KNOWLEDGE-BASED SIMULATION OF A GLASS ANNEALING PROCESS: AN AI APPLICATION IN THE GLASS INDUSTRY

Richard A. Herrod, Ph.D.
Jeff Rickel
Texas Instruments Incorporated
P.O. Box 225474 MS 439
Dallas, Texas 75265

ABSTRACT

This paper describes a knowledge-based simulation system for a glass annealing process. The long ovens, known as lehrs, in which annealing takes place are not well understood by their operators. In fact, only a few experts can predict the effects of a change in the lehr controls. Attempts to simulate the behavior of the lehr using conventional methods have not been successful due to the size and complexity of the lehr. Our knowledge-based approach is capable of both simulating the glass temperature curve in an annealing lehr and planning the necessary lehr control settings to achieve a desired curve. It consists of two cooperating expert systems, one rule-based and the other frame-based. The system also includes a high-bandwidth graphics display which allows operators to interactively test control-setting changes and ask for the control settings which meet desired specifications. A description of the domain, a history of the development, and details of the design are all presented, along with lessons learned from the experience.

I INTRODUCTION

The Texas Instruments (TI) Industrial Systems Division, located in Johnson City, Tennessee, is involved in the development of advanced industrial control products. Part of this thrust involves investigating the applicability of AI to a variety of industrial areas. One approach being used is to enter into development projects with key industries which are involved in bringing high technology solutions to bear on their businesses. This mutual interest resulted in an agreement between TI and Corning Glass Works to develop a knowledge-based system to address an important aspect of Corning's production operation. Corning provided the process (domain) expertise and TI embedded this knowledge in an expert system.

II THE DOMAIN

Corning sees AI technology as a major factor in process control in the future. They want to start training people in the application of this technology to manufacturing problems. In choosing a problem, they looked for one that was representative of their many manufacturing operations, and one that would provide real economic benefits. They also felt that finding a cooperative and enthusiastic expert would raise the chances of the project's success.

The application chosen was the capture of expertise in the process step of annealing glass to remove residual stresses that originate during the forming operation. In particular, the project would focus on the annealing of television picture tubes at Corning's State College, Pennsylvania, plant. Since upsets at the State College plant were costing the company money, they had a strong interest in finding a solution. This process is also used in Corning plants in Mexico and South Korea, and by all of the Consumer and Lighting Products plants; therefore, there is a potential for wide use of the system. Improvements to the process typically reduce the product losses of the plant and are immediately translated into increased operating margins. Finally, Corning's experts in designing and trouble-shooting the lehrs in which annealing takes place are nearing retirement age, and capturing their knowledge is very important. Thus, this problem area met all of Corning's criteria for problem selection.

The annealing process takes place in a very long oven known as a *lehr*. The lehr provides a controlled heat treatment cycle that softens the glass sufficiently to remove stresses built into the glass during forming. These residual stresses are the result of rapid cooling of the outer part of the glass when it comes in contact with the forming mold. The annealing temperature cycle must be cool enough that the glass will not lose its strength and change shape. Once the stresses have been removed, the glass must be cooled slowly to prevent regenerating them. The lehr which accomplishes all this at the State College plant is approximately 180 feet long.

The temperature profile in the lehr is produced by hot gas generated by burners just inside the front end of the lehr and guided down the lehr through ducts below a steel mesh belt that carries the glass. *Dampers* located along the duct control the amount of hot gas introduced into the sections of the lehr. The hot gas is introduced into the lehr chamber by a series of openings designed to prevent the gasses from directly impinging on the glass. The hot gas is then recirculated through the burner system at the front of the lehr. The temperature of the inlet gas stream is directly controlled by two *thermocouples* in the lehr, one near the front and one about 60 feet inside.

Generally, product defects caused by other processing steps become noticeable when the glass leaves the lehr. At that point the glass has cooled, and any defects which give rise to stress concentrations will cause breakage. When this happens, an expert is usually called in to determine the cause of the problem. If he determines that the lehr is at fault, he makes adjustments to the firing and airflow systems until the annealing process is back "in tune". The expert is also frequently needed when a change in product characteristics necessitates adjustments to the lehr, or when the lehr is restarted after a period of inactivity.

To achieve a desired temperature curve through the lehr, the expert has several controls he can modify. The thermocouples can be set just like a household thermostat; if the temperature at the probe drops below the set temperature, the temperature of the gas produced by the burners is increased until the temperature reaches the desired setting again. Dampers throughout the lehr control where the hot air is directed. By opening all the dampers, the heat is distributed fairly evenly through the lehr. By closing a certain damper, heat is blocked from all sections past that damper. Dampers can also be partially closed. *Ports* and *louvers* throughout the lehr act as air vents, letting hot air escape and thus lowering the temperature in the lehr. The difficult part is using all these controls to raise the temperature in some parts of the lehr while lowering it in others.

Usually, an expert can successfully re-tune the process within two or three attempts. Corning wanted the project to focus on transferring the burden of assisting operators in adjusting lehr firing and airflow systems from its human experts to a computer system. A secondary objective of the project was to develop in-house expertise so that AI techniques and procedures could be applied to other Corning processes.

III A HISTORY OF THE DEVELOPMENT

The development of the Lehr Simulation System took about nine months. Although the expert was skeptical that his methods could be reproduced by a computer, he realized the usefulness of such a system and was therefore very cooperative. Meetings between the TI knowledge engineers, who had no prior knowledge of glass annealing, and the lehr expert were held at Corning approximately every four to six weeks. These discussions lasted about three days each.

The initial goal of the project was to develop a breakage diagnostic which would determine the cause of glass breakage in the lehr; however, after the first interview with the expert, it was obvious that his real expertise lay in his ability to predict how various adjustments to the control settings of the lehr would affect temperatures at various points within the lehr. Because of the immense size of the lehr and the complex interactions between the various controls, operators at the plant cannot predict the effect of control setting changes, nor do they know what controls to modify in order to produce a desired temperature curve through the lehr. Corning had tried to simulate the behavior of the lehr using principles of thermodynamics and heat transfer, but the results could not be correlated with the actual lehr control settings. Thus, changes to the lehr required the expert, who had acquired rules of thumb and knowledge of cause and effect in the lehrs through his many years of working with them.

Although the breakage diagnostic expert system was to be developed first, it was determined that the real payoff would be in a knowledge-based simulator which could determine the temperature curve through the lehr given the control settings. This simulator would be useful not only for ordinary lehr operators, but also for the expert himself. Each time the expert proposes a change in the lehr control settings, about twelve hours elapse before he can see the result of this action as annealing curve changes. Since it typically takes him two to three tries before he achieves the desired effect, the immediate feedback which the simulator could provide would result in significant time savings.

The development philosophy was to get a prototype system up and running as quickly as possible for early evaluation. TI returned to Corning after one month to review the breakage diagnostic and the initial design of the simulator, and again one month later with the first working simulator program. Although this simulator prototype contained a number of deficiencies, it served as a catalyst for uncovering a vast wealth of lehr knowledge previously thought irrelevant by the expert. The entire system was reviewed at that time, with many misconceptions uncovered and a great deal of new knowledge acquired. Having something tangible to critique was enormously beneficial.

The system was iteratively enlarged and refined over the next few months. Each review resulted in refinement of the knowledge and the user interface. Verification of the simulator's accuracy was carried out by comparing its predicted temperature profiles with profiles measured by a thermocouple that was sent through the lehr. In general, the simulator's results were within the repeatability of the measurements. After about four months, the system was demonstrated to Corning process engineers, and, soon after, to Corning executives. It was favorably received with the consensus being that it would be useful to have at the plants.

About halfway through the project, the development of the diagnostic expert system was frozen and the range of the simulation expert system was expanded to include a planning component. While the simulator allows a user to estimate the effects of a control-setting change on the annealing curve, the planner allows a user to input characteristics of a desired curve and receive the necessary control settings.

After one month, the planner prototype was demonstrated to the lehr expert. Several further refinements later, the initial prototype was determined to be inadequate and its knowledge was encoded in a new version which more accurately models the expert's thought process. This new version went through several refinements before finally being packaged up with the simulator to form the Lehr Simulation System, which has been installed in one of Corning's plants for evaluation.

With this new system, an operator can predict the effects on the temperature profile when the control settings are changed. He can use the planner to determine which control settings will provide the desired annealing parameters. The operator can then use the simulator to modify the temperature profile for other processing concerns. For example, if it is not possible to meet the desired annealing parameters, the operator can use the simulator to decide what trade-offs can be made to ensure an adequate temperature profile.

IV THE SIMULATOR

Since the nature of the simulator project did not neatly fit into any of the usual expert systems paradigms, TI and Corning agreed to develop the system from scratch. Corning operates a number of VAX computers, so it was decided to develop the initial prototype of the Simulator on a VAX 11/780 using a public domain version of Lisp (NIL) available from MIT. The system must provide graphic output to the lehr operators, so Tektronix 4107 terminals were chosen as the output device. To encourage modularity, the program relies heavily on Flavors, an object-oriented programming language embedded in NIL. Prototyping the systems would have been easier and quicker using a Lisp Machine like the TI Explorer, but the availability of the VAX for both groups dictated the choice.

In deciding what control-setting changes to make, the expert does a mental simulation of the lehr. He can remember certain temperature curves and their associated control settings, so he uses this as a starting point. He also has knowledge of the general effect a particular control-setting change will have and which sections of the lehr will be affected. This approach directly led us to our simulation strategy and knowledge representation.

To encode the knowledge, we created structures much like frames that can represent both the magnitude and range of effect for each relevant lehr control. There is one such cause-and-effect structure for each damper, port, and louver. Each structure has slots which associate the various valid settings for that control with their corresponding effect (magnitude and range) on the slope of the temperature curve.

Other knowledge is embedded in the calculation routines. For instance, if a damper is only X% open, then the dampers after it cannot behave as if they were more than X% open. Thermocouples are accounted for by the knowledge that they represent points through which the temperature curve must pass, since any change in the temperature at these points causes the burners to compensate in order to bring the temperature back in line. Finally, the speed of the conveyor belt is used to determine how closely the glass temperature follows that of the air temperature.

To determine the air temperature curve through the lehr, we start with a default curve and default settings, just as the expert does. We then multiplicatively combine the effects of all control-setting deviations from the defaults with regard to the sections of the lehr they affect. The default curve and the resulting curve are expressed in terms of a series of slopes. By now using the temperature at which the air enters the lehr (a given), the temperature at each thermocouple, and knowledge of what the peak temperature will be and where it will occur (based on the control settings), we can propagate these slopes to arrive at the estimated air temperature curve through the lehr.

Based on sessions with the expert, it was determined that the glass temperature curve basically follows the air temperature curve with a slight lagging effect. The slower the belt speed, the less the glass temperature lags. However, to further complicate things, the responsiveness of the glass to changes in the air temperature also depends on the section of the lehr through which the glass is passing. Our knowledge representation scheme was expanded to allow the lehr to be sectioned off and the responsiveness of each of these sections recorded.

To maintain modularity in the system, the lehr is represented as a flavor object with a large number of instance variables. These instance variables record physical characteristics of the lehr as well as the lehr's associated cause-and-effect structures. It turns out that the particular cause-and-effect knowledge structures remain constant within various classes of lehrs. Thus, by simply defining the physical features of a new lehr, the system is capable of simulating that lehr, even though its physical structure may be different than previous lehrs. This fact was verified when we obtained temperature curve data for a lehr for which we originally had no data. The curves estimated by the simulator matched quite closely with those actually measured at the plant.

V THE PLANNER

While the simulator is able to estimate the temperature curve based on the control settings, the planner takes as input the desired curve parameters and produces the control settings which most closely achieve these parameters. Again, the expert's approach to the problem provided a model for the system. The expert has a sort of bag of tricks he uses to achieve a certain response in the lehr. For instance, if he wants to lower the cooling rate in some section, he may raise the back thermocouple. Surprisingly, there are also cases when *lowering* the back thermocouple temperature may lower the cool rate. He therefore uses his bag of tricks by pulling out of it the particular trick which applies to the current situation. By iteratively using these tricks to get closer to his goal, he eventually determines the necessary control settings.

The planner emulates this style of reasoning with forward-chaining rules. Given a goal of, say, lowering the cool rate, it begins trying rules which are known to accomplish this goal. It tries the rule by asking the simulator what the effect of the proposed change will be on the current temperature curve. If the change is in the right direction, it is made, resulting in a new temperature curve. If the change is in the wrong direction, the idea is abandoned and the next applicable rule is tried. Thus, through cooperation, the simulator program and the planner program together arrive at the desired control settings.

Forward chaining was chosen (as opposed to backward chaining) because the rules are not used for logical deduction. Instead, they guide the system through the search space of possible control-setting combinations. Each rule consists of a context in which this rule is applicable (i.e. raising the hold time), a set of actions to perform (i.e. changing a control setting or creating a subgoal), and conditions for success. The actions transform the current state of the control settings to some new state. The new state is then evaluated with the success conditions. If these conditions are met, the new state becomes the current state; otherwise, we proceed to the next applicable rule. Conflict resolution is handled through rule ordering; this allows the system to try the changes which are most likely to succeed first.

Because most rules are very similar, a small, extensible rule language is provided to allow non-programmers to modify the knowledge-base. This rule language does not, however, limit the flexibility of the rules; although all current rules are expressed in the language, new rules can be arbitrary Lisp functions. The rule language supports all the power that was found necessary in the current rule set, including forward chaining, the creation of subgoals, the use of Lisp predicates, constraint-posting, the setting of multiple controls, and much more. The design of this rule language was driven by the fact that the knowledge base will be maintained by people who are unfamiliar with Lisp.

Besides the separation of the planner's knowledge base and inference engine, the planner includes one more level of modularity: the *strategy* section. In talking with the expert, we found that achieving some temperature curve parameters is more important than meeting others. Also, some requirements can be relaxed in order to achieve others. The strategy section houses this knowledge.

The strategy section provides the goals which the inference engine and knowledge base will try to meet. It can also post constraints in order to protect previously achieved goals. These constraints act as implicit conditions for success for all rules. For instance, the strategy section might direct the knowledge base to meet some cool rate specification without raising the temperature at which the glass exits the lehr by more than 5%. Based on the results of trying to meet this goal, it can then modify its strategy and send the next appropriate goal. Since the specification of each goal (and thus the whole strategy) is fairly simple, we hope that this strategy section will prove to be just as modifiable as the knowledge base. The strategy section is a simple yet effective way of handling multiple interacting goals.

VI THE USER INTERFACE

Since the system will be used by people who are not familiar with computers, care was taken in developing an operator interface that would be easy to use. In the simulator, particularly, much time was spent on giving the user as much relevant visual information as possible. Using a color Tektronix terminal, an interface was developed which shows a graph of both the air temperature curve and the glass temperature curve in the lehr (see Figure 1). By using the space bar and the backspace key, the user can position the cursor over a particular control setting, change its value, and watch the resulting change in the two curves. As shown in Figure 1, the user can change the temperature settings of the front and back thermocouples (denoted **tc3** and **tc7**), open or close dampers (an integer represents the state of a damper, with **1** indicating completely closed and **5** indicating completely open), change the speed of the conveyor belt, or move the position of the first open port or louver (indicated on the graph by **P** and **L** respectively). Also, a curve analysis component was added which can give detailed

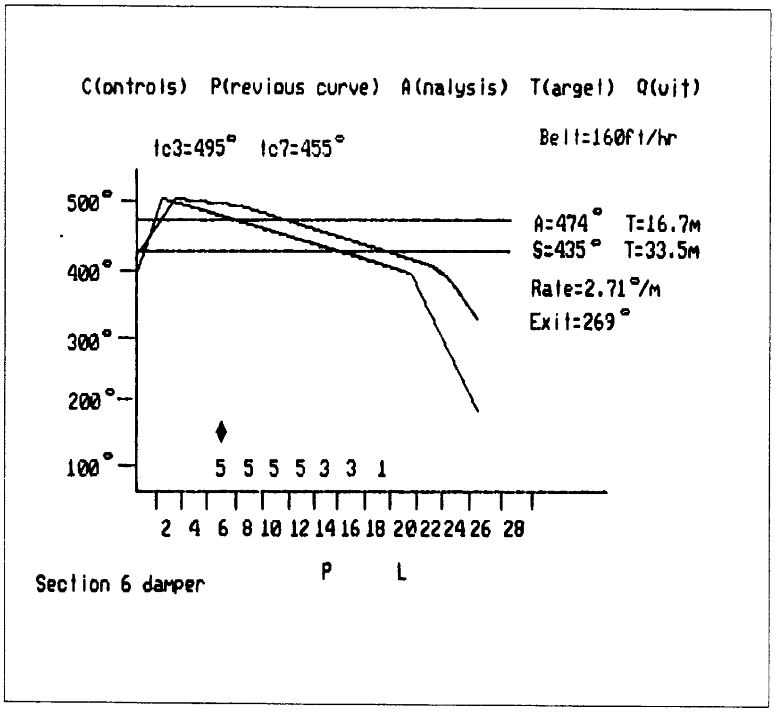

Figure 1: A typical screen from the heuristic simulator showing the air temperature and glass temperature curves along with the various control settings.

information and statistics about the glass curve.

The command line at the top of Figure 1 shows the user his available options. Pressing a **T** allows the user to specify the parameters of a "target" curve; this is how the planner is invoked. **A** takes the user to an analysis screen, which shows a detailed analysis of the glass temperature curve. Typing **C** takes the user to a screen which displays both the current control settings and the previous control settings. This is customarily used to determine what changes the planner made. The **P** option is closely related to this; it displays the previous glass curve (as a dotted line) on top of the current temperature curves so the operator can graphically see the changes. Finally, **Q** terminates the session.

VII SYSTEM PERFORMANCE AND EVALUATION

System response is fairly prompt. The simulator takes about one second after being invoked to return the estimated temperature curves. The planner is more variable, since it must invoke the simulator for each proposed change. On most problems, it takes less than one minute to run. In extreme cases, where the desired curve is very different from the current one, it may run for several minutes.

Maintainability and modifiability were high priorities in this project. It is easy to add new lehrs to the system if they are fundamentally similar to existing ones; only the physical properties need be recorded. To add a lehr which is drastically different from existing ones, the knowledge structures of cause and effect will also have to be modified. The planner is also easy to expand. New rules can be added with no real programming skills, and the strategy section can also be easily changed to reflect new trade-off considerations. We think that the maintainability of this system is one of its strong points.

VIII LESSONS LEARNED

The field of expert systems is burgeoning. Reflecting on the experience gained from this project, some important lessons can be listed:

- One of the keys to success is rapid prototyping of the system. It is more important to do the first attempt quickly than it is that the first attempt be complete. Experts are not necessarily aware of their thought process and it takes time for them to explain what they know. Having a prototype system to work with is a great help in uncovering the necessary knowledge. Each review of the system will elicit new information.
- The process of finding out what experts know is a very difficult one for knowledge engineers. They must be willing to ask experts to go over something time and again until they understand it. There are times of great discouragement in this process. Conversely, the experts must be understanding enough to be fully cooperative even if they are skeptical about what is going on.
- Strong management commitment to a project of this type is absolutely essential to its success. An expert's time is in short supply and the project must have a high enough priority to have sufficient access to that expert.
- Early demonstration to potential users is important both in getting feedback on possible deficiencies and in gaining their support. After all, what good is an expert system if no one uses it?
- Expert systems must continue to grow, so attention should be given to the people who will maintain the system. The method for adding new knowledge should match their capabilities. Don't expect process operators to become Lisp programmers.

IX SUMMARY

This project demonstrated that it is possible to build a model of a manufacturing process based on the knowledge of an expert. The resulting system, which relies on the expert's intuitive feeling of heat transfer rates and control variable interactions, is able to predict the measured system response. Consequently, the lehr operator is now able to make better changes in the control variables and to minimize process upsets caused by making the wrong control change. Also, the plant has a tool that allows them to plan changes to the lehr control variables for different products rather than waiting until a problem arises and the solution is more expensive.

ACKNOWLEDGEMENTS

The authors would like to thank the following people for their contributions: Rick Saenz, one of the original knowledge engineers; Jacky Combs, who helped with the user interface; G. Thomas Holmes and Joseph Hurley, who managed the project for Corning; and Bob Spadinger, our cooperative expert. We are also grateful to all the people who reviewed this paper.

REFERENCES

Adams and Williamson, "Annealing of Glass," *Journal of the Franklin Institute*, Nov. and Dec. 1920.

Burke, Glenn S., George J. Carrette, and Christopher R. Eliot, "NIL Reference Manual," MIT LCS Publications, Cambridge, Mass., January 1984.

Lillie, H. R., "Stress Release in Glass, A Phenomenon Involving Viscosity as a Variable with Time," *J. Am. Ceram. Soc.*, 19(1936)45.

Moon, David, Richard M. Stallman, and Daniel Weinreb, "Objects, Message Passing, and Flavors," *Lisp Machine Manual*, Sixth Edition, MIT LCS Publications, Cambridge, Mass., June 1984, Chapter 21.

Rajagopalan, Raman, "Qualitative Modeling in the Turbojet Engine Domain," *Proc. AAAI-84*, Austin, Tx, 1984, pp. 283-286.