PITCH EXPERT
A Problem-Solving System for Kraft Mills

Allan Kowalski, Diana Bouchard, Lawrence Allen, Yves Larin,
and Oliver Vadas

■ PITCH EXPERT was developed to make expertise available to mill-site engineers to solve pitch problems in kraft pulp mills. These problems have been estimated to cause losses to the Canadian pulp and paper industry in excess of $80 million each year. The design of the system took into account not only the complexity of the process interactions and the need for accuracy and completeness of recommendations but also the ongoing need for training mill personnel and the requirement that the system be maintainable and expandable without the constant involvement of the developers. PITCH EXPERT is now accessible by modem, and the savings achieved through use of the system covered the development costs within six months of release.

PITCH EXPERT is a large knowledge-based system that analyzes and diagnoses problems in kraft pulp mills that are related to pitch deposition and pitch dirt. By making scientific expertise directly available to mill engineers, the system augments their problem-solving capability and relieves the human expert of much of the burden of routine problem solving and training. The development and delivery of this system also served as a case study to determine the usefulness and relevance of knowledge-based technology in solving real problems in the pulp and paper industry.

PITCH EXPERT is now serving as a significant productivity aid to industry. At the time of writing, there were already 36 mills, representing over 20 companies, registered as active users. Current estimates (see Performance) are that these mills will save approximately $22 million a year by using PITCH EXPERT. Its performance has been documented in an initial 12-month study as well as by records of more recent use. To our knowledge,

PITCH EXPERT is the largest deployed knowledge-based system anywhere in the pulp and paper industry. The system can now be used directly by modem by mill process engineers.

Although previous publications (Kowalski and Gauvin 1992a, 1992b; Kowalski and Lebensold 1989) have reported on the promise and potential of PITCH EXPERT, the past 18 months have seen this promise fulfilled, and the proof of its positive impact, obtained from this initial 2-year period of industrial use, constitutes a main area of focus for this article. In addition, a retrospective analysis of the underlying reasons for success is provided in more depth than was previously possible. Many factors contributing to the success of the project could be assessed only after the system had been used for a significant length of time (more than a year). This article highlights these factors with (1) a quantification of actual savings, rather than mere projections of expected results; (2) a determination of which parts of the system design proved to be most useful or even crucial to its success and why; (3) an analysis of how useful and beneficial the aspects of the design related to system maintenance and update were over an extended period of time; and (4) an examination of the magnitude of the benefits that were derived from the addition of extra capabilities to the system beyond those that normally form part of a knowledge base (that is, conflicting information identification, compact explanation capabilities, extensive metaknowledge).

The successful development and delivery of PITCH EXPERT were largely the result of certain important design decisions. In particular, the structures and mechanisms used to build the system were specialized and customized to fit the specific needs of a system operating in the

… kraft mill pitch problems are estimated to cost, on average, several million dollars each year for each mill in North America …

real-life pulp-mill environment. Moreover, the system was designed from the beginning to be maintained on an ongoing basis without the involvement of senior knowledge engineers.

Kraft Pulp Mills

In the manufacture of paper, wood is first pulped to separate its fibers. One of the predominant pulp processes is done in a kraft pulp mill and consists of cooking wood chips at elevated temperature and pressure in the presence of certain chemicals (alkali and sulfide), washing the resultant brown pulp, bleaching to make the pulp white, and drying the pulp for shipment to a paper mill.

Pitch

Pitch, or wood resin, is the material in wood that is insoluble in water but soluble in organic solvents. It usually makes up 14 percent of the weight of wood after the bark is removed and is often a sticky material. With this gluelike material passing continually through a kraft mill, it is not surprising that under some circumstances, a certain amount deposits on the surfaces of the process equipment.

In addition to the pitch, other depositable materials can be present, such as hard-water soaps and defoamer components. In addition, other materials have a tendency to become entrained in deposits and increase their rate of growth. These materials include calcium carbonate from poor white-liquor clarification, carbon particles from fly ash or green-liquor dregs, bark particles, and sand introduced with wood chips. Although they contain many materials besides wood resin, the deposits are still usually referred to as *pitch deposits*.

Kraft Mill Pitch Problems

Kraft mill pitch problems can take a number of forms. In the most frequently encountered problem, pitch is partially liberated from the wood during pulping and tends to codeposit with other materials on the surfaces of the process equipment. These deposits grow in thickness until they reach a size at which they break away from their surfaces of attachment. When they break away, the chunks of deposit are carried with the pulp and are broken up by the pulp agitators and pump impellers. The result is small dirt particles in the final product that can result in the sale of the pulp at a discounted price or even the loss of a customer. Other pitch problems

include the plugging of screens and cleaners, deposits on the pulp machine, sticking problems on press rolls, and excessive resin.

Economic Significance

It is difficult to place an exact figure on the cost of pitch problems, even for a given mill. However, kraft mill pitch problems are estimated to cost, on average, several million dollars each year for each mill in North America; in Canada alone, this figure translates to $80 million a year.

Components of the cost include sale of off-grade pulp contaminated with pitch dirt, premature replacement of machine clothing, time lost for cleanups, and the cost of additives to control the problem. The additives can include detergents and solvents for cleaning surfaces, pitch dispersants for stabilizing the resin in suspension, and talc for detackifying the deposits. Even if a mill is successfully controlling pitch with the use of these additives, substantial savings can often be achieved if the addition rate can be reduced.

Technology of Pitch Control

Many factors are important in determining whether a kraft pulp mill has pitch problems. These factors include species and storage time of the wood being processed, degree of bark removal, purity and concentration of the cooking liquor, thoroughness of pulp washing, use of foam-control agents, process pH, water hardness, temperature of the process, and the use of additives for pitch control.

In solving pitch problems, chemical analysis of pitch deposits is often an important key. Numerous analytic chemical methods can determine the composition of the deposits. In a troubleshooting situation, this analysis leads to a better knowledge of what is depositing and usually points to a course of action to solve the problem. Mill personnel can also perform a number of diagnostic tests on the pulp suspension to obtain useful information for determining how to solve a pitch problem at a mill.

Why Knowledge-Based Technology Was Required

All this information must be considered when attempting to diagnose a pitch problem in a kraft pulp mill. Furthermore, the information and the conclusions drawn from it carry various degrees of certainty, reliability, and subjectivity. Missing, incomplete, and even inaccurate data are also a fact of life in pulp mills, as is the use of ambiguous and synonymous terminology. The expert must

consider all these factors when determining further questions and making recommendations. A conventional programming language would clearly be incapable of describing this situation adequately.

Furthermore, the questioning strategy that one must follow in diagnosing pitch problems is dynamic in nature. One must consider all the information to date before determining which question is the most appropriate one to be asked next. This situation is especially true because a given piece of information (such as species of wood, seasoning time, pH values) might be relevant to several possible causes of pitch deposition (Kowalski and Gauvin 1992a, 1992b).

The mechanism for handling the questioning logic of the system had to be flexible enough to allow for the easy addition of more sophisticated features at a future date, without the need to redesign a part of the system. Such features could include an answer-retraction capability (Reiter and de Kleer 1987) as well as a mixed-initiative mode of use that would allow the user to enter information immediately rather than wait for the appropriate question (Kowalski and Gauvin 1992a; Kowalski and Lebensold 1989).

PITCH EXPERT was designed to meet all these functional requirements and satisfy the previously mentioned constraints. However, it was equally important for the system to be easily maintainable and updatable so that as new methods of pitch control are discovered, the expertise could quickly and efficiently be incorporated into the system. In this way, state-of-the-art expertise could always be accessible to mills across the country. Similarly, PITCH EXPERT could be kept up to date with regard to changes in kraft bleaching technology. This requirement was carefully considered from the initial design to the final implemented version, and in retrospect, the high priority accorded this issue largely contributed to the system's success (Kowalski and Muise 1990).

Finally, one additional requirement necessitated building the system as a knowledge base. PITCH EXPERT had to be able to explain itself clearly and concisely in relation to every aspect of user interaction, including questions it might pose and conclusions and recommendations it might provide. This requirement was important for two reasons.

First, before a mill manager will authorize the implementation of a recommendation that might result in a significant expenditure (such as adding or increasing the feed rate of a chemical additive), he or she must be confi-

dent that such a measure will have a positive effect on mill production. This confidence can be gained only if this same mill manager is able to obtain a logical justification for the action from the system. Furthermore, our experience has shown that such a justification must be short and concise, so that the user can grasp the essential reasoning without getting lost in a multitude of rules and facts.

Second, PITCH EXPERT was to be used as a training tool. In this respect, allowing the users to ask the system questions—such as why is this question important, what do you mean, how do I find the answers to this question, and why do you make this recommendation or conclusion—is extremely important. Also, offering this same user the ability to specify the level of detail with which the answer is displayed further enhanced this use of the system.

The Need for Training of Mill Personnel

The experience of the domain expert, Lawrence Allen, in solving kraft mill pitch problems now spans over 20 years. Although there are still some grey areas in our understanding of the various interacting phenomena of pitch deposition and dirt formation, experience suggests that we have sufficient practical knowledge to provide expertise for solving most pitch problems in kraft mills. Because of this existing knowledge and the complexity of pitch problems, a considerable amount of consulting time has been spent helping mill personnel with pitch problems. At the mills, over time, certain personnel learn through reading, consulting sessions, and experience about how to avoid pitch problems. Nevertheless, one reality of life in a kraft mill is that over the years, there is a constant influx of new engineers and technical staff members. Thus, there is a constant need for training. No textbooks currently deal with this subject.

Implementation

PITCH EXPERT was implemented using the ART (automated reasoning tool) expert system development tool on top of the Lisp programming language. The hardware platform was a SPARCSTATION 1 with 32 megabytes (MB) of random-access memory and 500 MB of hard-disk storage. The system contains approximately 1200 rules and 3000 schemata. In addition, about 200 functions and *demons* (procedures attached to knowledge structures) are present in the system.

… one reality of life in a kraft mill is that over the years, there is a constant influx of new engineers and technical staff members. Thus, there is a constant need for training.
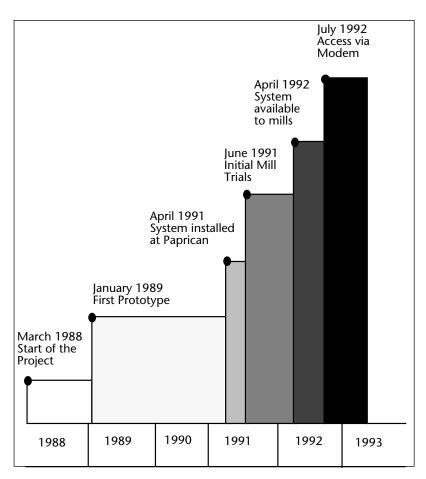
Figure 1. Important Dates in the PITCH EXPERT *Project.*

PITCH EXPERT was built over slightly more than four years, from March 1988 to July 1992, when the completed system was transferred to the Pulp and Paper Research Institute of Canada (Paprican). Figure 1 is a time line of important dates in the evolution of the PITCH EXPERT project. During this time, some 21 person-years of effort were expended on the project (15 for scientific and managerial personnel plus 6 for support staff). This time represents the lion's share (85 percent) of the approximately $2.5 million project cost ($900,000 million direct cost). The remaining expenditures were mainly for project-related equipment and software.

A smaller ongoing effort is also foreseen for system maintenance. This project is expected to represent about 1-1/2 person-years for each of the next few years at an anticipated cost of about $140,000 a year. Figure 2 shows the division of PITCH EXPERT into its four main reasoning modules: the domain module, the NTK module, the evidential reasoner, and the library module. Each module incorporates one or more specific reasoning strategies and uses a set of specialized knowledge representation structures.

The NTK Module

An expert's knowledge includes not only facts and reasoning but also the questions to ask and their order and circumstances. The goal is to get to the best and most complete understanding of the problem as possible (to make good recommendations) asking the
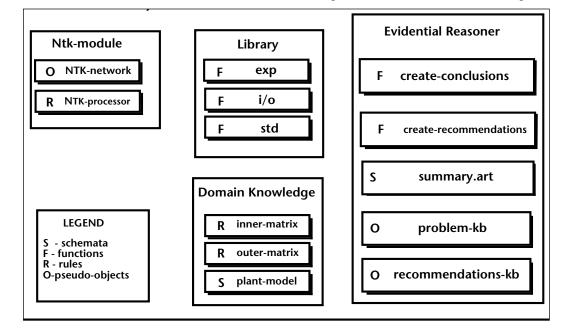

Figure 2. Major Components of PITCH EXPERT.

```
(defschema ntk-126
IS-A: ntk-question
TEXT:
HAS-BEEN-ASKED:
CAN-BE-ASKED:
ANSWER-CONSTRAINT-TYPE:
ANSWER:
SAVED-ANSWER:
COMES-FROM:
LEADS TO:
WHAT-DO-YOU-MEAN:
HOW-DO-I-FIND-OUT:
WHY-ARE-YOU-ASKING:
```

Figure 3. Typical NTK Question Description.

fewest possible questions (to avoid wasting time and effort). To complicate matters further, the answers to earlier questions can change the relevance or importance of later ones.

NTK stands for "need to know" and refers to the function of this module: to determine and implement dynamically the best questioning strategies to obtain information from the user, as needed, for the reasoning.

The NTK module is composed of an NTK network and an NTK processor. The module consists of several hundred nodes connected by a set of specialized relations to form a complex network (the NTK network). Each node represents either a question (NTK), a fact that is inferred directly from the answer to a question (NTKF), or a *normalized* fact (that is, a relative value) that is inferred through a demon from an NTKF (NTKIF). A typical question description is shown in figure 3.

These nodes are connected by multiple links to form a bidirectional network. In turn, these links combine to form two separate threads of reasoning running through the network. The first thread involves *comes-from links,* which trace back the questioning path or paths that can produce a given piece of information (that is, fact). The second thread, which operates in the opposing direction, can be composed of several types of links depending on the nature of the two nodes linked together. More specifically, a particular link can specify which answer is required to produce a given fact, which fact is required to make a given question relevant, or which fact and value set is required to produce a normalized value.

The NTK processor accepts a request from the diagnostic module to search out a given piece of information and then navigates through the NTK network trying to find the information. Figures 4 and 5 illustrate the
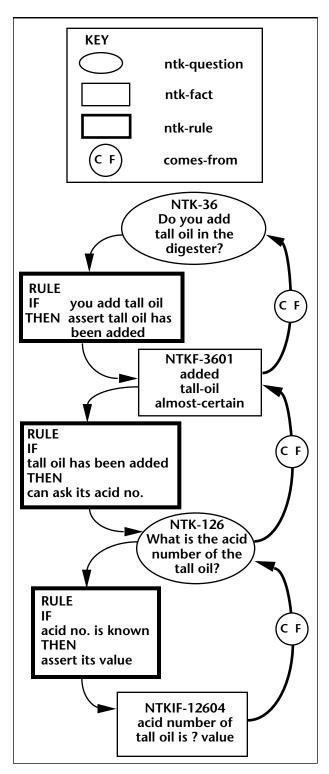

Figure 4. Operation of Reasoning in the NTK Network.

KEY

☐   **ntk-fact**

▣   **comes from/ leads to relation**

⬭ D   **deamon**

→   **assertion of one of possible NTKIF's**

NTKF-12601 acid number of tall oil is ?value

ASSERTS ONE OF FACTS BELOW

D

comes-from

NTKIF-12602 quality of tall oil is very good

NTKIF-12603 quality of tall oil is good

NTKIF-12604 quality of tall oil is bad

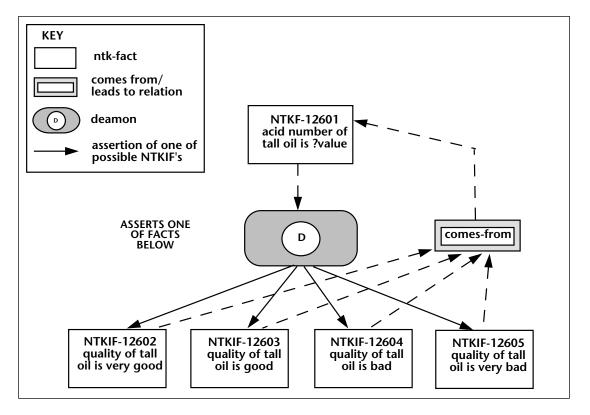**NTKIF-12605 quality of tall oil is very bad**

Figure 5. Operation of a Typical Demon.

operation of the NTK network.

This separation of questioning logic and control into a distinct module (Hughes 1987) and the use of a network of object-type structures (as opposed to the implementation of the logic completely in the form of rules) has greatly simplified the task of maintenance and development (Kowalski and Lebensold 1989). As new questions were acquired from the expert, they were added to the system with relatively little work; all we were required to learn was the appropriate position of the question in the NTK network, the relevant facts that could be derived from answers to the question, and the relevant connections between these facts and other facts already existing in the network.

Questioning paths or strategies can also be changed easily because these changes need to be reflected in only one part of the system (the NTK network). This easy mode of modification was especially important for PITCH EXPERT because many questions and the relevant information derived from their answers play an important role in more than one cause of pitch deposition. As a result, the answers to questions related to one cause of pitch problems can dynamically affect the questioning path for a second cause.

In the two years that have passed since PITCH EXPERT was delivered, the NTK network has undergone almost continuous modification, aimed at improving the questioning strategies of the system. These modifications have been performed exclusively by the Paprican maintenance team (as opposed to the senior knowledge engineers), always using the NTK diagrams and method; this maintenance approach has proven itself to be straightforward and simple. Clearly, it is one of the aspects of system design that is responsible for the continuous and smooth improvement in the system's performance since its delivery.

The nodes of the NTK network also offer an ideal centralized repository for information on the status of each question (that is, has it been asked, can it be asked, if asked what was the answer, and so on). Without these structures, it would be necessary to have a set of complex and difficult-to-manage rules to determine when and if it was necessary and meaningful to ask a given question.

Furthermore, these NTK nodes provide an easy attachment point for various kinds of deep knowledge and even metaknowledge, for example, online help facilities (what do you mean, why are you asking, how do I find
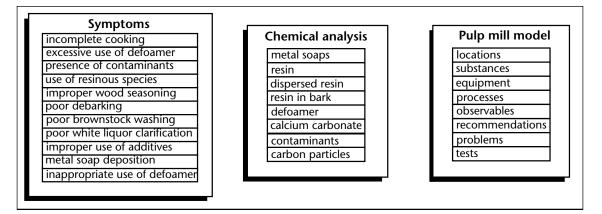
Symptoms

incomplete cooking
excessive use of defoamer
presence of contaminants
use of resinous species
improper wood seasoning
poor debarking
poor brownstock washing
poor white liquor clarification
improper use of additives
metal soap deposition
inappropriate use of defoamer

Chemical analysis

metal soaps
resin
dispersed resin
resin in bark
defoamer
calcium carbonate
contaminants
carbon particles

Pulp mill model

locations
substances
equipment
processes
observables
recommendations
problems
tests

Figure 6. The Modules Containing PITCH EXPERT's Domain Knowledge.

the answer, what are acceptable answers, and so on). These help facilities have been found to add greatly to the system's ease of use, thus gaining the confidence of the pulp-mill personnel. They have also made it possible to use PITCH EXPERT as a training tool, which is important because of high turnover in the mills. This training capability is seen by the users (that is, mill personnel) as a feature of great importance and is one of the reasons for its enthusiastic acceptance. Our analysis of the use of the system in its first two years shows a common trend among initial users to make heavy use of this feature. In fact, many first-time users have run a complete session of well over 100 questions asking why for each one. Such an exercise has allowed them to gain a deeper understanding of the true complexities of pitch-deposition problems.

The NTK network also allows PITCH EXPERT great flexibility in handling incomplete information. When a question is answered with *unknown,* alternate questioning paths are already explicitly laid out in the NTK network and can quickly and efficiently be activated.

Other potential features that can be implemented easily and efficiently using this approach include (1) multilingual operation and more advanced help and explanation capabilities; (2) answer-retraction capability, where the truth maintenance issue is not affected, but the ability of the system to read-just the questioning strategy automatically and dynamically would be of great help; and (3) mixed-initiative mode capability. The installation of such a capability would be simplified greatly by the already-existing model of the relationship of each question-answer set to all its immediately related facts and questions.

Domain Module

The *domain module* contains the actual domain knowledge relating to pitch problems. It consists of two major components: the mill model and the diagnostic module.

Mill Model   As shown in figure 6, the *mill model* consists of five semantic networks, each of which describes a typical kraft pulp mill from a particular perspective (locations, substances, equipment, processes, and observables). These five networks are connected, as appropriate, at various points. Together they provide a model of the pulp mill with both a lateral and a hierarchical frame of reference that is critical to the proper and efficient functioning of the rules in the diagnostic module. In addition, three semantic networks relating to problems, tests, and recommendations are used in the overall reasoning process.

The eight networks combined form a model consisting of over 2000 nodes (implemented as schemata) and over a dozen relations, most of them customized and complex. Figure 7 shows a small portion of one of these networks, describing the locations in a pulp mill. Although not a model-based system as formally defined (Winston and Shell 1990; de Kleer and Williams 1989), the mill model provides some of the same advantages (de Kleer 1991) by supporting and supplementing the diagnostic heuristics of the system.

Although the use of such a structure allows the option of customizing the models to match any particular mill exactly, our experience to date has shown the need for such a fit to be less than urgent. The generic model of the mill has served well in approximating the configuration of the actual mills that have used the system. It remains to be seen if the customization option becomes more important as use of the system grows, and kraft process operations change.

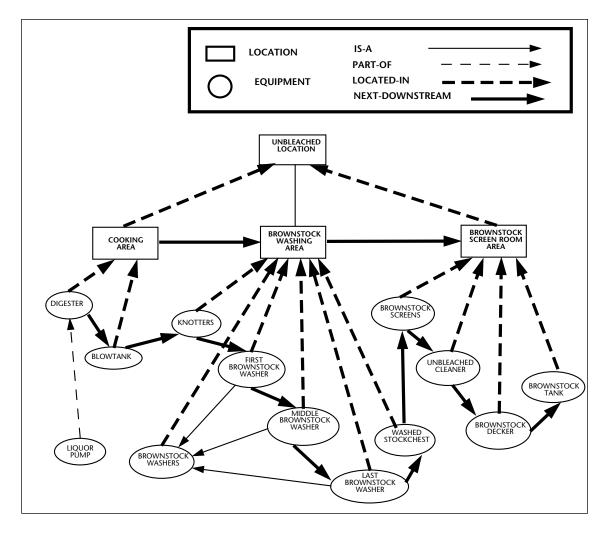**Diagnostic Module**   The *diagnostic module* implements the diagnostic reasoning strate-

Figure 7. Portion of the Pulp-Mill Model.

gies related to pitch problems. As shown in figure 6, it consists of 19 components, each of which addresses one particular aspect of pitch diagnosis, either symptoms of pitch problems or chemical analyses performed to diagnose pitch problems. Each component uses five types of structure to produce its diagnoses:

First, *NTK facts*, which were described in the preceding section, are mentioned here only because they serve to satisfy the conditions of diagnostic rules.

Second, *metaasserts* are commands that send a message, when appropriate, to the NTK module, instructing it to track down a particular piece of information that is needed for the diagnosis. Note that the sending of this message is the only involvement of the diagnostic module with the questioning strat-

egy (and is comparable to message passing for encapsulated objects).

Third, *inferred facts* are created from the diagnostic rules when the conditions of the rules are satisfied by NTK facts and previously inferred facts; they serve, in turn, to satisfy other diagnostic rules (creating the inference chains of the system).

Fourth, *evidence structures* are created by the diagnostic rules. They represent conclusions about the likelihood of a particular cause of pitch deposition based on the certainty and reliability of one source of information. Naturally, for any given cause of pitch problems, there might be several sources of information, each of which produces a separate piece of evidence. When all the available information has produced all the relevant pieces of evidence, the evidential reasoner (see next
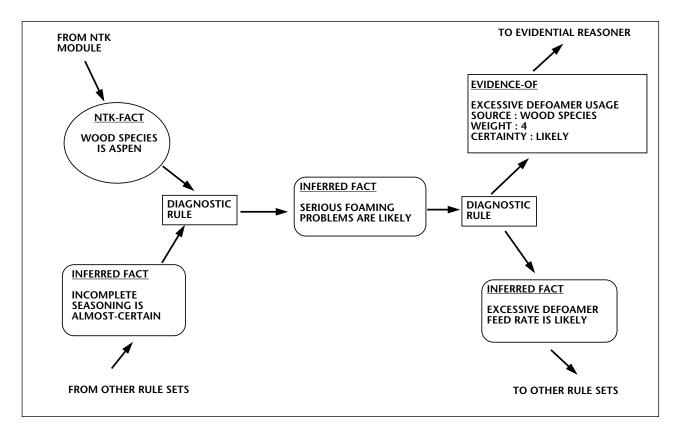
Figure 8. A Chain of Inference Leading to a Conclusion.

subsection) collates, sorts, and prioritizes them to reach its final conclusions.

Fifth, *rules* are standard, diagnostic if-then rules that use facts (both NTK and inferred) to create other inferred facts as well as pieces of evidence.

Figure 8 shows a portion of the rule set pertaining to the problem of excessive defoamer use. This example shows how the various knowledge representation structures and modules of PITCH EXPERT work together in a coordinated fashion.

A message is sent to the NTK module indicating the need to know the species of wood being used. If the species is aspen, and it has been inferred in another rule set that the wood is incompletely seasoned, then it will be inferred that serious foaming problems are likely to occur. If so, a second rule will infer that it is likely that excessive amounts of defoamer are being used. This same rule also creates an instance of evidence specifying the appropriate weight and certainty. Both of the inferred facts created in this rule set can be used to satisfy rules in other rule sets. The instance of evidence (along with all other instances of evidence) is processed at the appropriate time by the evidential reasoner.

This example clearly shows how this line of reasoning, using the species of wood and information about its seasoning, is unaffected by any new lines of reasoning using other information. In this way, maintenance requirements related to the addition of new lines of reasoning are kept to a minimum.

Figure 9 diagrams a second line of reasoning from the same rule set, which uses information from the NTK module related to the feed rate and the location at which defoamer was added, as specified by the user, to produce an inferred fact and an instance of evidence with the appropriate parameters.

Evidential Reasoner Module

Because of the functional requirements of the system (see Why Knowledge-Based Technology Was Required), it was clear early on in the design of PITCH EXPERT that a standard approach to dealing with uncertainties, such as the certainty factors of MYCIN or the implementation of the Dempster-Schafer theorem, would have been inadequate (Buchanan and Shortliffe 1985). Instead, a customized evidence-handling strategy was developed.

Figure 10 shows a typical piece of evidence, which consists of a specification of an associated cause of excessive pitch deposition, the source of the information (the specific piece
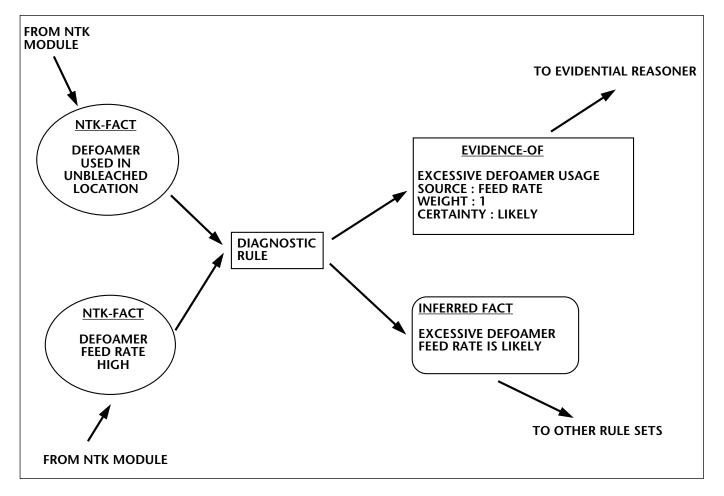
Figure 9. A Second Chain of Inference Leading to a Similar Conclusion.

```
Schema 2311.5.3

(instance of              excessive-defoamer-usage        )
(source                   defoamer-feed-rate              )
(weight                   1                               )
(certainty                unlikely                        )
```

Figure 10. Instance of Evidence.

or pieces of information obtained from the user that led to the conclusion of the evidence), a weight representing the reliability of the source of information from which the evidence was derived, and a value specifying the certainty associated with the conclusion.

Each instance of this evidence structure serves as a central storage place for information pertaining to one particular line of reasoning. This compact storage of information offers three advantages:

First, it allows the system to treat the conclusions of each line of reasoning separately.

This approach is important because it keeps the rules simple and easy to maintain. New lines of reasoning can easily be added, creating new pieces of evidence, without the need to modify existing lines of reasoning. Since delivery, many new lines of reasoning have been added in this manner, without any resulting errors in the existing lines of reasoning. The absence of errors is the result of separate handling, collation, and analysis by the evidential reasoner of the collection of evidence (as is explained later).

Second, it allows a flexible and concise system of explanations and justifications (Millet 1989) to be offered to the user (as explained later). This ability to present explanations has been recognized as a critical component in more recently built knowledge-based systems (Cawsey 1991; Suthers, Woolfe, and Cornell 1992) and was a key factor in gaining the confidence of mill personnel. A large majority of the sessions run by the mills make heavy use of this facility. Interestingly enough, the users have shown little interest in the facility for backtracking beyond these concise explana-

tions to view the specific facts and rules composing the chain of logic for a given conclusion.

Third, the creation of a collection of evidence for each possible cause of pitch deposition has made possible the development of a customized approach to the issues of uncertainty, reliability, and even accuracy of the information. By explicitly separating out the process of combining conclusions from different sources of information, PITCH EXPERT can perform a sophisticated analysis and comparison of all the individual pieces of evidence. This approach greatly enhances its ability to arrive at highly accurate final conclusions and also flag critical input inconsistencies.

The evidential reasoner gathers up all the instances of evidence, then sorts, collates, and analyzes them, finally presenting them to the user in a variety of ways. The reasoner includes a network of nonencapsulated objects that represent the set of problems that can lead to pitch deposition, a network of nonencapsulated objects that represent the set of possible recommendations for solving problems leading to excessive pitch deposition, and a set of nodes representing the various tests that can be performed on either pitch deposits or pulp and the problems these tests shed light on.

The evidential reasoner sorts evidence by problem. Evidence is then divided on the basis of whether it supports the existence of a problem (positive evidence) or refutes it (negative evidence). Then the evidence both for and against is summed up, taking into account the certainty and reliability of each source of information in a manner specific to each problem, to produce final positive and negative weights. These two weights are then combined to produce a final conclusion for the particular problem. In addition, if both the positive evidence and the negative evidence are strong, then the system flags this inconsistency as important, informs the user that inaccurate information has been provided, and advises the user to recheck the relevant information. This feature has proven to be essential for dealing with the real-life inaccuracies that exist in the mills.

The summary conclusions (indicating that one or more problems have been found) activate the appropriate nodes in the problem network, which, in turn, activates the appropriate nodes in the recommendation network, as well as the chemical test network (indicating what should be done about the problem(s)), at which point the relevant output is displayed. Such an approach makes it

PITCH EXPERT has deduced that the following has occurred with the certainty shown below

> very strong
>> inconsistent suspended solids concentration
>> poor baking

> strong
>> poor white-liquor clarification

> very weak
>> excessive defoamer use ← very important conflict here

PITCH EXPERT has deduced that the following has not occurred with the certainty shown below

> very strong
>> metal soap deposition in the pulp machine area

> standard
>> poor brownstock washing

Figure 11. Typical List of Prioritized Conclusions.

easy to offer the user the chance to peruse in more detail any conclusion or line of reasoning. Such a flexible presentation style has convincingly proven its worth because virtually all real sessions with the system result in several problems being identified, with several degrees of certainty. It has been observed that users often like to peruse many of the conclusions and recommendations in more detail, often coming back to look at certain information repeatedly.

As shown in figure 11, the user is first offered a concise list of all the conclusions drawn by the system, prioritized by strength of supporting evidence: The problems most likely to be causing the excessive pitch deposition are listed first, and those least likely are listed last. In addition, any conclusion for which a serious conflict of information exists is tagged with a warning message.

At this point, the user can request a more detailed explanation of any or all of these conclusions. This feature is especially useful in obtaining a better understanding of why important conflicts exist. An example is shown in figure 12. In this example, two pieces of evidence relating to the problem of excessive defoamer use have been created by the system. Each piece of evidence is derived from a particular source that is considered reliable, specifically "test results" and the "defoamer feed rate."

The result of chemical tests performed on the deposit show it to contain 29.8-percent

User input=> Why Excessive-defoamer-usage
   This was based on the following information:

Very Reliable
   Test-Results, specifically the 29.8% defoamer in the dry deposits, indicates that it is almost certain that it has occurred.

Very Reliable
   Defoamer-feed-rate, specifically the 0.2 Kg/ton, indicates that it is almost-impossible that it has occurred.

Based on this information PITCH EXPERT can conclude that it has occurred with a very-strong confidence but based on this information PITCH EXPERT can also conclude that it has not occurred with a very-strong confidence.

Figure 12. Typical Explanation of Conflicting Evidence.

PITCH EXPERT makes the following recommendations with the indicated levels of importance

very strong recommendations

Improve the efficiency of the brownstock washers. This will lower the concentrations of dispersed wood resin-dissolved soaps-defoamer-calcium carbonate particles and other depositable materials.
   This was recommended for correcting the following problem(s):
         poor-brownstocking-washing
         metal-soap-deposition-in-brownstock screen-room

   hit return to continue or why for more details

why

choose one of the following problems:

1 poor-brownstock-washing
2 metal-soap-deposition-in-brownstock-screen-room
3 next recommendation

Figure 13. Typical Recommendations Display.

defoamer. Such a high content of defoamer in the deposit indicates that almost certainly an excessive amount of defoamer must have been added to the pulp during production. However, the feed rate of the defoamer (0.2 kilograms a ton) is considered extremely low, leading to the conclusion that it is almost impossible that an excessive amount of defoamer was used.

These two lines of reasoning conflict strongly with one another and lead PITCH EXPERT to identify this situation as physically impossible, indicating that at least one of the lines of reasoning stems from incorrect data input to the system. Having identified such a conflict, PITCH EXPERT can now relay a warning to the user to recheck the information.

Such inaccurate data could have been caused by many factors, but their presence in any given session is a realistic possibility given the number of complex processes that are running simultaneously in a pulp mill and the variable quality of available process information.

If the user so desires, an even more detailed step-by-step backtrace of the chain of reasoning is available as well, with the relevant rules and facts at each step in the reasoning process shown in full detail. A similar display can be obtained for the recommendations, as shown in figure 13.

Library Module

The *library module* contains a collection of Lisp functions and ART demons that perform tasks such as normalizing measurements, testing value constraints, and collating and analyzing pieces of evidence. These functions and demons ensured a custom fit between the system and the knowledge and greatly enhanced the system's performance.

Maintenance

As mentioned earlier, the design of PITCH EXPERT explicitly addressed the issue of ease of maintenance. Although PITCH EXPERT is now in full operation, changes continue to be made to the knowledge and rule bases, to both make the knowledge representation more complete and accurate and cover new contingencies. These changes mainly affect the NTK module, rules, and functions, with occasional minor changes to the mill model.

Most of the changes now being made involve the addition of new knowledge to improve or extend system performance (Allen and Kowalski 1992). In some cases, the new knowledge is meant to fill a specific gap in the knowledge base; in other cases, it is a matter of adding knowledge as it becomes available. In other words, as research identifies new information that is relevant to solving pitch problems in kraft pulp mills, it is incorporated into PITCH EXPERT. In this way, the system is kept up to date, and obsolescence is avoided. At this time, the maintenance is performed on an incremental scale, and there are no immediate plans to add large amounts of new knowledge to the system.

Specialized Knowledge Structures

Specialized structures have played a major role in ensuring the maintainability of the system. Foremost among these structures are the NTK structures (see NTK Module), the evidence construct (see Evidential Reasoner Module), and customized relations such as that described in the following example.

Example of Customized Relations and Maintenance: The specialized next-down-stream relation is used extensively in PITCH EXPERT to build the model of pulp-mill locations (see Mill Model). This relation, made possible by the sophisticated relation-building capabilities of the ART shell, describes the relative position of two locations in a mill in terms of pulp or water flow between them. The resulting directionally linked chains can be used to make locational comparisons, even between different levels of abstraction. This factor also proved to be critical in the success of the system for three reasons:

First, the relation allowed easy insertion of new or additional locations just by modifying a few links. No other rules or functions needed to be changed. In this way, the pulp-mill model could be constantly and easily refined.

Second, because the relation is transitive (through a more general downstream relation), diagnostic rules can use it to compare locations to determine which is downstream of the other. For example, consider the location of a pitch deposit and the location at which talc was added. If the talc was added downstream of the deposit, it is unlikely that the talc contributed to the problem (because the talc is unlikely to swim upstream). If, however, the talc was added upstream of the deposit, it is possible that the talc addition influenced the deposition problem. If the talc was added directly upstream of the deposit, then it is probable that the addition of talc influenced the deposition problem.

Third, the relation allows comparison of locations at different levels of specificity. For example, the first brownstock washer is downstream of the cooking area as a whole because it is downstream of each component of the cooking area. This sort of concept is simple for human beings but can be extremely difficult and complex to represent effectively in a computer program. The ability to incorporate this knowledge into a single customized relation speaks volumes about the utility of powerful knowledge base shells, and the relevance of such a feature to ease of maintenance is self-evident.

Knowledge Representations

A second strategy that has facilitated both development and maintenance is the use of three coordinated pools of knowledge: a database, a collection of graphic blueprints, and the ART source code. Each pool contains the exact same information but in a form that offers particular benefits.

In the database, each instance of knowledge is stored as a record. These records can be accessed by queries and sorts based on various indexes to identify, for example, all the places where a fact, fact element, or class of facts appears. This approach makes it possible to change or augment the knowledge base quickly and without errors or omissions.

The graphic blueprints contain the same knowledge as the database but in a form that visually highlights links between knowledge structures such as rules, questions, and evidence. Both the domain expert and the system developers have found it useful to be able to look at the diagrams and grasp the essence of the logic without having to wade through masses of source code and system commands. Each blueprint can contain a large amount of information of various types, ranging from natural language versions of the rules to specialized metaknowledge commands. Different people can focus on the information of current interest to them, yet all work with a single diagram, which helps to ensure consistency from concept to design to implementation and, finally, through testing, debugging, and maintenance.

The code itself serves both as the ART implementation of the system and as a repository of knowledge. Code for each rule, schema, and function includes a history of when each was created and modified, who created it, and why. Furthermore, each function and demon contains extensive documentation in a standard form, explaining its purpose and approach and all parameters used.

Because the knowledge in PITCH EXPERT exists in three forms (diagrams, databases, and code) that are not automatically integrated, procedures were established to make sure that consistency was maintained among these three repositories of knowledge. Each instance of each structure (except for the rules) must be entered into the database. Each instance of each structure, including the rules, must of course be coded in ART. To create the blueprints, printouts of individual structures from the database or the source code files are physically pasted onto the diagrams and manually linked together. Further checking includes, for example, obtaining a

printout from the database of which facts exist as input (left-hand side of a rule) and which as output (right-hand side of a rule) and checking this printout against the diagrams. All this checking virtually assures that the blueprints represent the actual state of both the database and the code.

The success of this approach is attested to by the significant improvements in accuracy and completeness (as outlined in Performance) that have been achieved by changes made by Paprican personnel with minimal supervision by the senior designers. Considering the day-to-day use that is still being made of this maintenance system to maintain and upgrade the knowledge base, its importance to the success of the project cannot be overstated.

This maintenance mechanism has withstood the test of real-life use and still serves today as the control and self-checking mechanism that has allowed an ongoing dynamic environment to be created in which constant modification of the knowledge serves as one of the system's attractions rather than a detraction. Response from the users has shown that the ability to have state-of-the art expertise always available is a strong source of motivation to use it.

The perception and development of PITCH EXPERT as an ongoing evolving system, rather than a static entity, was a key factor in its success. The provision of such a maintenance capability, in whatever form proves workable, should be given serious consideration in the development of any major industrial knowledge base.

Use of PITCH EXPERT

PITCH EXPERT resides on a SPARC workstation at Paprican and is accessed by modem. The system operates in a question-and-answer mode, posing questions to the user one at a time and accepting typed answers until it has all the information it needs. At important decision points during a run, PITCH EXPERT also displays numbered menus of available options.

Each answer is checked for acceptability as soon as it is entered. The criterion can take the form of a list of specific acceptable answers, a class specification, or a requirement that the answer be a number. If an unacceptable answer is entered, PITCH EXPERT displays a message and then asks the question again.

Any question can also be answered with one of PITCH EXPERT's keywords. These keywords activate special features such as online help, access to intermediate conclusions and recommendations, and the saving of an incomplete session for later retrieval.

Missing information is handled with two keywords, *later* and *unknown*. *Later* is used when an answer is expected to be obtainable but is not available right now and *unknown* when no better answer can ever be expected. When answers become available for questions originally answered by *later,* the session can be retrieved and rerun, and these questions are asked again.

Each answer, as it is processed, can create NTKFs and, indirectly, NTKIFs. In turn, the NTKFs and NTKIFs activate more NTKs, or they can lead domain rules to fire, creating inferred facts or pieces of evidence. In either case, the chain of rules eventually leads to the creation of a collection of evidence to be evaluated by the evidential reasoner.

At any point in a PITCH EXPERT session, the conclusions and recommendations that are supported by the evidence asserted to this point can be viewed on the user's screen by use of the *recc* keyword. These conclusions and recommendations are, of course, not final and change during a session to reflect new information as it becomes available.

Three kinds of help are available for questions and can be accessed with keywords. *What* provides a reworded version of the question in case the original was unclear. *How* gives further detail on how to obtain the information required to answer the question. *Why* expands on why it is important to ask this question in the context of pitch control.

Because a typical session can take an hour or more, a facility has been provided (with the *save* and *restart* keywords) to save the current session at any point to be picked up again at a later time. A common use of this feature would be to store a session in which *later* answers were given to one or more questions, with a view toward retrieving and continuing the session when answers to these questions became available.

When PITCH EXPERT has an answer to all the questions it judges necessary to its reasoning, it completes its reasoning process and then offers to display on the screen the conclusions it has reached and the recommendations it has made. Conclusions and recommendations are shown one by one, in order of priority, each with its supporting evidence in order of reliability.

Recommendations are displayed in a similar manner. By typing *why* in response to a recommendation display, the user can obtain a summary of the reasoning leading to this recommendation.

PITCH EXPERT can produce a summary report

> *As the system was fine tuned with each new mill test run, the system's performance was continuously improved. At the time of this writing, it stands at 90-percent accurate and 93-percent complete.*

of the session just completed or of any earlier completed session. This report includes all conclusions reached, recommendations made, and conflicts noted. The report is displayed on the screen and incorporated into the transcript of the session.

Performance

There are two fundamental measures of expert system performance: the quality of the expert reasoning and the benefits (economic and other) provided by use of the system.

The most basic requirement of expert system performance is that it reproduce acceptably well the reasoning of the human expert. In the case of PITCH EXPERT, its conclusions and recommendations must be accurate (that is, the same ones the human expert would have come to in the same situation) as well as complete (nothing the human expert would have concluded or recommended is left out).

Initial Mill Trials

To evaluate the benefits of the system, financial and otherwise, a set of initial mill trials were conducted from May 1991 to April 1992. The trial runs were conducted with the cooperation of 13 kraft pulp mills across Canada (Allen and Kowalski 1992; Turney 1992).

During this test period, the sessions were conducted with the aid of an intermediary at Paprican. This person was responsible for manning the telephone and acting as the contact between the mills and the system. When mill staff members called in, they would verbally convey the relevant information about their particular mill to the intermediary, who, in turn, would enter all the information into the system. The corresponding output would be checked for accuracy and completeness by Allen and then would be sent out (with corrections if necessary) to the mill. With the results of each session, modifications were made to the knowledge base to address any errors or inconsistencies that were discovered.

The early sessions achieved an accuracy of 60 percent and a completeness of 70 to 80 percent. As the system was fine tuned with each new mill test run, the system's performance was continuously improved. At the time of this writing, it stands at 90-percent accurate and 93-percent complete. Further progress is expected. All 13 mills were to perform a complete evaluation of the benefits of the system. As of October 1992, three had completed their evaluation, seven more were in progress, and three had suspended the study for various rea-

sons stemming for the most part from severe constraints on resources.

The three completed evaluations show a combined annual savings of almost $1.9 million, for the most part because of an overall reduction in the amount of pitch-contaminated pulp produced as well as a decrease in the quantities of additives used. This figure is especially impressive in light of the fact that it represents savings for just three kraft pulp mills. With more than 40 such mills in Canada alone, these savings almost equal the $2.24 million development cost of PITCH EXPERT to date. The 36 mills already using PITCH EXPERT can expect to save $22.4 million a year.

Discussion of the Three Completed Evaluations

Given the importance that these savings represent, it is worth examining them in greater detail. In the interest of brevity, we focus on those mills whose evaluations of the effects of PITCH EXPERT were complete. For reasons of confidentiality, the company names and locations of the particular mills cannot be specified. They are referred to as mills 1, 2, and 3. They are, however, real and certainly typical kraft pulp mills, and the sessions of PITCH EXPERT that they ran involved real data representing real pitch problems. The cost savings are outlined here, and the results are summarized in table 1.

Case Study of Mill 1: Mill 1 has an annual production rate of 170,000 tons of pulp. The production of off-target (that is, pitch-contaminated) pulp represents a loss of $40 a ton. The defoamer used in this mill costs $0.88 a kilogram.

Before running a diagnosis with PITCH EXPERT, this mill had 5 percent of the pulp it produced contaminated with pitch. In addition, it was using 3.8 kilograms of defoamer for each ton of pulp.

After running a full diagnostic session with PITCH EXPERT, entering all the required information and implementing the resulting recommendations, mill 1 had reduced its off-target pulp to just 1 percent of total production. In addition, the rate of defoamer use was reduced to 2.8 kilograms for each ton of pulp.

Therefore, the annual savings from the reduction in off-target pulp is $272,000. The annual savings from the reduced use of defoamer is $150,000. Thus, the total yearly savings is $420,000.

Case Study of Mill 2: Mill 2 has an annual production rate of 360,000 tons of pulp. The production of off-target pulp represents a loss of $40 a ton. The defoamer used in this mill

The three completed evaluations show a combined annual savings of almost $1.9 million ...

Mill Number	1	2	3
Off target before	5%	2%	N/A
Off target after	1%	0%	N/A
Annual production (tons)	170,00	360,00	450,00
Annual estminated savings	$420,000	$871,000	$593,600

Table 1. Summary of Case Studies.

costs $1.08 a kilogram.

Before running a diagnosis with PITCH EXPERT, mill 2 had 2 percent of the pulp it produced contaminated with pitch. In addition, it was using 4.5 kilograms of defoamer for each ton of pulp.

After running a full diagnostic session with PITCH EXPERT, entering all the required information, and implementing the resulting recommendations, mill 2 completely eliminated the presence of pitch-contaminated pulp. In addition, the rate of defoamer use was reduced to 3.0 kilograms for each ton of pulp.

Therefore, the annual savings from the reduction in off-target pulp is $288,000. The annual savings from the reduced use of defoamer is $583,000. Thus, the total yearly savings is $871,000.

Case Study of Mill 3: Mill 3 is an excellent example of how PITCH EXPERT can help a mill achieve significant cost savings even when pitch deposition is already successfully controlled.

This mill has an annual production rate of 450,000 tons of pulp. The defoamer used in this mill costs $1.00 for each kilogram.

In this mill, there was no significant pitch deposition at the time of system use. However, the rate of defoamer use was 2.7 kilograms for each ton of pulp, and other additives were being used to prevent possible pitch deposition.

After running PITCH EXPERT and implementing the resulting recommendations, mill 3 reduced the defoamer feed rate to 1.5 kilograms for each ton of pulp.

Therefore, the annual savings from the reduced use of defoamer is $540,000. The projected annual savings from the reduction in use of other additives is $36,000. The projected annual savings from the reduction in the need for changes in machine clothing is $17,600. Thus, the total yearly savings in the mill is $593,600.

Modem Access to PITCH EXPERT

As a result of this evaluation exercise, there was an almost universal agreement among the mills involved that the system should be made directly accessible to the mills by modem rather than through a Paprican technician. Mill personnel would then be free to use the system as a training tool by making extensive use of the what, how, and why facilities as well as being able to evaluate numerous what-if scenarios.

However, it was decided not to install PITCH EXPERT at individual mill sites. Given the size and complexity of the system and the fact that it continues to evolve, it was judged that updates would be easier with a single copy stored at Paprican. Having a single copy at a central site also greatly facilitated the restriction of access to authorized users. Direct modem access to the system was implemented in July 1992. What is of particular interest here is that the initiative and enthusiasm for modem access came from the mill personnel, highlighting the enthusiasm with which PITCH EXPERT was received by industry.

Once PITCH EXPERT was made accessible by modem, its use grew quickly. Eight months after its release (March 1993), 80 percent of Canadian kraft pulp mills, representing 90 percent of Canadian production capacity and over 20 companies, had accessed the system. In fact, to this point, acceptance and use of PITCH EXPERT by the mills had been limited only by the necessity for mill personnel to receive introductory training on proper use of the system. With limited personnel available for this task, it has been impossible to provide this training to more than one mill at a time.

Given the rapid growth in the number of users and the limited resources available, it has not been feasible to perform an exhaustive analysis of system use beyond the initial 12-month study. Rather, ongoing day-to-day observations are made of use patterns. Having a centralized system, accessed by modem, has made this observation possible; studying the use of a number of copies of the system installed on site at each mill would have been much more difficult.

Based on these observations, after two years of use, the following patterns have been observed:

First, initial use of the system is usually for

educational purposes rather than troubleshooting, including heavy use of the help facilities to learn why the system asks its questions and why it makes its recommendations as well as what-if exercises using hypothetical data. Such use illustrates to the user just how complex and difficult accurate pitch-problem diagnosis really is. This period lasts anywhere from several weeks to several months.

Second, following the initial exploration of the system, a more methodical use begins with real-life data pertaining to the particular mill. Gathering these data can often take several weeks, depending on the availability of personnel at the mill. As more data are entered into the system, the updated recommendations are usually produced with stronger certainty and are perused in more detail, making heavy use of the explanation facilities.

Third, over the long term, roughly half of the users access the system at least once a month and half less often.

To date, PITCH EXPERT users have all been members of Paprican. As members, they are considered to have paid for development of the system with their industrial membership fees. Making PITCH EXPERT available for use by nonmember companies is now being considered. There has already been strong interest from several countries. Although the fee structure has not yet been determined, it is clear that if nonmember companies are given access to the system, they will be charged for its use.

Other Benefits of PITCH EXPERT

A number of less quantifiable but equally real benefits have also been realized in participating mills. PITCH EXPERT provides training in pitch-control techniques to mill engineers. Many pulp mills are in less desirable remote locations and tend to be staffed by engineers who are inexperienced or new to the industry and have a high turnover rate. The help and explanation facilities are particularly useful here. During development of the system, the expert commented that when he visits mills, he spends much of his time educating mill personnel as well as investigating pitch problems. It is important that PITCH EXPERT be able to perform both these functions.

Using PITCH EXPERT also encourages mill personnel to do more information gathering and testing on a regular basis, thus improving their overall problem-solving capability. The sequence of questions serves to highlight information that should be obtained when solving pitch problems, and the help facilities instruct users on what to look for in the mill and how to perform tests and procedures.

The session transcripts also serve as a record of the problem-solving process. This record can be important because pitch problems in a given mill can be intermittent, and a successful solution procedure can be forgotten by the time the problem recurs.

Conclusions

The PITCH EXPERT project is a success. The system is now used by an ever-growing number of kraft pulp mills from many companies all over Canada. Its possible use by mills in other countries is now being considered; some companies outside Canada have already expressed interest.

The analysis of system performance shows that knowledge-based technology can indeed have a dramatic impact on productivity and costs. The savings realized by mills using PITCH EXPERT and following its recommendations clearly justify the expense of development. Although the continued introduction of this technology into the pulp and paper industry in the form of other knowledge-based systems will have to proceed in a methodical, stepwise manner, it is clear that the success of PITCH EXPERT has realized the first step toward this goal.

From the technical standpoint, it is clear that the success of the project was a result of the philosophy of custom fitting the system to the knowledge (Bouchard et al. 1991) as well as the careful attention that was paid to the issue of ongoing maintenance. This process involved the selection of a large, sophisticated, and powerful shell (in this case ART) followed by customization to achieve a perfect fit with the needs prescribed by the domain knowledge.

Finally, PITCH EXPERT serves to highlight the fact that large and sophisticated expert systems can and do provide distributed and up-to-date expertise in a readily available and accessible fashion, which translates into improved productivity and a more competitive industry.

The major achievements of the PITCH EXPERT project can be summarized as follows:

First, a large knowledge-based system is being regularly used by mill personnel to solve real industrial problems. Few systems in the pulp and paper industry reach this stage of practical usefulness. Mills already using the system can be expected to save a total of $22.4 million a year.

Second, a sophisticated question-asking mechanism (the NTK module) was developed to enable PITCH EXPERT to obtain the informa-

tion it needs without asking unnecessary or irrelevant questions.

Third, a flexible strategy was developed for combining pieces of evidence of various strengths and weights to reach conclusions about the existence and importance of problems.

Fourth, a powerful set of customized relations and knowledge structures was developed for modeling an industrial process.

Fifth, thanks to its maintenance-oriented design and associated strategy, PITCH EXPERT is being successfully maintained with minimal involvement by AI specialists.

Acknowledgments

References

Allen, L. H., and Kowalski, A. 1992. PITCH EXPERT System: A Case Study for a Knowledge-Based System for Pitch Control in the Pulp and Paper Industry. In Proceedings of the Forest Sector 2000 Conference, 98–103. Ottawa, Ontario, Canada: Industry, Science, and Technology Canada.

Bouchard, D. C.; Kowalski, A.; Larin Y.; and Vadas, O. 1991. Acquiring and Structuring Expert Knowledge. In Proceedings of the Tappi Engineering Conference, 81–93. Atlanta, Ga.: Technical Association for the Pulp and Paper Industry.

Buchanan, B., and Shortliffe, E. 1985. *Rule-Based Expert Systems*. Menlo Park, Calif.: Addison Wesley Publishing Co.

Cawsey, A. 1991. Generating Interactive Explanations. In Proceedings of the Ninth National Conference on Artificial Intelligence, 86–91. Menlo Park, Calif.: American Association for Artificial Intelligence.

de Kleer, J. 1991. Focusing on Probable Diagnoses.

In Proceedings of the Ninth National Conference on Artificial Intelligence, 842–848. Menlo Park, Calif.: American Association for Artificial Intelligence.

de Kleer, J., and Williams, B. C. 1989. Diagnosis with Behavioral Modes. In Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, 1324–1330. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Hughes, S. 1987. Domains, Tasks, and Questions: An Approach to Knowledge Analysis. In Proceedings of the Third International Expert Systems Conference, 37–44. Medford, N.J.: Learned Information Inc.

Kowalski, A., and Gauvin, D. 1992a. PITCH EXPERT—An Engineered Collection of Specialized Knowledge Structures. In Proceedings of the Ninth IEEE Conference on Artificial Intelligence for Applications, 20–27. Washington, D.C.: IEEE Computer Society.

Kowalski, A., and Gauvin, D. 1992b. PITCH EXPERT: A Productivity Improvement Tool for Pulp Production. In Proceedings of the Eighth International Conference on CAD-CAM: Robotics and Factories of the Future, 1854–1866. Metz, France: International Society for Product Enhancement.

Kowalski, A., and Lebensold, L. 1989. A Diagnostic Aid to Pulp Production. In Applications of Artificial Intelligence VII, 857–866. Bellingham, Wash.: International Society for Optical Engineering.

Kowalski, A., and Muise, J. 1990. Toward a Self-Managing Knowledge Base. In Proceedings of the Third International Symposium on Artificial Intelligence, 117–123. Mexico, D.F.: Editorial Limusa.

Millet, C. 1989. A Study of the Knowledge Required for Explanation in Expert Systems. In Proceedings of the Fifth IEEE Conference on Artificial Intelligence Applications, 83–96. Washington, D.C.: IEEE Computer Society.

Reiter, R., and de Kleer, J. 1987. Foundations of Assumption-Based Truth Maintenance Systems. In Proceedings of the Sixth National Conference on Artificial Intelligence, 183–188. Menlo Park, Calif.: American Association for Artificial Intelligence.

Suthers, D.; Woolfe, B.; and Cornell, M. 1992. Steps from Explanation Planning to Model Construction Dialogues. In Proceedings of the Tenth National Conference on Artificial Intelligence, 24–30. Menlo Park, Calif.: American Association for Artificial Intelligence.

Turney, P. 1992. PITCH EXPERT: Canadian AI Success Stories. *Canadian Artificial Intelligence* 30:14–17.

Winston, P. H., and Shell, A. R. 1990. *Artificial Intelligence at M.I.T.: Expanding Frontiers. Model-Based Reasoning: Troubleshooting.* Cambridge, Mass.: The MIT Press.

Allan Kowalski is currently a consultant specializing in knowledge-based systems at the Centre de Recherche Informatique de Montreal in Canada. He received his B.S. and M.S. in computer science from Concordia University in 1984 and 1988, respectively. His area of interest is the strategies and methods of development and the long-term maintenance of large, deployed knowledge bases.

Diana Bouchard is an associate scientist at the Pulp and Paper Research Institute of Canada. She received her B.A. in math and philosophy from McGill University in 1968. She also received an M.S. in geography and an M.S. in computer science from McGill University in 1972 and 1979, respectively. Her area of interest is industrial computer applications, with a special focus on knowledge-based systems.

Lawrence H. Allen is the director of research in chemical sciences at the Pulp and Paper Research Institute of Canada. He received his B.S. in physical sciences at Carlton University in 1965 and a Ph.D. in chemistry from Clarkson University, Potsdam, New York, in 1970. Since 1972, his research has focused on pitch control for both pulp and paper mills, making him one of the paper industry's most sought-after consultants. He is a fellow of the Chemical Institute of Canada and the Technical Association for the Pulp and Paper Industry.

Yves Larin is an associate technical specialist at the Pulp and Paper Research Institute of Canada. He received his B.S. from University of Montreal in 1979. He has spent the last three years working in the field of knowledge-based systems.

Oliver M. Vadas is a senior technical specialist at the Pulp and Paper Research Institute of Canada. He received a B.S. in electrical engineering and an M.S. in computer engineering in 1963 and 1965, respectively, from the University of Technology of Budapest, Hungary. He joined the Pulp and Paper Research Institute of Canada in 1972, where he works on evaluating emerging technologies for industrial applications.

Innovative Applications of Artificial Intelligence 4

Proceedings of the IAA1-92 Conference.

Edited by A. Carlisle Scott & Phillip Klahr.

"Far from being isolated in the halls of academia, AI is finding its way into our lives and saving us time and money—usually when we are most unaware of it. This book is a great eve-opener to AI's possibilities through the work of others, largely in global operations."
-Library Journal

Contents:

- Smart: Support: Management Automated Reasoning Technology for Compaq Customer Service / *Timothy L. Acornand Sherry H. Walden*
- A Knowledge-Based System within a Cooperative Processing Environment / *Dale B. Danilewitz, and Frederick E. Freiheit IV*
- Help Desk: Using AI to Improve Customer Service / *Debra Logan and Jeffrey Kenyon*
- Making Sense of Gigabytes: A System for Knowledge-Based Market Analysis / *Tej Anand and Gary Kahn*
- TPF Dump Analyzer / *R. Greg Arbon, Laurie Atkinson, James Chen, and Chris A. Guida*
- Marvel: A Distributed Real-Time Monitoring and Analysis Application / *U. M. Schwuttke, A. G. Quan, R. Angelino, C. L. Childs, J. R. Veregge, R. Y. Yeung, and M. B. Rivera*
- Pharos—The Single European Market Adviser / *Ebby Adhami, Michael Thornley, and Malcolm McKenzie*
- The Credit Assistant: The Second Leg in the Knowledge Highway for American Express / *James Dzierzanowski and Susan Lawson*
- MOCCA: A Set of Instruments to Support Mortgage Credit Granting / *Steve Hottiger and Dieter Wenge*
- Cresus: An Integrated Expert System for Cash Management / *Pete Shell, Gonzalo Quiroga, Juan A. Hernandez-Rubio, Eduardo Encinas, Jose Garcia, and Javier Berbiela*
- DMCM: A Knowledge-Based Cost-Estimation Tool / *Norman Crowfoot, Scott Hatfield, and Mike Swank,*
- SlurryMinder: A Rational Oil Well Completion Design Module / *E. Brent Kelly, Philippe Caillot, Robert Roemer, and Thierry Simien*
- An Application of Model-Based Reasoning in Experiment Design / *Andrew B. Parker and W. Scott Spangle*
- A Truly Magic Solution / *Rita C. Kidd and Robert J. Carlson*
- Adjudipro / *J. P. Little and Mark Gingrich*
- HUB SIAASHING:A Knowledge-Based System for Severe, Temporary Airline Schedule Reduction / *Trish Dutton*
- Arachne: Weaving the Telephone Network at NYNEX / *Elissa Gilbert, Rangnath Salgame, Afshin Goodarzi, Yuling Lin, Sanjeev Sardana, and Jim Euchner*
- Knowledge-Based Code Inspection with ICICLE / *L. Brothers, V. Sembugamoorthy, and A. Irgon*
- Automatic Programming for Sequence Control / *Hiroyuki Mizutani, Yasuko Nakayama, Satoshi Ito, Yasuo Namioka, and Takayuki Matsudaira*

Published by The AAAI Press / The MIT Press

ISBN 0-262-69155-8. $30.00
To order, call toll free: 800/356-0343 or 617/625-8724

MasterCard and Visa accepted