

Plenary paper presented at: International Federation of Automatic Control (IFAC) Symposium on Computer Aided Design in Control Systems, July 15-17, 1991, Swansea, UK

EXPERIENCES USING KNOWLEDGE-BASED REASONING IN ONLINE CONTROL SYSTEMS

G.M. Stanley

Gensym Corporation, 125 CambridgePark Dr., Cambridge, MASS., 02140 (USA) *

Key words: Expert systems, Process control, Supervisory control, Knowledge engineering, Industrial control, Control applications, Artificial Intelligence, Real time computer systems, Automation, Sensor failures

ABSTRACT

This paper gives an overview of industrial applications of real-time knowledge based expert systems (KBESs) for process control. After a brief overview of the features of a KBES useful in process control, several case studies are reviewed. The lessons learned 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. 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). Overviews of KBESs are available elsewhere, so we will focus on aspects of particular relevance to control systems.

A general overview of object-oriented programming is given by Stefik & Bobrow (1986). Overviews of KB techniques in process control are given by Stephanopolous (1990), and Bristol (1989). Arzen (1990) gives an overview of the issues for unifying control systems and knowledge-based systems. 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).

KBES and control technologies are complementary, rather than competitive technologies. 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 control 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
- 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 really a graphical language. Continuous control system engineers generally think in terms of data flow languages consisting of processing blocks and signals, as described earlier. GRAFCET (Baker and others, 1987) is an example of a graphical language for sequential control which can be built in a KBES (Årzen, 1990). In GRAFCET, objects representing actions are connected by directed arcs specifying sequential or concurrent execution.

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 CONTROL

Some roles of expert systems in process control have been outlined by Stephanopolous (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: detection, root cause analysis, repetitive problem recognition
- 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)
- 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

The following applications are likely candidates for future industrial use:

- Predictive maintenance
- Process validation
- Intelligent supervision of adaptive control, model identification, parameter estimation, data reconciliation, and state estimation
- Automated design of control systems (and implementation)
- Intelligent supervision for optimization

In the area of adaptive control, Åstrom, Anton & Årzen (1986) describe an experimental expert system implementation of an autotuner. Foxboro sells an autotuning control system implemented in hardware, trademarked as the EXACT™ controller (Bristol, 1989). It has been successfully applied in numerous applications. It is advertised as an expert system, capturing extensive knowledge on controller tuning. In that sense, it is acting as an expert system. However, the controller does not appear to have been implemented using expert system techniques. It is described by Bristol as "10-15 pages of FORTRAN or C, illegible to outsiders (even to the concept inventor!)". It appears that the usual advantages of simple, readable, maintainable representation in the form of rules and objects in a KBES were not achieved. EXACT has proven the existence of a market for intelligent low-level control and adaptive control for at least some applications. Major issues besides robustness may well be packaging and cost.

CASE STUDIES OF INDUSTRIAL APPLICATIONS

Monitoring and closed-loop control of salt-water desalination by Reliable Water

The Reliable Water Company builds unattended salt water desalination plants. The plants are typically installed in remote locations where no operator is routinely on site. Technicians are often not conveniently available on site, either, so the system has to be very reliable, and diagnose itself when it breaks down. The system requires more control than the typical reverse osmosis plant, due to a unique energy-recovery technique. There is a significant amount of sequential control, and some continuous control.

The knowledge-based control system is at the forefront of automation practice (Yankee Conveyer, 1989). The plant only runs with the expert system: there are no "manual" operations, manual valves, or independent dials or gauges. All sensors are computer inputs, and all manipulators are computer-controlled. In addition to control, the KBES also includes maintenance management, inventory management, performance monitoring, and repair advice.

The input sensors include pressure, flow, pH, salinity, conductivity, switch position, and various equipment statuses. Outputs include valves, pump power, circuit breakers and similar equipment. A typical system has about 50 analog and digital inputs, and 50 outputs. The system has about 350 variables, 700 rules, 470 procedures, 70 functions. and about 150 generic formulas.

The KBES is constructed hierarchically. There are rules at a high level, such as "if water is too-salty then ...", which trigger lower-level rules. There are explicit high level goal rules built in to consider the effects of control actions, and to prevent damage. The highest level goal is "make drinking water". This decomposition makes it easy to enter high-level rules without worrying about the operational details of the low-level systems.

To eliminate the need for operator intervention, a significant portion of the system is dedicated to recovering automatically from extreme conditions. For instance, the system can recover from a temporary power failure

caused by lightning. First, the plant is shut down properly. The system finds the tripped circuit breaker, physically resets it, checks all the equipment, and then restarts the plant.

The system detects failures, and substitutes redundant equipment where possible. Sensors are routinely calibrated and cross-checked. When a sensor fails, a calculated value is substituted instead. The system runs the plant at partial capacity, or in some suboptimal mode if necessary.

The knowledge-based approach makes the plant more reliable and less costly to operate. The complete knowledge-based system took about five man-years to implement. Plants are in successful operation throughout the world.

Plant experiences in monitoring, simulation and control at DuPont

DuPont has numerous online KBESs for diagnostics, data reconciliation, scheduling, optimization and control (Rowan, 1989; Rehbein and others, 1990; Schreiber, 1991). They have also combined object-oriented modelling with object-oriented expert systems.

FALCON, DuPont's first online expert system, performed fault diagnosis system at a nylon plant in Victoria, Texas (Rowan, 1987,1988). It was a forward-chaining rule-based system written in custom LISP code, containing about 650 specific rules, and based on 31 input variables.

An extensive dynamic simulation was built to test the expert system during development, and verify the fault detection for at least 25 likely faults. Faults included failures of pumps, heat exchanger fouling, off-gas flow restriction, and control system component failures. Only single failures were successfully diagnosed.

Diagnosis based on quantitative, first principles models was combined with empirical, "heuristic" plant-specific knowledge and historical trend data. Overall redundancy of data combined with models was essential, so that basic techniques associated with data reconciliation could be applied to detect the faulty sensors. For instance, the imbalance of material flow about one or more process nodes helped pinpoint measurement failures, as shown in Mah, Stanley & Downing (1976).

Fault diagnosis during abnormal extremely dynamic operations such as startup was difficult. A significant effort was put into making the system robust enough to recognize that the plant was in a safety interlock mode, or even shut down. This alarm filtering was necessary to avoid erroneous or nuisance advice.

The diagnostic system detected 5 out of 6 actual plant faults of the defined types during the first months of operation. However, at first it also made nuisance announcements due to noise, sensor errors, and spurious process conditions. It was necessary to suppress the announcement of faults until they had been repeatedly detected. To handle problems of noise and slow sensor drift, it was necessary to introduce filtering on the sensors, and reason based on both raw and filtered data. In particular, slow thermocouple drift led to many false announcements due to errors in the energy balance calculations. It was necessary to periodically calculate a filtered bias term for each energy balance to help avoid these errors. This model residual bias approach to dealing with slow sensor drift was developed by Stanley (1982). Engineering judgement was required to tune the system to avoid nuisance announcements, but still detect real faults fast enough for operators to prevent a dangerous condition.

DuPont was encouraged by the technical success of the project, and implemented other online systems based on G2. The custom-coding approach of FALCON in LISP was not economically justified, due to high development and maintenance costs.

DuPont has described a newer KBES, based on G2, installed at the same nylon plant. The incentives for this quality control system are uniformity of actions across all operations, reduction in product variability, improvement in adipic acid purity, reduction in waste generation, and ISO 9000 (safety) certification. The

KBES supervisory control includes the balancing of process flows and inventories, and controlling stream compositions. This system is an integral part of the plant's quality management effort. An important part of the system is the graphical interface, which automatically displays easy-to-follow response plans.

This system has been running online since March, 1990. The system has improved production capability and quality control, and reduced losses. Also, system has proved to be a valuable tool for technology transfer. That is, the expert system helps the operators understand complex process behavior and improves the uniformity of their actions. The KBES distributes expertise, so that the best operator and engineering team is in effect always available, and always vigilant. (Normal estimated expert availability is 15% in a round-the-clock plant operation.) Due to the current success, the system is being further expanded during 1991.

DuPont is not widely publicizing details of its current online systems, as it considers them to be a competitive advantage. They have stated that they "routinely" see returns on investment as high as 10 to 1 (Rehbein and others, 1990; Rowan, 1989). The economic impact is typically in improved yield, quality, and utilities.

Diagnostics at the Monsanto's Krummrich Chemical Plant

Monsanto has multiple online installations. One major success was in building a generic library for fault diagnosis at the Krummrich chemical plant (Mertz, 1990; Spang Robinson Report, 1989; Rehbein et al, 1990). Portions of the library have been installed at other sites. The rules focus on low-level failures such as sensor or valve failures, which are common to all plants. The application takes advantage of the class hierarchy and direct analysis of the process schematic. The system also includes online troubleshooting documentation.

The system performs transmitter and control loop validation, analyzing sensor and control loops. Sensors and valves are the weakest links in the overall control process. Quantitative material and energy balance models are used as part of this validation, leading to the same sensitivity to noise and drift experienced by DuPont. Some failures, such as frozen valves or plugged tank vents, cause no immediate alarm, but can be detected. The application at the Krummrich Plant analyzed 200 of 475 sensors, and 100 out of 135 control loops. The rules for flow, level, amperes, and control loop validation were all done generically. The logic for pressure and temperature had to be customized beyond the generic library in some cases.

Noise was identified as a problem at this installation, leading to false alarms. One technique used is delay of any operator notification of problems until a fault conclusion has been reached a number of times, as done at DuPont.

A major goal is to find faults either not detectable by the operator and control system, or find them sooner. By finding the minor problems quickly, you can prevent their effects from propagating and becoming major problems. In addition to finding the fault, the system gives assistance during the upset. Monsanto internal studies had shown that operators' inattention to small problems (not generating immediate alarms, such as stuck valves) often led to major process upsets. The operators also had trouble understanding process data and responding properly during major process upsets.

The benefits are safety, quality control, environmental performance, on-stream time, and the ability to safely accomplish further optimization. Monsanto's own evaluation was that process interruptions cost \$500,000 each year, and that half of that is saved by the expert system.

The system at the Krummrich plant has been online continuously since September, 1989, shut down only a few times due to power failures. The first project paid for itself and other systems have been installed.

Monitoring, simulation, and operator-in-the-loop control in a paper mill - EPAK

EPAK provides decision support to the operators in a paper mill run by Norske Skog AS of Norway (Yeager, 1990; Opdahl, 1989). The highly-automated plant is Norway's largest paper mill. The system is a supervisory control system, with a human in the loop to approve the control changes.

EPAK, built on top of G2, is available from ABB AFORA of Finland. The system assesses paper quality, recommends control actions, and uses simulation to predict the effects of the recommended actions. The incentives are to improve quality, minimizing variations between shifts. The initial system took about 6 months to build, although the commercialization continued after that time.

There are three types of knowledge captured in the system: "assessment knowledge", "correction knowledge", and "process knowledge". The assessment knowledge is used to symbolically classify the quality as high, low, or OK. The correction knowledge recommends the control actions, accounting for interactions. Rules are used to select among combinations of directional changes (steps) in one to three control parameters, balanced to avoid unwanted side effects. Some numerical modelling is used. In short, the correction knowledge maps quality changes into appropriate manipulated variable changes. The inverse of this knowledge is called process knowledge. It predicts the impact of manipulated variable changes on the product qualities. Process knowledge also includes limits on manipulated variables and various calculations.

Some of the process knowledge is represented as a "Quality Matrix". It is used for online simulation to predict states and the effects of changes. This is a linear model with time delays, with different coefficients for each paper grade. The knowledge is provided both by engineers and by operators. The quality matrix is directly accessible as a set of graphical objects and connections, for inspection and change. There is a set of input objects representing up to 10 control parameters, and a set of output objects representing quality parameters. The connections between the objects have an attribute containing the linear effect coefficients. Other process knowledge is represented by rules and functions.

The possible (multivariable) control actions are stored as individual objects. These objects have attributes containing the name of the quality property, the direction of the desired change, the applicable control parameters, and the direction of change of those control parameters. The actions are prioritized, and the priorities can be changed online. Priorities are changed for instance, for each customer, or when constraints are violated, or to reflect changes in business goals, or to reflect new insights. Rules decide which of these objects will generate a manipulated variable change.

The control actions frequently include iteration: a small change is made, to see what the effects will be. If the process is improved as result of those changes, further steps are taken. The system tracks all changes made and predicts their effects, so that control actions are based on predicted deviations from targets. By iterating, the system uses both correction knowledge and process knowledge. The simulation presents the predicted effects to the operator. The operator can run a "what if" analysis to simulate the effects of his own proposed actions. This prediction is very popular with the plant operators, and was a major contributor in building confidence and obtaining operator acceptance.

Operators' proficiency improved, and more timely decisions are made. By speeding decisions, out-of-specification product per unit of production is reduced. Economies are achieved, because the tighter quality control allowed closer approaches to constraints. For instance, use of expensive chemical pulp is held to a minimum. Due to better color control, downgraded and repulped paper has been reduced. The system became the most knowledgeable source of information on quality as related to paper machine operation. Quality-related decisions are now always as good as the best operator's decisions. A framework is in place for continued addition of operating knowledge as well.

Planning and closed-loop control of composite curing at Alcoa

An object-oriented, rule-based control scheme was found to be very effective for controlling the curing of composite materials in autoclaves and presses (Manzini & Roehl, 1990). Parts are created with woven graphite fibers and various epoxy resins. The parts must be mixed following a recipe, then heated under pressure to stimulate the chemical reaction, and then cooled when the reaction is done. During the chemical reaction, heat is generated as well. Significant gradients of temperature across the thicknesses of the material

can occur, which must be carefully controlled to maintain strength of the final material. Traditional open-loop ramp-and-soak controls can compensate for thermal lags, but this leads to excessive production times.

The cure state is inferred indirectly from measurements of temperature and pressure, and dielectric properties. As a resin cures, it progresses through a sequence of distinct states, each of which requires a different type of control. A major problem is detecting the state transitions, so that the control strategies can be switched. Detecting the state transitions requires sensor values and estimates of their derivatives. There is a strong incentive to recognize the state transitions and minimize the processing time. The process monitoring system is an integral part of this control scheme.

The control of the part manufacturing is based on its geometry and material, the process equipment used, and constraints on non-uniform temperature distribution. Each part can have a variety of different features, such as thick and thin regions, further complicating the control. Extreme flexibility is needed, due to low production volumes, and large variations in the materials, part geometries, and processing rules. Traditional control techniques have been inadequate for materials manufacturing of this type.

A knowledge-based approach is used, which integrates a variety of expertise in the form of models, data, and heuristics. Objects describe the parts, the status of the various portions of the parts, the process requirements, the behaviors of the resins, the processing equipment, and the sensors. Connections between the objects are used to represent the physical relationships between the objects. The class hierarchy with inheritance is used extensively to simplify specifications of equipment and sensor. The class hierarchy with inheritance is also extremely useful in representing the processing requirements and recipes of the resins, which fall in several broad categories.

Based on the complete set of requirements for the part, a plan for control of the manufacturing process is automatically developed using a set of rules which analyze all the information contained in the specification objects and connections. The plan covers the complete phases of startup, running, and shutdown. The system then makes the part, running on a 60-second cycle of data acquisition and control.

The system is successful, with processing time savings of an average of 30% over the previous conventional open-loop strategies.

Waste incineration monitoring at 3M

The KBES described by Kinoglu (1991) monitors a solid waste incinerator plant, including the following subprocesses: combustion in a rotary kiln, steam generation, and pollution control. The system makes recommendations such as cycle completions and shutdowns. For instance, pressure drop trends of the energy recovery equipment are monitored to watch for metal build-up, which prompts a shutdown. This system combines quantitative process and control models with expert system techniques. The simulation portion of the system makes extensive use of the object-oriented style of modelling based on the plant schematic. The operators or engineers can also use the simulation in a "what-if" mode to evaluate possible changes or enhancements.

There is a significant variation in the unknown contents of each drum of waste material fed to the incinerator. So, part of the monitoring system estimates combustion properties of the drum contents, and then predicts via online simulation the emission levels at the precipitator and stack. The system tells the operators how long a drum should be burned in the kiln (cycle time). This system helps reduce the time spent on each drum, and also ensures adequate incineration. By eliminating conservative, fixed cycle times, the overall plant capacity is effectively increased.

The system is running online 24 hours a day. The early online results of the system were very encouraging. In particular, two accidents were averted in a fairly short time. The system has been detecting problems and recommending actions much faster than previously possible, and also detecting problems indicated only by

slowly-changing variables which were unnoticed by the operators. Due to accident prevention and better selection of shutdown times, the system paid for itself in a short time. The system is being expanded to model more of the plant, and to expand the advisory rules. The system is being expanded to monitor the tank farm, and provide step-by-step instructions to the operators for their manual activities. Closed-loop control of some parameters is anticipated by the end of the 1991.

Closed-loop environmental control, analyzer control, and monitoring at Biosphere II (Space Biosphere Ventures)

Biosphere II is a 3-acre closed system which can be completely sealed off from material exchange with the world. 8 people, plants, and animals will be sealed in the airtight greenhouse in Arizona, maintaining their own atmosphere for 2 years. A variety of biomes, each with their own climate, are maintained. Knowledge-based systems are being used in 4 major areas: climate control, monitoring, analyzer sequencing control, and design simulation (Girard, 1990). The heart of the system is a network of 6 computers with expert systems sharing information for distributed control and monitoring, and exchanging data with other computers.

Both control and monitoring of the environment are very critical. A breakdown in systems control and monitoring, leading to high temperatures, plant death, and a CO₂ runaway, could leave the environment uninhabitable in a day, ruining the efforts of many people at a cost of many millions of dollars. Tight control is desired so that experiments can be run on the effects of the controlled variables.

The KBES controls temperature, humidity, and air velocities. The control problem is more difficult than in the case of office buildings, where the areas are physically separated and the target temperatures and humidities are roughly the same throughout the building. Also, since the entire outside of the system is glass, variations in sunlight have a stronger effect. Each biome requires different targets. There are significant interactions between the biomes, because they are not isolated from each other. The control system is distributed among several knowledge bases and computers, with feedback and feedforward controls outputting to low-level control hardware.

The analyzer/sampler control uses a graphical, sequential control language developed in the KBES. This flexible system coordinates the analyzer with automated sample location changes, and performs automatic calibration.

At the time this paper is being written, the G₂-based supervisory controls have not been tested, pending completion of the lower-level control hardware. The control system KBES was implemented in about 6 man-months. The sequential control of the analyzer and its sampling system have been operational online for many months. Its implementation took 1 - 2 man-months. Portions of the monitoring systems needed before complete closure of Biosphere II have been in use. That implementation took about 6 man-months.

Fermentation reactor supervisory control at the University of Newcastle and MIT

At the University of Newcastle upon Tyne (U.K.), knowledge-based systems have been built which span supervisory control, scheduling, and fault detection/diagnosis for fermentation reactor control. (Aynsley and others, 1989, 1990; Morris and others, 1991)

Fermentation control is difficult because it is nonlinear, time-varying, and models are poor. It is usually a batch process. Measurements of living processes are mostly indirect, such off-gas analysis and mass balancing. There are large measurement uncertainties. There are many interacting variables, through different phases (fast growth vs. slow growth, etc.).

The Newcastle KBES applies a set of rules to the online data to determine the current phase of operation (microbial growth). The current phase then determines which other rule sets apply to operations within that phase. The system also checks for data consistency using overall redundancy of models and measurements.

For the diagnostic portion of the KB, the hardware fault detection is fairly generic, while the "metabolic faults" are more specific to the organisms in the given fermentation. In addition to diagnosis and growth phase determination, the knowledge-based system is responsible for feed scheduling, and modifications to the schedule in response to faults or other process changes. The system determines the best time to terminate the fermentation and harvest the product.

The initial control scheme was tested first on a dynamic simulation. The system was also tested on offline industrial data, and detected a contaminated batch in only 17 hours, compared to the 75 hours it took people in the actual industrial plant. Had the system been online at the time, a considerable saving in substrate feed costs and reactor time would have been achieved.

The system currently supervises pilot-scale Bakers Yeast fermentations online at the university. An industrial application is now being developed. Similar efforts are in progress at MIT's Bioprocess Engineering Center, and closed loop control has been operational. Results of that work will be published.

OTHER NOTEWORTHY APPLICATIONS

Operator advice at an ISI Agroindustriale Sugar Company (Italy) sugar plant

A KBES has been installed online at an ISI Agroindustriale Sugar Company in Italy (Filippini et. al, 1990). The system advises operators in the following areas: evaluation of regulatory control loops, prevention of malfunctions, standardization of quality and yield, and optimization. The system has been developed by 3 engineers over a three month period. It has been well-received, and is considered an essential tool in plant management. It is being expanded.

Closed-loop supervisory control of a concentrator at Noranda

Noranda built a knowledge-based closed loop supervisory control system for Brunswick Mining & Smelting, controlling a concentrator (Northern Miner Magazine, 1990). The system reads and/or changes the values of over 500 instrument signals. The first application was online dynamic material balance control. The control is based directly on analyzing the process schematic. The system paid for itself in the first four months of operation.

Distributed Control System configuration validation at Simons Eastern

The flexibility of modelling within a graphics-oriented KBES was used to develop a system to test distributed control systems (Anderson, 1990). The system simulates a process and validates the DCS configuration. The system has successfully validated several DCSs at Simons' staging facility. It is being expanded to use more "AI" capabilities to automate more of the overall test procedure. For instance, the system will plan process disturbance tests and analyze the controller behavior.

Space Shuttle monitoring at NASA's Johnson Space Center

NASA uses knowledge-based systems to monitor the Space Shuttle at the Johnson Space Center (Muratore, Heindel and others, 1990; Pohle, 1991; Montgomery, 1991; Girard, 1990). The goal is to improve the speed and quality of Mission Control decisions. There are about 5 different applications online, and another 5 are in progress. One expert system monitors the status of 38 primary Reaction Control Jets on the shuttle. The system can detect current problems, predict potential problems, and simulate the consequences. The KBES replaces a 20-page procedure manual, and does necessary calculations to speed decisions by 20 - 45 minutes. Another application is the Booster Flight Controller expert. It also can detect failures and simulate different scenarios. During ascent or entry, the flight controller has less than 20 seconds to detect a problem, analyze it, and take action. The KBES has demonstrated the ability to correctly identify all known booster failure scenarios, and simulate additional potential ones, in a few seconds. Typical applications have taken two to six months to build.

Noteworthy projects in progress

HALDRIS is an online expert system for supervisory control and diagnosis of aluminum smelters at Hydro Aluminun, Norway (Rolland et al, 1991; Fjellheim, 1990). After operating in simulated mode, the prototype was put online at the plant in November, 1990. Full scale implementation and validation will be completed in 1991. Average yields are expected to be increased by .5%.

The OECD Halden Reactor Project is a European nuclear safety project. The SAS-II KBES for safety assessment and post-trip guidance (Nilsen, 1990) is currently running with a detailed simulation. Plant installation is in progress at the Forsmark-2 power plant in Sweden. The knowledge is based on physical model and logic diagrams, consisting of interconnected objects such as AND gates, OR gates, and NOT gates, which then drive alarm objects. The gate objects and interconnections define a graphical language replacing the equivalent rules.

An online emergency procedure management system has been developed for a nuclear power plant in Belgium (Foret, 1990). The project took 3 months, and was tested with operators using an accurate online training simulator. Operator training was found to take just a few minutes.

As part of the Advanced Solid Rocket Motors (ASRM) program, Lockheed is building a pilot plant and then a full-scale plant for solid fuel propellant for NASA. Lockheed has been using an object-oriented simulation of the plant and control system for the design phase (Braunstein, Brown and others, 1989). Implementation will use the same tools. That is, the controls developing using the knowledge-based tool are tested on the simulation, and then will be used in the plant when it is built. The complete implementation integrates process simulation, process control, monitoring, and fault diagnosis. The plant model is based on the process schematic. Simulation statements and rules analyze the schematic.

GENERAL LESSONS LEARNED

Real-time KBESs are robust enough to have succeeded in numerous applications, including closed loop control.

The summary of applications at DuPont, Reliable Water, Alcoa, University of Newcastle, and others makes this clear. As time goes on, more of the systems will become closed loop. 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, and operators routinely approve the recommendations. These current systems are doing control -- the human acts as a random deadline in the loop.

Also, many of the open-loop applications will migrate to closed loop someday, as people build up confidence, following the earlier pattern of migration of computer open-loop monitoring to closed-loop control. ABB's EPAK may be example of that.

Except in the Reliable Water case, there is generally a DCS or PLC at the lowest control level in KBES applications. These systems perform the simple tasks adequately, run on higher-reliability hardware than most computer systems, and are already in place before most of the KBES projects are started.

Significant benefits are derived in areas complementary to conventional controls, such as diagnosis, quality management, and abnormal operation

Significant benefits have been achieved, as shown in the case studies. Most 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. Diagnostics are needed as part of the overall control system to catch major problems 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. These were major issues in the composite curing process and the fermentation control systems.

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.

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. This was especially apparent for Reliable Water and Alcoa. 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 help 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 (as was done to some extent at Alcoa). 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.

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.

Maintainability is a major issue

The FALCON system, and similar early 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.

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 (as in the case of FALCON), 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 does not need to garbage collect at run time.

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. This type of high-level representation was a key to the Alcoa and Reliable Water examples.

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.

KBESs complement control systems by dealing with diagnosis and abnormal conditions

KBESs provide new tools which extend & complement existing techniques, rather than replacing them. This is natural, because a KBES emphasizes qualitative, symbolic methodologies rather than quantitative methodologies. It emphasizes representation of concepts at a high level of abstraction or multiple levels, instead of the low-level plant models.

As an example, conventional continuous control schemes usually depend implicitly or explicitly a small-signal system model linearized near some nominal value. The multivariable controls normally assume that all sensors are good. These controls work well during normal operation. However, in case of a process upset, sensor failure, or other fault, many control systems are put in manual mode by the operators so that more extreme corrections can be made directly. The normal models and controls break down during the extreme operation. There, the more effective models or actions are likely to be simpler, but based on first principles or else on heuristics. These alternate controls are easier to build in a KBES than in conventional systems.

KBESs also augment batch control systems, by providing better recognition of the transitions between different operating phases.

KBESs complement SPC techniques by earlier problem detection and root cause analysis, achieving Real-Time Quality Management

SPC tools are sensitive detectors of problems. However, they offer no guidance as to the root causes of problems, or how to correct the problems. This is a fundamental limitation, because SPC techniques cannot capture process model knowledge and use it. A KBES can apply SPC to detect problems, and integrate a diagnostic system to 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. Monsanto provided a classic motivating example similar to the following: if a valve sticks, a reflux drum will empty, ending reflux to a distillation tower, finally causing product to go out of specification half an hour after the fault. Diagnostics could detect the stuck valve in much less than a minute.

Organizational and implementation issues

Some of the factors in successful applications are highlighted by Rehbein et al (1990). Factors contributing to overall success are generally similar to those for process control projects, such as end-user and expert involvement, management commitment, and so on.

Basic software engineering and project management techniques are still needed, although the process usually follows more of a prototyping and iterative refinement approach with extensive end-user review.

Simulation is useful both for model-based reasoning and for testing

LESSONS LEARNED IN KNOWLEDGE REPRESENTATION

Graphical specification of knowledge is effective

Users like developing graphical problem-specification languages. For example, Exxon is building a toolkit based on logic networks made up of AND gates, OR gates, and so on (Weber, 1991). The Halden project uses a similar approach, as does a nuclear plant monitoring system being developed in Japan. The Alcoa application, EPAK and others made extensive use of relationships between objects as a form of graphical specification language. Many of the applications, e.g., at Monsanto, derive the needed information directly from process schematics.

Generic knowledge libraries shortening development time

Many users (e.g., Monsanto, ABB, and others) 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. Since most industrial real-time expert systems are primarily forward-chaining, the processing load depends on when the system inputs have changed enough to justify propagating new information (Washington & Hayes-Roth, 1989).

To reduce the impact of noise, you can filter heavily and accept the phase lag 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, for either analog or symbolic data, can be extremely useful, even though they might not be normally be used in closed-loop feedback control.

In addition to conventional filtering techniques, 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. Various symbolic forms of filtering, such as latching and event counting have a significant role to play as well. Monsanto and DuPont both found it necessary to delay fault alarms until the condition had been true for a period based on time or event counts. Forward chaining itself is often configured to only propagate truly new state information, and this provides a form of filtering.

The importance of filtering of various types has been reported for most of the applications. Filtering is considered an important part of the toolkit being developed by Exxon (Weber, 1991). SPC techniques are now being thought of as a form of filtering for input to the rest of the expert system. This has been the case in the DuPont, Exxon, and others.

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

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 qualitative and quantitative techniques. This is intuitive, because the system is taking advantage of more knowledge about the process. Furthermore, in most 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. The sensitivity of model-based approaches is good for significantly increasing 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 pinpoint the exact cause). A good example would be 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. The developer must pay extra attention to filtering in this case.

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

This has been a major finding of the DuPont projects and EPAK, for example. 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.

USER AND DEVELOPER INTERFACE LESSONS LEARNED

Operators can use a KBES

Operators can indeed learn to use a system based on windows, graphical objects, a mouse, and menus. Some had previously expressed doubt about this, just as skeptics once doubted that operators could use a CRT and keyboard rather than panel boards.

User and end-user interface is important, and well-supported by KBES

Good end-user interfaces are important, since the results are often advice. The user wants to be able to understand the recommendation and the current process condition as quickly as possible.

The development of graphics tied closely to the objects has led to much progress in user interfaces. In fact, modern "direct manipulation" graphics packages are generally object-oriented, so that one overall technology now spans the needs of control, simulation, and graphics interfaces.

There is now less of a gap between representation schemes for developers and end-users. If the developer is specifying the system in terms of graphical objects directly understandable by the end-user, it is natural to try to make the objects perform duties for both development and end-use. This approach has major advantages for the developer. For instance, debugging is significantly simplified. Also, the end users themselves can

recognize mistakes quite easily, for instance, in specification of the plant schematic. This is in great contrast to the early days when developers worked in assembly language or FORTRAN, totally unreadable by the end-user. The KB becomes a repository of common information available for users with different needs.

A factor in success is that the operator is always given an opportunity to override the advice or actions of the expert system, even in closed-loop applications. The operator in the DuPont applications also has the ability to enter a comment indicating the reason for rejecting the advice. The information is periodically reviewed by the local engineer, and the knowledge base is updated as needed. In general, it has been found that the expert system learns indirectly from the operators, and that the operators learn from the expert system as a tutor as well.

Users would like ranking of possible faults, or some numerical measure of probability

This has been highlighted in the DuPont experience, and is built in to the Exxon toolkit.

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. To capture the benefits of these lessons, so that future implementations will be simpler, Gensym had developed a product, the Diagnostic Assistant™ (Stanley, Finch, & Fraleigh, 1991). The product is a primarily a graphical language. It combines both the data flow and sequential control aspects of other graphical languages.

REFERENCES

- Anderson, J.P. (1990). Distributed Control System Testing. *Proc. Instrument Society of America International Conference (ISA-90), Oct. 14-18, 1990, New Orleans, Louisiana, USA*, 791-797.
- Årzén, K.-E. (1990). A sequential function chart toolbox. *Proc. European Gensym Users Society meeting, Munich, Germany, Dec., 1990*.
- 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.
- Åstrom, K.J, J.J. Anton, and K.-E. Årzen (1986). Expert Control. *Automatica*, 22 (3), 277-286.
- Aynsley, M., D. Peel, and A.J. Morris (1989). A Real-Time Knowledge-Based System for Fermentation Control. *Proc. ACC, Pittsburgh, USA*, 2239-2244.
- Aynsley, M., A.G. Hofland, G.A. Montague, D. Peel, and A.J. Morris (1990). A Real-Time Knowledge-Based System for the Operation and Control of a Fermentation Plant. *Proc. ACC, San Diego, USA*, 1990, 1992-1997.
- Aynsley, M., D. Peel, A.G. Hofland, G.A. Montague, and A. J. Morris (1990). Real-Time Knowledge-Based Systems in Fermentation Supervisory Control. *IEE Vacation School on Knowledge-Based Systems for Industrial control, Strathclyde, Sept. 1990*, Chapter 11, Peter Peregrinus Ltd.
- Baker, A.D. and others(1987). GRAFCET and SFC as Factory Automation Standards: Advantages and Limitations. *Proc. American Control Conference (ACC-87), Minneapolis, Minn., USA*, 1725-1730.
- Braunstein, D.M., A.S. Brown, D.L. Drew (1989). A Knowledge-Based System for Process Monitoring and Control in Propellant Manufacturing. Paper #89-0486, *Proc. ISA (Instrument Society of America), October, 1989*.

Bristol, E.H. (1989). Objects, Rules, and Process Control. *Proc. American Automatic Control Conference (ACC-89)*, 1368-1376.

Elmqvist, H. and S.E. Mattsson (1986). *Proc. IEEE Control Systems Society Third Symposium on Computer-Aided Control System Design (CACSD), Arlington, VA, Sept. 24-26, 1986.*

Filippini, F., S. Urbinati, and M. Solimano (1990). An Expert System for the Prevention of Malfunctions and Technological Optimization in a Beet Sugar Plant. *Proc. European Gensym Users Society Meeting, Munich, Germany, Dec., 1990.*

Fjellheim, R. HALDRIS: An Operator Support System for Aluminium Production. *Proc. European Gensym Users Society Meeting, Munich, Germany, Dec., 1990.*

Forêt, J.-M., and E. Abbeloos (1990). On-line Procedure Management System: A Layer Above G2. *Proc. European Gensym Users Society Meeting, Munich, Germany, Dec., 1990.*

Girard, J.E. (1990). You Bet Your Life It's a Network. *New Science Associates AI Newsletter*, May 7, 1990.

Girard, J.E. (1990). NASA Applies Gensym to Real-Time Problems. *New Science Associates AI Newsletter*, July 16, 1990.

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.*

Kinoglu, F. (1991). A Real-Time Expert System for Monitoring the Waste Incineration Process. *Proc. 1991 Simulation Multiconference, New Orleans, Louisiana, USA, April 1-5, 1991.*

Mah, R.S.H., G.M. Stanley, and D.M. Downing (1976). Reconciliation and Rectification of Process Flow and Inventory Data. *Ind. Eng. Chem. Process Design & Development*, 15, (1), Jan., 1976, 175-183.

Manzini R.A. and E.A. Roehl (1990). Flexible Control of an Organic Matrix Composite Cure Process Using Object-Oriented Control Concepts. *Proc. American Automatic Control Conference (ACC-90)*, San Diego, CA, 1980-1985.

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.*

Montgomery, R. (1991). Jet Control. *Proc. Gensym Users Society Annual Meeting, March 6-8, 1991, Houston, Texas, USA.*

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.

Morris, A.J., G.A. Montague and others (1991). Towards Improved Process Supervision -- Algorithms and Knowledge-Based Systems. *Proc. Fourth International Conference on Chemical Process Control, South Padre Island, Texas, USA, February 17-22, 1991.*

Muratore, J.F., T.A. Heindel and others (1990). Real-Time Data Acquisition at Mission Control. *Comm. ACM*, Dec., 1990, 18-31.

Nilsen, S. (1990). Experiences Made Using the Expert System Shell G2. *Proc. 2nd International IEEE Conference on Tools for Artificial Intelligence, Washington, DC, Nov. 6-9, 1990.*

Northern Miner (1990). High-Tech Operations: The Manless Mill. *The Northern Miner Magazine*, May, 1990, 22-23

Opdahl, P.-O. (1989). EPAK: An Expert System for the Support of Paper Quality Control. *Proc. Gensym User's Society Fall 1989 Meeting, Cambridge, Mass, USA.*

Pohle, G. (1991). DATACOMM - Recorder Operations Expert System. *Proc. Gensym Users Society Annual Meeting, March 6-8, 1991, Houston, Texas, USA.*

Pohle, G. (1991). Data Acquisition for G2 using RTDS. *Proc. Gensym Users Society Annual Meeting, March 6-8, 1991, Houston, Texas, USA.*

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.

Rolland, W.K., A. Steinsnes, A.S. Larsen, K.A. Paulsen, P. Opdahl, H. Smith-Meyer (1991). HALDRIS - An Expert System for Process Control and Supervision of Aluminum Smelters. *TMS (The Minerals, Metals & Materials Society) Meeting, New Orleans, USA, February, 1991.*

- 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. (1987). Using an Expert System for Fault Diagnosis. *Control Engineering*, February, 1987, 160-164.
- Rowan, D.A. (1988). AI Enhances On-Line Fault Diagnosis. *InTech*, 35 (5), 52.
- Rowan, D.A. (1988). Applying On-line Expert Systems: Technical Issues and Business Concerns. *Chemical Manufacturers Association Sixth Process Control Users Forum, April 10-13, 1988, New Orleans, Louisiana, USA*.
- Rowan, D.A. (1989). On-Line Expert Systems in the Process Industries. *AI Expert*, August, 1989, 30-38.
- Schreiber, L.A. (1991). Utilization of G2 for Production and Quality Control. *Proc. Gensym Users Society Annual Meeting, March 6-8, 1991, Houston, Texas, USA*.
- 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. (1982). Online Data Reconciliation for Process Control. *AIChE Annual Meeting, Nov., 1982*.
- 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*.
- Stefik, M. and D.G. Bobrow (1986). Object-Oriented Programming: Themes and Variations. *The AI Magazine*, 6, (4), Winter, 1986, 40-62.
- 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.
- Washington, R. and B. Hayes-Roth (1989). Input Data Management in Real-time AI Systems. *Proc. Eleventh International Joint Conference on Artificial Intelligence (IJCAI-89), Detroit, Michigan, USA*, 250-255.
- Weber, R. (1991). Real-Time Diagnostic Toolkit. *Proc. Gensym Users Society Annual Meeting, March 6-8, 1991, Houston, Texas, USA*.
- Yankee Conveyer* (1989). Artificial Intelligence Takes Over Process-Control Plant. *Yankee Conveyer, The Yankee Group, Boston, MA, USA*, May, 1989.
- Yeager, R. (1990). Real-time Experts Come to the Paper Industry. *PIMA magazine*, April, 1990, 45-47.
- Yeager, R. (1990). Norske Skog Uses Expert System to Optimize Production Quality. *Pulp & Paper*, December, 1990, 112-115.

* The author, Greg Stanley, may now be contacted at:

<http://gregstanleyandassociates.com/contactinfo/contactinfo.htm>