

Applying Case-Based Reasoning to Manufacturing

David Hinkle and Christopher Toomey

■ CLAVIER is a case-based reasoning (CBR) system that assists in determining efficient loads of composite material parts to be cured in an autoclave. CLAVIER's central purpose is to find the most appropriate groupings and configurations of parts (or loads) to maximize autoclave throughput yet ensure that parts are properly cured. CLAVIER uses CBR to match a list of parts that need to be cured against a library of previously successful loads and suggest the most appropriate next load. CLAVIER also uses a heuristic scheduler to generate a sequence of loads that best meets production goals and satisfies operational constraints. The system is being used daily on the shop floor and has virtually eliminated the production of low-quality parts that must be scrapped, saving thousands of dollars each month. As one of the first fielded CBR systems, CLAVIER demonstrates that CBR is a practical technology that can be used successfully in domains where more traditional approaches are difficult to apply.

Lockheed manufactures many parts for aerospace applications from multiple layers of graphite-threaded composite materials. The use of composite materials, especially in aerospace applications, is on the increase because of their unique weight and strength qualities. Depending on the orientation of the graphite fibers, a part can be extremely flexible in one direction but rigid in another. In addition, a part made from composite material is both lighter and stronger than aluminum. The increased use of graphite parts, as well as the high cost of a spoiled part (as much as \$50,000 for a single part), has put greater reliability and efficiency demands on a relatively new and complex manufacturing process. CLAVIER is a fielded advisory system that Lockheed shop floor personnel use to improve the efficiency of the composite-fabrication shop and simultane-

ously ensure that high-quality parts are produced. CLAVIER's central component uses case-based reasoning (CBR) (Redmond, 1990; Rissland, Kolodner, and Waltz, 1989; Kolodner, Simpson, and Sycara 1985) to recommend collections of parts and appropriate spatial configurations for curing in a large pressurized convection oven known as an *autoclave*.

The following section describes the composite-fabrication domain. The section entitled The CLAVIER System discusses the CLAVIER system, its CBR component, and the rationale behind the selection of the CBR problem-solving method. The next two sections discuss the development, deployment, use, and payoff of CLAVIER. Finally, the last section presents some of the important lessons learned in developing and fielding CLAVIER that extend to other AI and non-AI application-development efforts.

Application Domain

Composite part fabrication requires two major steps: lay-up and curing. *Lay-up* is the painstaking process in which multiple layers of graphite and fiberglass composite material are fitted by hand on the exterior of a contoured mold. The lay-up of a single mold takes from two to seven days, depending on the size of the mold and the skill of the technician. In the second step, *curing*, the molded composite material is hardened by pressurized heating in a large convection autoclave.

The length of the curing cycle (six to eight hours), the limited number of available autoclaves (two in Lockheed's Sunnyvale, California, facility), and the high part-production rate require the shop to cure multiple parts in each autoclave load. However, for the parts to




Figure 1. Thermodynamic Profile of a Load While Curing in an Autoclave.

be cured effectively, all the parts in a load must heat up at approximately the same rate.

In particular, during the ramp-to-dwell stage (figure 1), the hottest part, the *leader*, and the coolest part, the *lagger*, must be within a 30° F delta, and all the parts must advance at least 1 degree per minute. Once the parts enter the first dwell phase, the operator has 40 to 80 minutes to get all the molds within a 20° delta. Then, pressure is added to the autoclave, and the parts must maintain the 20° delta during the second ramp phase. Once all the parts have reached the cure stage, they must be cured at this temperature for two to three hours. If any of the molds don't follow the correct thermodynamic profile, a discrepancy report must be issued, and the part must be inspected for flaws. If the part has been damaged or weakened, it must be scrapped.

Thus, optimal autoclave loads are those that maximize the number of parts that are cured while all the molds are kept within the thermodynamic engineering specifications. The chief technical problem faced by a composite-fabrication shop—and the primary problem addressed by the CLAVIER system—is determining a set of autoclave loads that will correctly produce a given list of parts.

Designing loads for the autoclave is a complex task that has few guiding principles and requires experienced personnel. Two major

factors must be considered when designing an autoclave load: (1) the particular molds chosen and (2) the spatial arrangement of the molds within the autoclave. Each mold has its own inherent heating characteristics, which are affected by factors such as the size of the mold, the shape of the mold, and the thickness of the material (figure 2).

The position of the mold within the autoclave is critical to the effective curing of the part. The shifting of a mold's position as little as 12 to 24 inches can cause it to fall outside the target thermodynamic profile. Within the autoclave, the heat is not uniform; spots within the autoclave are naturally warmer or cooler than others. For example, the front of the autoclave is generally warmer than the back. Furthermore, because an autoclave is a convection oven, the placement of molds in the front of the autoclave influences the air currents reaching the molds in the back, creating relatively warm and cool spots that cause molds to heat up either more quickly or more slowly. These heating characteristics and temperature variations must all be taken into account when determining the grouping and configuration of a set of molds for a load.

For example, figure 2 shows a load in the autoclave. This particular load has four molds: S-455, D-144, D-145, and D-337. The S-455 mold is large and heats up slowly. The other molds are smaller and heat up more

quickly. Although these molds do not seem compatible, they are. Because the large, slow mold is in the warmer front of the autoclave, it heats up more quickly. In addition, the small, faster molds are behind the large mold, which partially blocks the airflow to the back of the autoclave, causing them to heat up more slowly. These factors compensate for each other, making the load compatible.

The CLAVIER System

CLAVIER is a case-based shop floor assistant that addresses the problem of properly grouping and spatially configuring sets of composite parts (loads) for loading into an autoclave. It is a stand-alone application that is written in Macintosh Common Lisp, runs on a standard MACINTOSH with eight megabytes of random-access memory, and has an extensive high-level graphic user interface to make CLAVIER's capabilities accessible to the shop floor personnel. The functional architecture of CLAVIER is shown in figure 3.

The primary objective of CLAVIER is to provide shop floor personnel with an intelligent load-selection aid that helps to ensure high-quality composite part curing yet maximize the quantity and priority of the parts processed through the autoclave. Thus, its central component is a CBR system consisting of a case base of previously used loads, a load retriever that suggests loads from the case base, and a graphic load editor and a new-load validator for use in maintaining the case base (figure 4). CLAVIER also has facilities for capturing and tracking pertinent shop floor data, such as the part production schedule that drives the shop, the number and work shifts of shop personnel, and the supply of material and other resources. Last, CLAVIER has a heuristic load planner-scheduler that uses the case base and the shop floor data to plan several days' worth of autoclave runs at once. Figure 5 shows the top-level graphic user interface that you use to interact with CLAVIER's planner-scheduler and that also provides access to the rest of the CLAVIER system.

The Case-Based Loading Adviser

CLAVIER's central component is a case-based loading adviser that assists the user in arranging composite parts inside an autoclave to achieve maximum throughput yet maintain part quality and minimize the effort needed to control heat-up rates. As shown in figure 6, a load-selection consultation with CLAVIER involves as many as three

steps: (1) case retrieval, (2) case adaptation, and (3) case validation.

Knowledge Representation One of the major advantages of CBR is that it is possible to build and field a system with a small library of seed cases and allow the knowledge base to be expanded and refined over time. Initial cases (past loads) were taken directly from the experts' notebooks that they were required to maintain as part of their regular job. Cases were annotated with text comments and classified valid or invalid for each autoclave (validity is context dependent; a load that is valid in one autoclave might not be valid in another, even if the autoclaves are the same size and have similar vent-airflow configurations).

A graphic editor was developed to enable users to edit and record their own cases. User ability to manage a nonmonotonic knowl-




Figure 2. Arrows Indicate the Airflow's Path through the Autoclave.




Figure 3. The CLAVIER High-Level Functional Architecture.

edge base was critical. Cases consist of the molds to be cured in the load, any tables that are needed to support the molds in the autoclave, and the spatial arrangement (two-dimensional coordinates) of the molds and tables.

The expertise of a CBR system is accumulated in a library of cases. A case represents both a problem's context (used to determine if a case is similar to a new problem) as well as the correct solution to the problem. In CLAVIER, the context explicitly represented in a case includes the tables used in a load and their positions, the molds on the tables and their positions, and information on the results of running the load in the autoclave (that is, valid or invalid). Implicitly represented in each case, through the association of the context with the solution, is the complex reasoning required to consider all the factors that affect the quality of the parts in a load. It is important to note that this information does not need to be stated explicitly (which would be difficult, if not impossible, in this domain). Currently, the case library is maintained by the experts themselves as a by-

product of their interaction with CLAVIER, again as part of their normal job.

Case-Retrieval Mechanism CLAVIER embeds CBR technology within a complete data management system for the manufacturing shop floor.

The retrieval mechanism has two input: (1) the case memory of previously run autoclave loads and (2) the list of parts that need to be manufactured (figure 6). CLAVIER recommends loads using three main criteria: (1) to maximize the number of needed parts that the load will manufacture, (2) to minimize the number of unmatched (extraneous) parts that the load contains, and (3) to maximize the quality of the load (determined by part compatibility). The retriever recommends several loads to the user in ranked order. If an exact matching case is found, this load is selected.

Case Adaptation and Validation If an exact matching case cannot be found, CLAVIER presents the closest matching cases. The user then decides how he/she wants to try to modify the case. After the user makes a modification, CLAVIER tries to validate the new con-




Figure 4. CLAVIER Load-Validation Interface.

figuration. Validation is done by comparison with similar valid and invalid cases. CLAVIER makes a recommendation about whether or not the new case is likely to be valid. If the load is predicted to be valid, the system then proceeds to generate the description of how to configure the molds within the autoclave. If the system predicts that the load might be incompatible (because of similarity to an invalid load), it suggests alternative configurations that are similar but valid. If the system strongly believes that the load will be incompatible, it suggests ways of breaking the single load into multiple valid loads. In this case, CLAVIER sacrifices some of the load's efficiency to decrease the risk of part defects. After the load has been cured in the autoclave, the operator tells CLAVIER whether or not the load was successful. The annotated new case is then stored in the library, allow-

ing the system to expand its expertise and, thus, learn.

Why Case-Based Reasoning?

The autoclave-loading domain is a particularly difficult one in which to apply traditional knowledge engineering techniques. In talking with the expert autoclave operators, it became clear to us that sometimes even they are forced to use trial-and-error methods. When they encounter a new situation (for example, a mold type they have never cured before), they cannot to predict what molds it will be compatible with without testing several possibilities in the autoclave. Once they gain some experience with a mold, they are able to reason about what other molds might be compatible with it, but they must still validate any hypothesis in the autoclave. Even the best experts are not sure if a load will be

Clavier Autoclave Loading Advisor

Current Plan:

Mon 11 Jan (127)		Tue 12 Jan (128)		Wed 13 Jan (129)		Thu 14 Jan (150)		Fri 15 Jan (151)	
First Run	(Rc1) 2am								
(Pending)	(Plannable)								
S-136		D-316	D-693	D-144		S-152		S-136	
S-136		D-141	D-693	D-144		S-152		S-136	
		D-141	S-814	D-145				S-136	
		D-606		D-145					
		D-606		D-606					
Second Run	(Rc2) 2am								
(Pending)	(Plannable)								
S-136		S-616		S-136		S-706-101		D-136	
S-136		D-316		S-023		S-660-103		D-136	
D-136		S-616				S-660-101		D-137	
D-137						S-706-102		S-136	
Third Run	(Rc1) 9am								
(Pending)	(Plannable)								
D-226	D-903	S-406		D-226	D-903	S-709-101		D-226	D-903
D-226	D-209			D-226		S-660-106		D-226	D-209
D-337				D-337		S-660-105		D-337	
D-337				D-337		S-709-102		D-337	
D-606				S-616				D-606	

Options:

Plan Creation/Extension: Add Next Day, Update/View Planning Data, GetPlan

Plan Tracking & Editing: Update Execution Status, Toggle Plannable->Fixed, Interchange Loads, Select/View Load Layout, View/Assess Completed Load

Molds: Add Mold(s), Delete Mold(s), Move Mold(s)

Other: Print Plan, Print Load's Layout, View Past Loads

Messages: Exit Clavier

Figure 5. CLAVIER Planner-Scheduler Interface.

compatible until they test it in the autoclave.

A constructive, rule-based approach to load generation was found to be infeasible because even the experts did not have the first principles needed for such an approach. When they were asked to explain how they determined the correct position of a mold within a load, they were unable to do so except within the context of a specific load that they had previously cured in the autoclave. With few exceptions, the experts' reasoning concentrates on the load as a whole rather than on the placement of individual molds.¹

Another approach that was considered was the use of thermodynamic modeling. With this approach, a mathematical model is constructed to simulate the thermodynamic properties of a mold. This approach has been

tried, with some success, in production facilities that are curing single parts at a time and that, typically, are manufacturing each part only once or twice. This approach, however, is not feasible in a continuous, high-volume production environment in which multiple parts must be cured for each load. When there are multiple parts to a load, it is not only the thermodynamic properties of the mold and the thermodynamic properties of the airflow that must be modeled but also the effect that a particular mold in a particular position has on the airflow reaching the molds behind it. This tremendous increase in complexity makes thermodynamic modeling prohibitively difficult and expensive when dealing with a manufacturing process such as Lockheed's, where mold interaction is a criti-




Figure 6. CLAVIER System Task Flow.

cal factor. In summary, there is no reliable way to accurately predict mold compatibility before testing the load in an autoclave.

Because of the difficulty of applying traditional expert system techniques to the domain, we decided to try machine-learning techniques. Clustering and induction techniques were explored, but they are not well suited to this domain. First, because of the spatial aspects of the domain, the total number of possible loads (that is, the *search space*) was extremely large, and our sample of past loads (that is, the *sample space*) was small, only a few dozen cases.


Second, although our sample of past loads was classified into two categories, valid and invalid, the actual situation is a little more complicated. Loads in which a mold goes outside the allowable thermodynamic profile are clearly incompatible (invalid). However, even loads that do stay within the thermodynamic profile can still be classified invalid based on how close they came to going outside the profile (that is, how risky they were). In addition, how much risk the operators are willing to take depends on whether some alternatives are less risky. For example, if a new load is developed that is similar to, but less risky than, a load currently classified as valid, the new load is added to the database,

and the old load is reclassified as invalid. Thus, in practice, the validity of a load is not strictly Boolean and can change over time.

Before CLAVIER, when the expert autoclave operators were trying to decide which parts to load in the autoclave and how to arrange them, they would look through their log books in search of a past load that was applicable to the current situation. After observing the operators, it quickly became clear to us that the human experts were reasoning from whole past experiences and that CBR was the most appropriate technology to apply to the problem.

Development and Deployment

Development of the CLAVIER system began in March 1989. The initial version was fielded in November 1989. Development of version 1.0 was completed in September 1990. From September 1990 through November 1991, the system was substantially expanded, including the load-validation and the planning modules as well as extensive data-entry, record-keeping, and report-generation capabilities. With version 2.0, the scope of the system was expanded to include virtually all aspects of the composite-fabrication process. Development time has been estimated at two person-years. In December 1992, CLAVIER was extend-



ed for use in Lockheed's composite-manufacturing facilities in Georgia. Maintenance and enhancement of the CLAVIER system continues.

Throughout the development of CLAVIER, we took a unified approach to system design, user feedback, and training. Joe Sferrazzo and Henry Rodriguez (both expert autoclave operators) were part of the development team from the beginning. Consequently, after a few days of coaching in the use of the fielded system, the operators were off and running, with only occasional queries. After we trained the initial users, they, in turn, trained all the other operators.

Use and Payoff

CLAVIER has been in continuous daily use at Lockheed's Composites Fabrication Facility in Sunnyvale, California, since September 1990. Two to three autoclave loads are cured each day in this facility, all of which are selected through operator consultations with CLAVIER. CLAVIER also generates hard-copy reports of the autoclave loads that are used for record-keeping purposes. The system has recently been expanded for use in other Lockheed manufacturing facilities, and negotiations are under way to license the software to other aerospace companies. The CLAVIER system is useful for any autoclave area with high-vol-

ume production and multiple parts to each autoclave load.

CLAVIER ensures that high-quality load configurations are used for manufacturing composite parts, even when the experienced autoclave operators are unavailable. This consistent level of expertise is critical to producing high-quality parts and maintaining the production schedule. There are now five operators and two support persons who regularly use the system as part of their daily routine to generate autoclave load configurations and other reports.

If a mold goes outside the correct thermodynamic profile, a discrepancy report is issued, and the part must manually be inspected at a cost of \$1000. If the part is flawed and must be scrapped, it costs an average of \$2000, but for some parts, it can cost between \$20,000 and \$50,000! Since CLAVIER came on line, discrepancy reports as a result of incompatible loads have virtually been eliminated, saving thousands of dollars each month.

One important additional benefit to CLAVIER is that it has clearly demonstrated—both to management and the technicians on the shop floor—the power of knowledge-based systems. Since CLAVIER's initial fielding, we have developed several other knowledge-based applications for use in other stages of the manufacturing process.

Lessons Learned

Several important lessons about building and deploying real-world AI applications were learned over the course of the CLAVIER project. Most of the lessons seem obvious in hindsight but are nevertheless easy to neglect and are important to keep reminding oneself of.

First, users do not care whether the application uses sophisticated AI techniques or random guesses to generate results; what they care most about is that the system is easy to use and provides tangible benefits. Thus, the user interface and other mundane components that simplify use or save labor are at least as important to the success of the application as the underlying algorithms; so, development effort should be allocated accordingly. CLAVIER's single most popular feature, for example, has probably been its ability to produce hard copies of the autoclave loads, which users must have for record keeping and which they formerly drew by hand.

Another lesson we learned is that because algorithm output can only be as good as the algorithm's input (garbage in equals garbage out), it is important to assess the quality of the data input before investing significant resources in developing sophisticated algorithms. In CLAVIER, for example, there is a high degree of uncertainty in many of the shop floor data input used by CLAVIER's multiple-load planner-scheduler. This uncertainty limits the extent to which load schedules can accurately be projected to approximately one week, although we designed CLAVIER's planner-scheduler to accommodate planning for several weeks into the future.

Conclusion

CLAVIER has shown that CBR can be an effective problem-solving method in complex, real-world domains, including those not amenable to other AI and non-AI techniques. CLAVIER also illustrates, however, that regardless of the sophistication and elegance of the underlying problem-solving technique, it is often the application's user interface and labor-saving features, as well as the quality of its data input, that determine the application's success as a fielded system.

Acknowledgments

We want to recognize Ralph Barletta, Linda Cook, Dan Hennessy, and especially Bill Mark for their contributions to CLAVIER's development. We are indebted to the members of the composite-fabrication shop—especially Joe Sferrazzo, Henry Rodriguez, Diana LaPietra,

and Byron Ravenscraft—for their cooperation. We would also like to thank John Creagan and Ram Sriram for their continued support.

Note

1. One exception that was found, which CLAVIER uses in validating layouts, was that a valid load that is modified strictly by removing molds (that is, it has a subset of the molds) will generally be compatible. The remaining molds, however, will typically have to be repositioned.

References

- Hennessy, D., and Hinkle, D. 1992. Applying Case-Based Reasoning to Autoclave Loading. *IEEE Expert* 7(5): 21-26.
- Kolodner, J.; Simpson, R.; and Sycara, K. 1985. A Process Model of Case-Based Reasoning in Problem Solving. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, 284-290. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Pursley, M., and Shelton, F. 1985. Exploratory Development on the Processing Science of Thick-Section Composites (Contract No. AFWAL-TR-85-4090), Lockheed Aeronautical Systems Company, Marietta, Georgia.
- Redmond, M. 1990. Distributed Cases for Case-Based Reasoning; Facilitating Use of Multiple Cases. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, 304-309. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Rissland, E.; Kolodner, J.; and Waltz, D. 1989. Case-Based Reasoning. In *Proceedings of the DARPA Case-Based Reasoning Workshop*, 1-13. San Francisco, Calif.: Morgan Kaufmann.



David Hinkle is a senior scientist at Lockheed's Artificial Intelligence Center, where he is currently investigating applications of case-based reasoning technology. He received his B.S. in computer science from Syracuse University in 1987 and his M.S. in computer science from Northeastern University in 1989. His research interests include case-based reasoning, machine learning, expert systems, robotics, and neural networks.



Chris Toomey is a senior scientist at Lockheed's Artificial Intelligence Center in Palo Alto, where he is currently investigating distributed information processing using software agents. He received his M.S. in computer science from Stanford University in 1992. Beyond his primary interest in agent-based information management, his research interests include information retrieval and filtering, planning, machine learning, and human-computer interfaces.