MISSION SAFETY CRITICAL SYSTEMS
RESEARCH & APPLICATIONS
ricis
SYMPOSIUM '92

EXPERT SYSTEMS IN THE PROCESS INDUSTRIES

DR. G. M. Stanley

ABSTRACT

This paper gives an overview of industrial applications of real-time knowledge based expert systems (KBESs) in the process industries. After a brief overview of the features of a KBES useful in process applications, the general roles of KBESs are covered. A particular focus is diagnostic applications, one of the major application areas. Many applications are seen as an expansion of supervisory control. The lessons learned from numerous online applications are summarized.

BIOGRAPHY

Dr. Stanley, Principal Scientist at Gensym Corp., joined Gensym in 1987. He developed methodologies for practical application of real-time expert systems. He worked on and led several major diagnostics projects in the nuclear industry and in environmental controls. As Director of Applications Development, he led the development of new products such as the Diagnostic Assistant. He is the author of many technical papers in expert systems applications, data reconciliation, Kalman filtering and estimation theory, and simulation. He worked at Exxon Chemical from 1976-1987, holding key technical and management positions in process control, process engineering, optimization, expert systems, dynamic simulation, and information systems. Dr. Stanley completed his Chemical Engineering Ph.D degree at Northwestern University in 1976. His work in estimation and control advanced data reconciliation, Kalman filtering, and fault diagnosis, especially for network problems. He obtained his MS Ch.E. from Northwestern University, with a thesis on adaptive control. His BS ChE. was from Purdue University in 1971.

EXPERT SYSTEMS IN THE PROCESS INDUSTRIES

G.M. Stanley

Principal Scientist
Gensym Corporation - Southern Region, 10077 Grogan's Mill Road, Suite 100, The Woodlands, TX 77380

ABSTRACT

This paper gives an overview of industrial applications of real-time knowledge based expert systems (KBESs) in the process industries. After a brief overview of the features of a KBES useful in process applications, the general roles of KBESs are covered. A particular focus is diagnostic applications, one of the major application areas. Many applications are seen as an expansion of supervisory control. Finally, the lessons learned from numerous online applications are summarized.

BACKGROUND

Knowledge-based systems overview
Artificial Intelligence (AI) techniques include rule-based expert systems and object-oriented systems. The emphasis is declarative representation: separating the description (the knowledge) of a process, from the subsequent analysis of that knowledge by an inference engine. The knowledge is thus made more explicit, visible, and analyzable, instead of being hidden inside of procedural programming code. The knowledge is built as much as possible to be independent of the immediate application. Good expert system tools are generally based on an object-oriented paradigm, and we call them knowledge-based (KB) systems, or knowledge-based expert systems (KBESs).

A more detailed review of KBES applications in the process industries, with an emphasis on process control and a large number of references, is given by Stanley(1991). Descriptions of the features needed in a KBES for real-time control are given by Rowan(1989), Moore and others (1988), Moore, Stanley & Rosenof(1990), and Hoffman, Stanley & Hawkinson(1989).

A KBES in the process industry is often an extension of supervisory process control. Control technology generally emphasizes quantitative processing, while KBESs integrate both qualitative and quantitative processing. A KBES provides a general framework for integrating technologies as diverse as control design and operation, neural nets, rule-based systems, symbolic cause/effect models, logic networks, differential equation solving, and scheduling algorithms.

Some features of Knowledge-Based Expert Systems useful for online systems
Current online industrial applications are generally built within shells, which package a combination of tools. Different KBES shells may include some of the following features useful for online control applications:
• objects with attributes
• class hierarchy for objects, with inheritance of properties and behavior
• associative knowledge, relating objects in the form of connections and relations
• structural knowledge (e.g., "part-of" relation)
• representation and manipulation of objects and connections graphically
• rules and associated inference engine
• procedures
• analytic knowledge, such as functions, formulas, and differential equation simulation
• real-time features such as a task scheduler for concurrent operations, time stamping and validity intervals for variables, history-keeping, and data interfaces
• interactive development and run-time environment

Not all shells contain all these features. This paper is based mainly on experiences of users of G2, a real-time KBES shell which does include all these features.

The emphasis in a KBES is in building up descriptions, or knowledge, independent of the subsequent use of that knowledge in multiple applications. For instance, the developer specifies the types of objects in the plant, and specifies conditions which might correspond to a fault. The easy buildup of this declarative knowledge, combined with the available graphical interfaces, encourages a rapid prototyping and iterative refinement approach to software development.

Users often use a graphics-oriented KBESs to create a graphical language by defining the behaviors of objects and connections. For instance, a system based on AND and OR gates is a typical graphical language. Continuous control system engineers generally think in terms of data flow languages consisting of processing blocks and signals. Another common approach is to define objects representing actions, connected by directed arcs specifying sequential or concurrent execution. The GDA (Gensym Diagnostic Assistant) product, built using G2, is a complete graphical language encompassing both data flow (filters, AND & OR gates, etc.) and sequential control

In general, users of KBESs are representing almost everything as objects. It fits well with the way they think.

GENERAL ROLES OF A KBES IN THE PROCESS INDUSTRIES

An overview of KBES applications in the process industries, with an emphasis on process control, is given by Stanley(1991). That paper includes a number of case studies, and numerous references to successful applications illustrating the summarized points made in this paper. Some roles of expert systems in process control have been outlined by Stephanopolous (1990) and Årzén(1990). An overview of some current and expected applications is given by Rehbein and others (1990). Rosenof (1990) has summarized some roles for KBESs in batch process automation. Many of the online applications span more than one of the areas defined below, exploiting the usefulness of a KBES as a general framework:

The following are proven successful application areas for a KBES:

- Fault diagnosis: early detection, root cause analysis/alarm filtering, repetitive problem recognition, test planning, alarm management
- Supervisory control
  - Complex control schemes
  - Recovery from extreme conditions
  - Emergency shutdown
  - Heuristic optimization, e.g., debottlenecking
  - Startup or shutdown monitoring
  - Batch phase transition detection and subsequent control mode switching
  - Process and control performance monitoring
- Statistical Process Control (SPC) and subsequent assignment of causes
- Real Time Quality Management (combination of the above)
- Online "smart" operator and troubleshooting manual
- Sequential or batch control
- Control system validation
- Object-oriented simulation of processes and control systems

The following KBES application areas are actively being developed and tested by industry:

- Scheduling
- Operator training, with real-time simulation
- Tank farm management
- Formalizing compliance with ISO-9000 (quality), government, or other standards

Some evolving and future roles of a KBES in the process industries:
- Predictive maintenance
- Process validation
- Intelligent supervision of adaptive control, model identification, parameter estimation, state estimation, data reconciliation, optimization, neural network training & run-time coordination
- Automated design of control systems (and implementation)

Economic justification for a KBES in the process industries
After some early experimentation, applications are now generally justified based on economics as well as safety. Some companies have published information quantifying substantial benefits. For instance, DuPont has stated that they "routinely" see returns on investment as high as 10 to 1 (Rehbein and others, 1990; Rowan, 1989). Monsanto's evaluation of its first online system showed benefits of $250K/year (Mertz, 1990; Spang Robinson Report, 1989, Rehbein and others, 1990). Other examples can be found in the review paper by Stanley (1991). The justifications for a typical application such as diagnosis, include:
- Safety
- Real-time quality management/quality control
  - Early problem detection with multiple variables
  - Determining the assignable causes of problems
- Yield/production
- Loss prevention
- Equipment protection
- Environmental Protection
- Sensor/model validation for successful plant optimization
- Formalizing compliance with ISO-9000 (quality), government, or other standards

GENERAL LESSONS LEARNED

Real-time KBESs are robust enough to have succeeded in numerous applications, including closed loop control.
Some current systems operate in a "closed loop" mode, manipulating valves or controller setpoints. Many of the current systems already occupy a "grey" area between open and closed loop control: the control goal is closed-loop, but an operator is in the feedback loop, routinely approving the recommendations. Many of the open-loop applications will migrate to closed loop, as people build up confidence.

Significant benefits are derived in areas complementary to conventional controls, such as diagnosis, quality management, and abnormal operation
Significant benefits have been achieved. Many of the credits are in the same areas as good process control, e.g., process repeatability, quality improvement, achieving best demonstratable operation, shorter batch time, lower waste or energy costs, and avoidance of accidents. However, the reasons for the benefits often complement those of process control, since they are often derived during periods of unusual operation, or from better planning of the normal control operations. KBES-based diagnostics are needed as part of the overall control system to catch major problems such as sensor failure, and then disable the fragile, "normal" control systems which only handle normal operations. Quality problems can be thought of as faults -- they are economic faults, just less severe than safety problems. Diagnostic techniques typically also are used in batch control systems to detect or plan the transition from one operating phase to another.

Normal models and controls break down during the extreme operation. There, the more effective models or actions under extreme conditions are likely to be simpler, based on first principles or on heuristics. These alternate controls are easier to build in a KBES than in conventional systems.

KBESs complement SPC techniques by earlier problem detection and determining assignable causes, achieving Real-Time Quality Management
SPC tools are sensitive detectors of problems. However, they offer little or no guidance as to the root causes (assignable causes) of problems, or how to correct the problems. This is a fundamental limitation, because standard SPC techniques do not capture process model knowledge and use it. A KBES can apply SPC to detect problems, and then pinpoint the cause of the problems. Thus, the broader problem of "maintaining product quality" can be addressed through a combination of SPC techniques, diagnostic techniques, and conventional control systems during normal operation. This broader approach to "Real Time Quality Management" has been successfully applied by DuPont and others.

Pure SPC systems also require the users to wait until faults have propagated and repeatedly caused off-spec products. By building in process knowledge, faults can be detected and corrected long before SPC techniques recognize a product problem. Diagnostic techniques implemented in a KBES can use SPC techniques as sensitive detectors of problems, but also provide a broader framework for building in the knowledge to determine the causes of problems and correct them. For example, if a valve sticks, a reflux drum may empty, ending reflux to a distillation tower, causing product to go out of specification half an hour after the fault, with detection at an online analyzer within another 10 minutes, and confirmation from the laboratory in 2 hours. Diagnostics could detect the stuck valve in less than a minute.

KBESs return significant economic benefits
The systems can now be justified for economic reasons, in addition to safety.

Significant benefits are derived from productivity in development
While the earliest expert systems were major efforts, a graphics-oriented real-time KBES now can significantly shorten development time vs. conventional coding. The ability to rapidly prototype and get user input is a major benefit. While any of these systems could be implemented in conventional code, it would be difficult, more time-consuming and error-prone, and harder to maintain.

KBESs reduce the gaps between specification, implementation, and run-time
KBESs encourage declarative representation of the information needed for design of a system, such as objects with attributes which are used to build models. The process schematic itself is part of the design basis, and can be used directly at run-time. The design procedure itself can be automated. For instance, goals and subgoals can be represented as objects suitable for deriving control strategies. Domain-specific heuristics on selection of controlled and manipulated variables can be explicitly represented as rules or objects.

In an integrated package, many of the objects (such as the process schematic) used by the designer can be used by the end user. Status indications via color are useful to both the designer and end-user. A programmer separate from the designer and end user is generally not needed. A separate software package for design and run-time use are not needed.

Maintainability is a major issue
Early custom-coded systems were not maintainable, and are no longer used. Maintenance is a major issue at plants, because they are always being modified, and related computer systems need to evolve with it. Modern KBES shells provide a better framework. Systems must be changeable in a natural way by the users, not just AI developers.

System integration is a major issue
A significant portion of the overall effort is in systems integration. Tools which build in extensive support for real-time data interfacing save significant development effort.

Graphics-oriented KBESs are an integrating technology
Due to their high-level ability to represent, manipulate, and display knowledge in various forms, graphics-oriented KBESs can be used as a tool to integrate other techniques. One KB representation

can be used for multiple purposes. Work is under way at various locations using a KBES to integrate such diverse technologies as neural networks, fault trees, databases, and expert system rules.

A KBES can fill the "CIM gap" between process control and planning & scheduling. For instance, once the KBES has a representation of the plant schematic, the recipes, and the processing sequence and estimated processing times, that same representation can be used both for planning purposes, and then to carry out the sequential control. The key is that the plant and product information is represented in a way independent of the application. In a continuous plant, a hybrid system can decide when it is time to do an emergency shutdowns, and carry out the shutdown. In a batch process, the hybrid system can detect the end of one phase of operation, and switch control schemes for the next phase of operation.

KBESs specialized for real-time use are needed for process control applications
Earlier attempts to extend the traditional static expert system shells, or to code a system from the beginning, were generally interesting learning experiences. These mostly ineffective attempts were generally driven periodically by batches of data placed in files.

However, for the dynamic industrial environment, these approaches generally proved too slow, too difficult to be economically justified or maintainable, and often too unreliable. A specialized real-time KBES uses an asynchronous processing model for data acquisition and task execution within the expert system. The necessary features for history-keeping, time stamping, and so on, are provided. Also, early LISP-based systems, without special memory-management provisions to prevent garbage collection, could suffer seemingly-random pauses during garbage collection (memory reclamation), unacceptable for real-time operation. A real-time KBES should not garbage collect at run time.

LESSONS LEARNED IN KNOWLEDGE REPRESENTATION

Graphical specification of knowledge is effective
Users like developing graphical problem-specification languages. In many cases, users have defined their own graphical languages, using blocks to represent numerical or logical operations, and to represent sequential action steps. Based on these experiences, Gensym has built the GDA product already mentioned. Many users have made extensive use of the information available directly from process schematics built out of objects with connections between them.

Generic knowledge libraries shorten development time
Many users are building libraries which can be reapplied at different sites, based on analyzing a process schematic. This is especially applicable in diagnostics, where low-level failures in valves and sensors are essentially the same in all plants. This results in rapid transfer of technology and development, uniformity, and maintainability. The knowledge libraries speed applications at the first site as well, because much of the configuration for the entire site is for repetitive elements such as valves and controllers.

Symbolic and numerical filtering, and evidence combination techniques for managing noise and uncertainty are important
Event and trend detection, with their associated upstream models and filtering, provide the interface between the continuous, external world, and the higher-level, usually symbolic states in the knowledge based system. Noisy signals can lead to transient false conclusions, obscure the true conclusions, and lead to excessive forward chaining.

To reduce the impact of noise, you can filter heavily and accept the resulting lags in many cases, since a fast feedback loop is not within the application. Nonlinear techniques such as various forms of hysteresis based on state or time can be extremely useful. The importance of filtering of various types has been reported for most of the applications. In addition, SPC techniques are now being thought of as a form of filtering for input to the rest of the expert system as well.

In addition to conventional filtering, other techniques can be quite useful. Conversion from numerical values to symbolic values of "high", "low", or "OK", significantly reduces the number of state changes in subsequent processing (forward chaining). Various symbolic forms of filtering, such as latching and event counting have a significant role to play as well. Some users found it necessary to delay fault alarms until the condition had been true for a period based on time or event counts.

Just as errors can be reduced by combining multiple sensor values using numerical models (such as Data Reconciliation or Kalman Filtering), sensor evidence can be combined using techniques more powerful than simple discrete logic (true/false values only). Industry has only begun to experiment with various models of evidence combination and fuzzy logic, which can also help address these problems. These techniques will become more prominent in future applications. Users have indicated a desire for ranked lists of possible faults, which requires better schemes than discrete logic.

Quantitative information and models are often needed
A significant amount of knowledge has been abstracted by engineers into mathematical models. The best systems are hybrids of qualititative and quantitative techniques. This is intuitive, because the system is taking advantage of more knowledge about the process. Furthermore, in many of these systems, the simulation is specified in an object-oriented form. The user often creates graphical objects with attributes, and the library equations directly derive the necessary mathematical representation from that structure

Diagnostic systems based on deviations from quantitative models tend to be very sensitive to faults of all types, even when operation is close to normal. Model-based approaches can significantly increase sensitivity to real faults. Also, the time to recognize those faults is shortened, because they are detected within the normal operating range, before significant harm is caused by the fault.

Approaches based on deviations from models (residuals) also have the advantage of detecting some faults which were not even anticipated, but which affect the variables in the model equations. (In that case, the system can alert the operators, although not necessarily identify the exact cause). A good example is material and energy balance equations which do not explicitly account for an actual leak or pipe break, since they are low-probability events. However, if a pipe does break, the resulting large material and energy balance equation residuals will quickly indicate a problem, even if the logic does not explicitly derive a specific conclusion beyond the initial problem detection. However, minor problems such as mild sensor drift, mild process upsets, slightly larger than normal noise, or modelling error can all lead to incorrect detection of faults.

Most useful industrial systems involving continuous variables are hybrids of the model-based and pure symbolic approaches. The models can generate residuals, which then feed into the symbolic logic.

Knowledge-based systems provide good repositories for process technology, improving the uniformity of operator responses
Since the embedded knowledge is visible to the operators, is testable and generally can be queried, the operators can use it as a learning aid, and can continue to refine it. Whether the operators take manual actions based on the system, or allow the system to directly manipulate the process, the results are higher uniformity of control actions.

CONCLUSIONS

Online KBESs are making significant contributions to process control and management. They are economically justified. The applications and benefits are often in areas which complement traditional process control technology, for instance, in handling abnormal situations, and in overall quality management. The KBES integrates new techniques with conventional controls.

Many lessons have been learned from the industrial experiences, such as the importance of filtering, the importance of integrating SPC tools, and the need for integration of quantitative models.

REFERENCES

Årzén, K.-E. (1990). Knowledge-Based Control Systems -- Aspects on the Unification of Conventional Control Systems and Knowledge-Based Systems. *Proc. American Automatic Control Conference (ACC-90), San Diego, CA*, 2233 -2238.

Hofmann, A.G., G.M. Stanley, and L.B. Hawkinson (1989). Object-Oriented Models and Their Application in Real-Time Expert Systems. *Proc. Society for Computer Simulation International Conference, San Diego.*

Finch, F.E., G.M. Stanley, and S.P. Fraleigh (1991). Using the G2 Diagnostic Assistant for Real-Time Fault Diagnosis. *European Conference on Industrial Applications of Knowledge-Based Diagnosis, Segrate (Milan), Italy, October, 1991.*

Fraleigh, S.P., F.E. Finch, and G.M. Stanley (1992). Integrating Dataflow and Sequential Control in a Graphical Diagnostic Language. *Proceedings of the International Federation of Automatic Control (IFAC) Symposium on Online Fault Detection and Supervision in the Chemical Process Industries, Newark, DE, April, 1992.*

Mertz, G.E. (1990). Application of a Real-Time Expert System to a Monsanto Process Unit. *Proc. Chemical Manufacturer's Association Process Control Conference, Miami, Florida, March, 1990.*

Moore R. and others (1988). The G2 Real-Time Expert System. *Proc. Instrument Society of America (ISA-88),* 1625-1633.

Moore, R., H. Rosenof, and G. M. Stanley (1990). Process Control Using a Real-Time Expert System, *Proc. 11th IFAC World Congress, Tallinn, Estonia, USSR*, 234-239.

Rehbein, D., D. Deitz, S. Thorp, and L. Tonn (1990). Expert Systems in Process Control, *Proc. ISA, New Orleans, Louisiana, USA, Oct. 14-18, 1990,* 1989-1995.

Rosenof, H.P. (1990). Expert Systems Come to Batch Process Automation. *Proc. Instrument Society of America International Conference (ISA-90), New Orleans, Louisiana,USA*, 1367-1376.

Rowan, D.A. (1989). On-Line Expert Systems in the Process Industries. *AI Expert*, August, 1989, 30-38.

Spang Robinson Report (1989). Real-Time Application Profile: Monsanto's Process Control System. *The Spang Robinson Report on Artificial Intelligence*, Nov., 1989.

Stanley, G.M. (1991). Experiences using knowledge-based reasoning in online control systems. *International Federation of Automatic Control (IFAC) Symposium on Computer Aided Design in Control Systems, July 15-17, 1991, Swansea, UK.*

Stanley, G.M., F.E. Finch, and S.P. Fraleigh (1991). An Object-Oriented Graphical Language and Environment for Real-Time Fault Diagnosis, *Proc. European Symposium on Computer Applications in Chemical Engineering (COPE-91), Barcelona, Spain, Oct. 14-16, 1991.*

Stephanopoulos, G. (1990). Artificial Intelligence ... What Will Its Contributions Be to Process Control? in *The Second Shell Process Control Workshop, ed. David M. Prett and others, Butterworth Publishers, Stoneham, MA,USA*, pp. 591-646.