delays for e.g., file operations (save, paginate, recalc). Reliability is also important; typical software has a MTBF (crashes) on the order of hours, and a hardware MTBF (system needs repairs) on the order of months.

Typical process control applications incorporate user interface hardware that is much less capable, but much more reliable, than that described above. Displays of 640 x 480 pixels with up to 256 colors are just beginning to give way to 1280 x 1024 x 16 bit displays. Systems can, however, be configured to support many such displays simultaneously. Multimedia use is minimal (sound = horn).

Software is highly customized: X-windows and other open environments are becoming available, but are being adopted only very slowly for critical operations. Graphical user interfaces lag those available elsewhere.

Reliability, however, is robust: Software MTBF is measured in months; hardware MTBF is often measured in years. The user-visible performance is also adequate: <1 sec scan rates, <1 sec updates, <4 sec screen call ups.

We conclude that while there are significant differences between the state of the art and current practice, they are not unexpected given the relatively small market, rigorous requirements, different nature of work in process control domains.

Decision support

A working definition of decision support systems is technology that presents data to users in ways that enable easier detection of significant patterns and easier, more rapid selection and execution of appropriate actions. Technologies that comprise and contribute to decision support systems include knowledge-based systems, model-based systems, empirical approaches (e.g., statistical process control and neural networks), and engineering process control (Davis, 1995).

Decision support can be augmented via on-line documentation, procedural advisors, scheduling assistants, and other technologies. Our overview of decision support in process control distinguishes between state of the available technology, which may include demonstrations of applications in restricted domains, and the typical application of that technology to the supervisory control domain.

Expert Systems

The promise of expert systems in the mid-80's was to insert the best operational and decision-making capability into a computer program that would then continuously give that advice on-line. In this way, the best knowledge about the process and the process operation would be available at all times, even to less experienced operating personnel. This promise has largely been left unfulfilled.

Corporate downsizing has led to the reduction or elimination of support for large expert system projects, due to lack of capital and engineering resources. Available capital and people have been focused on very basic opportunities using conventional technologies. Management focus is on quick pay back applications.

The concepts of intelligent systems and expert systems are being applied in small focused applications. There have been very few "expert-in-a-box" applications. In many cases, the expertise is less well understood than initially thought and may be disbursed among various people in many groups. It can be difficult for these people to get commitment to work on a large expert system project.

The current cost of abnormal situations?at least $20B to the U.S. economy as a whole?is mostly derivative; that is, it does not directly impact the process industry, and the costs that do represent a relatively small proportion of their total operating budgets. Significant incidents occur infrequently, in which case the somewhat reduced process utility does not significantly affect the economics of the operation. Only in sold-out conditions do short-term outages have a major impact on economics [However, the continuing consolidation of the industry and increasing plant utilization are leading to such impacts more frequently.] A more serious economic consequence is degraded operation, which can produce off-specification product with attendant recycling or waste disposal cost. Even then, using cost reduction as a justification for investment is difficult: Competing opportunities leading to increased revenue are typically viewed more favorably. Even when cost-effectiveness is proven, expert systems must also compete with other approaches for dealing with detection and correction of abnormal operation. Chronic or serious abnormal operation may be addressed by redesigning the process or process equipment to eliminate the possibility of occurrence. In some cases, this more cost effective than developing an expert system to detect or predict such an occurrence.

Most current expert system applications therefore address specific, high pay back situations, where degraded performance can be detected early, and prior to the process becoming out of specification. These small process "watchdogs" are being embedded in some distributed control systems (DCS) as well as in higher level software environments such as host computer systems or Gensym's G2 tool. The advantage of the higher level network-based platforms is the ability of the applications to communicate information via electronic mail and provide report information which is easily accessible to users.

More complex diagnostic systems have been developed by Davis (1991) of Ohio State University in collaboration with a number of petrochemical companies. In DuPont, more complex expert system applications have been targeted for process transition management (Rowan, 1992) and complex diagnostic advisory systems. One of these diagnostic systems is a control loop performance monitor which tracks several hundred loops and reports any deviation from "normal" performance statistics. The application has proven valuable in detecting several instances of loop instability caused by marginal tuning, transmitter degradation, and blocked impulse lines; in one case, using the statistics generated by the application, an engineer gained a key insight into process equipment. These results would not have been possible without this sophisticated loop monitoring. The techniques in this monitor are similar to those reported by Jofriet (1994).

Another promise of expert systems technology was that once applications were built for one location, they could be easily leveraged to other locations with the same process, therefore distributing the best knowledge throughout the business. In practice, this has been difficult. Cultural barriers exist as these sites tend to be operated independently. Also, due to control/computer systems infrastructure being different from site to site, these applications must be redesigned for a new site. This is a significant cost, but some companies have nevertheless successfully leveraged applications. Air Products & Chemicals, Inc. claims to have broadly leveraged their expert system applications to their industrial gas plants, which are similar in design (and minimally staffed), and therefore a particularly good fit for these applications.

Though large expert systems applications are infrequent, the technology is recognized as a useful component of other types of systems. For example, in new energy management technology that DuPont is evaluating, expert systems capability is needed to provide equipment diagnosis prior to optimization. Expert systems technology has also been found to be useful in other types of optimization applications (Faccenda, 1995). Finally, this capability is necessary to support robust performance in advanced model predictive control (MPC) applications.

Modeling

Model-based diagnostic and optimization techniques compare plant data with a model of a plant and identify disparities of potential interest. Some, like APACS (Benjamin, 1990), attempt to identify potential causes of abnormal situations by manipulating the model until better agreement with plant data is achieved (and then announcing the manipulation that most economically explains the plant data). Others, such as Formentor (Pennings & Saussais, 1993), use knowledge-based approaches to identify discrepancies. Strictly empirical methods use neural networks, fuzzy logic, or statistical methods to categorize plant data as normal or abnormal, and, if abnormal, to further determine what the problem is.

Process models can be used to study process behavior and therefore used as an analysis tool during development of decision-support systems. These models tend to be dynamic, mechanistic models based on first principles. Empirical or data models can also be used to detect and announce abnormal patterns in process data. These models are developed from existing process data containing the patterns which must be detected.

An alternate approach to decision support technology is to use technology to train operators to themselves detect abnormal situations and react appropriately. There is much interest at DuPont in improving operator skills and abilities. One plant is adding a fifth shift so that operators are in training as much as one week per month. A number of projects have been undertaken to create small dynamic simulations of critical processes (including abnormal situations). These simulations can then be connected to the plant DCS and operators can be trained in a "flight simulator" mode. This has been particularly effective for the startup of new processes and the training of inexperienced operators (Bauer, 1995; Hawthorne, 1995). Operators who are confident that they can recognize and correct abnormal situations are a great asset.

Empirical approaches

Empirical approaches are also useful as decision support technology (Piovoso, 1991). Empirical models are linear or nonlinear correlation models derived from actual process data. These models include multiple linear regression, partially least squares, principal component regression, neural nets, and time-series models. There is a great deal of interest in industry in using existing process physical measurements to supply models which can then predict process parameters or product properties. In some cases, such as with neural nets, the "model" thus derived is the quintessential black box: It generates correct outputs for specific inputs, but it does so using techniques that do not map well onto human understanding of such processes.

Generally, a large amount of historical process data is available from which to develop these models. In practice, however, often this historical data is insufficient to make accurate models. Often, the process must be disturbed with step or pulse testing in order to have sufficient movement in all input variables in order to identify the model structure. These tests generally must be agreed upon by manufacturing management and be scheduled in advance. Therefore, these applications take more effort than might be initially expected. Nonetheless, there is a great deal of interest in generating these models because of the high perceived value.

As industry moves to on-line product release capability, these types of models will be essential in determining that the process is within specification on a continuous basis. In order to cope with the complexities of real process data, many of these models tend to be small multiple input, single output models. In order to capture process dynamics for model identification and control, time series techniques are frequently used.

Statistical/Engineering Process Control

Cooperation between Control Engineers and Statisticians over the last few years has led to better control techniques for regulating continuous process variables with infrequent laboratory sampled measurements (MacGregor, 1988). The exponentially weighted moving average (EWMA) filter has been shown to be superior to other statistical filters where regulatory action can be taken without a cost penalty. Because of the low level of complexity, these applications are of wide interest within DuPont. Also, Tennessee Eastman has reported a great deal of activity in this area (Paulonis, 1995). The EWMA-based controller can substitute for a Cusum or Shewart controller. Reductions in process variability for 30 to 60 percent have been seen.

Application of Future Technology

Networks

Computer and networking hardware and software are advancing rapidly and will add new dimensions to decision support system applications. One DuPont business is porting an existing batch scheduling, ingredient monitoring application to a client/server architecture. The user interface will be revised from a character cell terminal to a Windows-based package. Windows-based user interface will allow for much easier data access and application navigation. Many of these legacy business support applications will move towards client/server and Windows interface technologies. Software tools such as VISUAL BASIC and VISUAL C++ make it feasible for corporate Information Systems programmers to develop Windows applications. For real-time applications, software environment such as Gensym's G2 allows for much more rapid software development. These tools also have the advantage of running across various hardware platforms and operating systems. Therefore, when computer hardware architecture changes, the application software will not have to be ported; this is a major cost savings.

Having real-time process data available on inexpensive PC packages allows for other inexpensive software packages to readily access the data. Applications such as spreadsheets, statistical packages, modeling and math packages are very cost effective tools.

Network technology is allowing process data and information to move over wide area networks. This is allowing for global distribution of decision support systems. Data can be pulled across the network, processed in a central computer system and the results can be returned. Current network bandwidth is sufficient for monitoring relatively slowly changing chemical processes. This allows for minimal hardware/software and application support at the local site and allows a central support staff to handle most of the work. Thus economies-of-scale are achieved via this architecture. This will be a huge advantage for global companies with sites operating in relatively remote low technology areas.

Collaborative systems

As the capabilities of decision support techniques improve, the fundamental principles of user-system collaboration are receiving renewed attention. In some cases, the demands of industry applications are driving the development of theoretical knowledge in this area, blurring the line between academic and applied research (Cochran, 1994). It is now apparent that the success of decision aiding systems does not depend merely upon their accuracy, but on their ability to accommodate specific and sometimes unrecognized user needs.

For example, human experts may be ignored by an operator in critical situations due to their unfamiliarity, their interaction style, or because the operator can simply not afford the additional effort required to pay attention. Decision support systems have fallen victim to the same fate: If they are in any way misunderstood or mistrusted by users, or even if they are merely inscrutable, they tend not to be relied upon when they are really needed. Esprit's Gradient project and PRECARN's IGI project are both attempting to address this specific issue.

Even if systems are mostly correct, and mostly paid attention to, other problems emerge: If operators grow to depend upon decision aids, such systems may lead to a decrease in the alertness of operators in the short term, and to a decrease in their proficiency in the long term. This problem has received significant attention in the aerospace industry with the introduction of glass cockpits and automated flight management systems; DCS vendors are applying the knowledge thereby gained?sometimes at great cost?to industrial DCSs as well.

The NIST-funded Abnormal Situation Management project, led by Honeywell, is a particularly ambitious effort in this area: It aims to provide an environment for user system collaboration, which specifies standards for the interaction of diagnosis, state estimation, scheduling, planning, and operator support systems. Such systems, once installed, will become part of a collaborative environment in which the expectations and needs of operators are explicitly supported during abnormal situations. If successful, this research project will prove the feasibility of combining the many current approaches to decision aiding in a comprehensive and effective way.

Conclusions

The state of the art of user interface hardware in the process industries is well behind that which is available in other markets, which is unsurprising given the extreme reliability demands placed on process-connected equipment. As a result, user interface displays and some other forms of user interface software are also lagging.

Nevertheless, the application in the process industry of decision support technology as a whole is not only at the state of the art, but is driving it, due to the same factors that drive the hardware reliability criteria?process safety, the criticality of the operator's role in the management of the process, the enormous capital investment in the process equipment and the materials being processed, and overall economic competitiveness. These factors have led to a continual search for better ways to support operations personnel, backed in some cases by significant, shared, investment in new approaches.

If these approaches pan out, the industry may finally be available to develop some comprehensive and cost-effective decision aiding systems which integrate a wide variety of the successful, but isolated, efforts to date.

References

Benjamin, M., Chou, Q., Mensah, S., and Mylopoulos, J. (1990). Expert systems for advanced process analysis and control of nuclear power plants. IAEA Specialists Meeting on Analysis and Experience in Control and Instrumentation as a Decision Tool, Arnhem.

Bauer, E. J., and J. T. Stainback (1995). History of an operator training simulator on a new polymer line project, AICHE 1995 Spring National Meeting, Houston.

Cochran, E. (1994) Basic versus applied research in cognitive science: A view from industry. Ecological Psychology, 6, pp. 131-136.

Davis, J.F. (1995). Intelligent support systems in the process industries: The critical importance of structure in large-scale systems. Abnormal Situation Management Consortium Spring Meeting, Cleveland.

Davis, J.F. (1991). A Framework for implementing on-line diagnostic advisory systems in continuous process operations, preprint, AICHE 1991 Fall National Meeting.

Faccenda, J.F. (1995). In C. E. Bodington (Ed.), Planning, scheduling, and control integration in the process industries, McGraw-Hill, N.Y. pp. 201-209.

Fray, R. (1995) Compact simulators for fossil-fueled power plants. IEEE Spectrum, 32, pp. 46-51.

Hawthorne, R., T. Fu, B. Tyreus, and B. Donovan (1995). Process dynamic simulation at DuPont's Corpus Christi plant, AICHE 1995 Spring National Meeting, Houston.

Jofriet, P., M. Harvey, C. Seppala, B Surgenor, T. Harris (1994). A G2 knowledge base for control loop performance analysis, 1994 Gensym Users Meeting, Cambridge, MA.

MacGregor, J. F. (1988). On-line statistical process control. Chemical Eng. Progress, 10, pp 21-31.

Paulonis, M. A. (1995), Closed-loop control of properties measured off-line, AICHE 1995 Spring National Meeting, Houston.

Pennings, R. & Y. Saussais. (1993). Formentor: Real-time safety oriented decision support using a goal tree-success tree model. Third International Conference on Artificial Intelligence, Avignon.

Piovoso, M. J., and A. J. Owens (1991). Sensor data analysis using artificial neural networks. In Arkun and Ray (Eds.), , Chemical Process Control CPCIV, AICHE, NY. pp. 101-118.
Rowan, D.A. (1992). Beyond FALCON: Industrial applications of knowledge-based Systems, IFAC Symposium: On-line Fault Detection and Supervision in the Chemical Process Industries, Newark, DE.