Expert Systems in Data Processing: California Travel Expense Claim System

Allyson Bartman-Gatt

The CALTREC project automated the process of creating travel expense claims for the state of California. This project is a good representative of the types of projects expected to be done by the state. The application is heavily embedded in traditional data processing systems, which added much additional work outside the scope of what is considered to be knowledge engineering.

The Travel Expense Claim Process

The state of California employs about 200,000 people who work for over 200 departments and belong to 20 collective bargaining units. Each year, approximately one third of these people travel on state business and incur expenses in the course of their trips. In addition, while conducting state business, many employees incur a wide range of nontravel-related expenses, including business expenses and overtime meal expenses. Everyone who incurs expenses on state business is entitled to be

reimbursed by the state. The vehicle through which the state of California reimburses its employees for state–business-related expenses is a form called the *travel expense claim*, which documents the expenses being claimed. It is completed either by the claimant or a secretary, and until recently, these claims were produced manually. Roughly 400,000 state of California travel expense claims are submitted each year.

Once a travel expense claim is completed by the claimant, it is submitted to the claimant's supervisor for approval, then forwarded for further auditing. Completed travel expense claims are subjected to two levels of audit: The first audit occurs in the department's accounting section, the second at the State Controller's Office (SCO). Once the accounting section approves the claim, the claimant is reimbursed from department funds, and the claim is forwarded to SCO, which audits the claim a second time. Once the claim is approved by SCO, the department is reimbursed.

The rules used to audit travel expense claims are developed by the state of California's Department of Personnel Administration (DPA) and distributed to the rest of the state departments. What is allowed for reimbursement is subject to a number of limitations, depending on such factors as the terms of the employee's current collective bargaining unit contract, the destination and length of travel, and the start and ending dates of travel. In addition, depending primarily on amounts expended for certain items, approvals by various individuals, as well as receipts, can be required.

These rules tend to be confusing and difficult to interpret. Although the rules are documented in a number of manuals, often a person submitting the claim does not know of the existence of such manuals and relies on other sources for the rules governing how to fill out a claim. Even employees who know where to access correct information often do not bother or misinterpret the documentation. To make matters worse, travel expense claim rules change relatively frequently. DPA notifies, by memo, each department's accounting section of changes, but invariably this information does not get distributed to everyone who needs it, and the error rate on subsequent travel expense claims tends to increase. Because people have trouble even remembering or understanding a single version of the travel expense claim rules, to expect them to be able to remember that a change has taken place (if they know about it) and understand how this change should affect the travel expense claim is hopeless.

Given these conditions, it should be no surprise that the error rate of manually produced travel expense claims is staggering; estimates by SCO of the error rates it experiences range from 50 to 80 percent, which suggests that even accounting sections are not correctly applying

travel expense claim rules. Because these rates are so high, SCO must audit every single travel claim it receives and is virtually swamped by a backlog of claims to process. Each time an error is detected, the claim is returned. It goes back to the claimant if the accounting section detected an error or to the accounting section and then to the claimant if SCO detected the error. Once the error is corrected, the audit process resumes, but delays result, either in getting reimbursement to the claimant or releasing funds to the department. Although estimates of the exact cost of this process to the state of California have not been made, it is easy to see that given the large numbers of travel claims involved, the cost of the extra handling, along with the delay in the transfer of funds, must be immense.

Travel expense claims can have errors that result in unnecessary cost to the claimant, such as when the claimant is entitled to be reimbursed for an expense but does not claim it because of a lack of knowledge of the rules. The audits performed by accounting sections and SCO are geared to detect cases where reimbursement is being claimed for expenses that are not allowed under the DPA rules, thereby saving money for the state. The auditors assume that the claimant knows the rules and has claimed all reimbursable expenses that were incurred and, conversely, that expenses that were not claimed were not incurred. However, often, the claimant does not fully understand the rules and assumes that an incurred expense is not reimbursable when it really is and does not claim the expense. Because the audit process has no interest in and really no way of protecting the interests of the claimant, the claimant's ignorance of the rules goes undetected. This situation represents another, probably large set of errors that, although not a direct cost to the state of California, is significant nevertheless.

Clearly, the extremely high error rates associated with the travel expense claim process represent a substantial cost to the state of California and its employees. One way to attack this problem would be to implement the travel expense claim rules in a centrally located and maintained computer application. Maintaining all the rules in one place would ensure that the correct rules would always be applied to every claim, thus travel expense claims produced by such a system would be more accurate, consistent, and complete. Equally important, all claims would reflect the most recent rule changes as soon as the changes were added to the software. In fact, several attempts were made to build such an application using traditional software. However, every attempt failed. These applications were no more than data-entry systems; in every case, the software used was inadequate to represent the number and complexity of rules required. Thus, the error rates associated with the attempted travel claim systems were no better than

with the manual process. Further, none of the persons involved with these attempts consulted with SCO or DPA in the course of building the applications, evoking the wrath of both entities, which refused to accept the travel expense claim forms produced by these systems.

The CALTREC Project

In 1988, Russ Bohart, the director of the state of California Health and Welfare Agency Data Center (HWDC), became interested in expert systems. It seemed to him that this technology would be widely applicable in state government because much of what goes on in the course of state business consists of judgmental processes based on policies, regulations, procedures, and guidelines, in other words, rules. These processes are continually getting more complex and more difficult for a person to remember and understand, inevitably resulting in errors and inefficiencies for which some cost must ultimately be borne. Furthermore, such processes do not lend themselves well to automation using traditional techniques because of the large number, the complexity, and the changeability of the rules involved. Because expert system technology is directed at just this sort of problem, Bohart decided that his organization should embark on an experiment to evaluate the potential usefulness of expert system technology to the departments that HWDC serves and to the state of California.

HWDC provides computing resources to 10 state departments, supporting about 45,000 employees and 12,000 terminals on a network that extends throughout the state. The data center itself employs a staff of 200 people who run the computer systems and provide technical support to users. The predominant mode of large-scale computing in the state of California is based on IBM System 370 architecture, that is, large mainframes running 1 or more of the 3 IBM 370 operating systems that support many large online and batch applications. Because HWDC is one of the largest IBM installations in the state of California, its expert system applications must run on the available IBM platforms. Thus, the expert system experiment would be conducted using an expert system shell that would run on an IBM mainframe system. HWDC negotiated with the (at this time, 2) vendors that were then offering such shells, and IBM's expert system environment (ESE) was chosen.

Nine volunteers—seven from the data processing divisions of five HWDC-supported departments and two from the data center itself—were selected to participate in an eight-week experiment with expert systems. The volunteers were joined by a project leader and a technical expert from IBM who worked full time on the project and by six other IBM staff members who participated for varying lengths of

time. During the course of the eight-week session, the group would spend a week in training on the ESE shell and then build a prototype expert travel expense claim adviser. At the end of the session, the group was to demonstrate the fruits of its labors to the department managers, report on its experiences, and recommend whether to proceed with expert systems.

Automating the travel expense claim process appeared to be an ideal application for a first attempt at building an expert system. First, there was a crying need for improvement in the travel expense claim process, and no other attempts had really succeeded. Although the travel expense claim rules were too complex to automate using traditional languages, expert systems are capable of handling complex problems and, therefore, should be an effective way to automate this process.

Second, virtually everyone who works for the state of California has at some time grappled with the mysteries of filling out a travel expense claim and, thus, would appreciate the benefits of an automated travel expense claim application. Further, the travel expense claim application is considered to be a good example of the kinds of things that ought to be automated in state government, and because the travel expense claim process would be so familiar to many employees, it should readily serve as an inspiration for other expert system applications. Therefore, a working automated travel expense claim adviser would not only have broad appeal statewide but would also serve as a convincing ambassador for expert system technology.

Third, Bohart had the support of DPA, which loaned the state's top expert on travel expense claims, Diane Hachey, to the project. Hachey was (and still is) enthusiastic about the project. She was available full time throughout the course of the eight-week session, has devoted a great deal of her own time in support of the project since this time, and has become one of the most outspoken proponents of the system. Travel expense claim rules are, for the most part, created by Hachey and are set forth in about 50 pages of the *State Administrative Manual*. Because Hachey was involved with the project, any problems with rule interpretation could easily be resolved. With travel expense claim knowledge so easy to access, the knowledge-acquisition phase of the project would be relatively straightforward, therefore reducing the risk of failure. Also, as the recognized expert on travel expense claims, Hachey would lend her unquestioned credibility to the system (in fact, she encourages its developers to promote the system as "Diane in the computer"); this approach would ease the task of selling the system to other departments as well as SCO.

Finally, the overall risk associated with this project was low. First, if the experiment failed to produce a viable application, the loss in terms

of time and effort invested would be acceptable. A major goal of the project was simply for the group to learn about the strengths and weaknesses of the technology; if the group had been unable to produce a good system, then the understanding that it would have gained about why would still justify the effort. Second, the application itself was not strategic to the state. If it did not succeed, business could continue as it had in the past; this case did not involve some external factor (such as legislation) that was forcing the development of the application. Third, the technical risk was low: The application did not require access to, or modification of, any other existing automated systems and, therefore, would not interfere with or jeopardize normal operations in any way.

Thus, implementing an expert system to automate travel expense claims was a good choice for a first attempt at building an expert system. The project was highly justified and relatively low risk. If it succeeded, its product would be extremely valuable, and its developers would be heroes, but if it failed, the memory of the attempt could fade quietly into obscurity as everyone returned to his(her) former duties.

The team set to work. Team members divided themselves into two subgroups because it was felt that nine people were too many to be working on the same application. One group continued with the travel expense claim project, and the other group selected another application to prototype. The first action taken by the travel expense claim group was to select a name for their application: CALTREC, which stands for California travel expense claim system. The group then selected a team leader, developed an overall design, and divided the work among the team members. By the end of the eight-week period, a reasonably stable prototype had emerged. Not only did the prototype contain a knowledge base, it also had interfaces to working external programs that printed the travel expense claim produced by the expert system and manipulated start and end dates and times given by the user into forms that the knowledge base could use. Also, the knowledge base interfaced to a DB2 database that stored personal information (name, address, working hours, and so on) about the claimant.

The prototype was demonstrated to management, and despite some rather shaky moments, the system was well received. Bohart considered it promising enough to found a knowledge-based system section at HWDC so that work could continue on CALTREC and (eventually) other expert system projects. One of the participants in the experiment, Jim Henderson, was chosen to lead the new knowledge-based system section, and the IBM technical expert who had been involved in the project (the author) was hired a few months later.

The knowledge-based system section set out to complete CALTREC. Although the knowledge base was substantially complete, there was much

work remaining on other aspects of the application. This work predominantly consisted of traditional data processing activities, such as writing programs in traditional languages, writing job control language, and building databases. Both members of the knowledge-based system section have extensive backgrounds in IBM mainframe data processing, and this experience proved to be essential to the completion of the project. About three fourths of the time spent on the CALTREC project was devoted to this work, which is outside the scope of what is considered to be knowledge engineering; these tasks are described later. The fact that so much of the work involved in this project fell outside the bounds of knowledge engineering is significant considering that this application is expected to be typical of expert system applications implemented in the state of California's IBM mainframe environments. This situation indicates that a major amount of traditional data processing effort will probably be required in any expert system project in this kind of environment, not only within the state of California but in any organization that uses this type of hardware and software configuration.

The first effort involved exhaustively testing, debugging, and enhancing the prototype system's interfaces to external routines, which were sketchy, incomplete, and not robust enough to support the release of the system to the general user population. Second, although the knowledge base was complete from the standpoint of containing most of the rules covering travel expense claims, it lacked sufficient control structure to allow the production of claim forms that include more than one trip for each claim (a common occurrence). Implementing this format entailed restructuring the knowledge base and enhancing the print routine used for the prototype.

Third, the system was tailored for delivery under HWDC's office automation system. This system, IBM's PROFS, runs under the VM operating system on a large mainframe and is available to over 3,000 users. Because IBM mainframes are the predominant mode of data processing in the state of California, with PROFS being the major office automation tool, deploying the application under PROFS assures the widest possible exposure to potential users. (Deployment on personal computers was never considered because of the problems anticipated with version control for such a volatile application.)

Integrating CALTREC with PROFS entailed a great deal of thought and some additional nonknowledge-engineering effort. Past exposure to typical users of data processing systems had equipped the knowledge-based system section with a good understanding of the type of person who would use the CALTREC system and what this person's needs would be. PROFS users and, in fact, most of the intended users of the CALTREC system fit into this category. These people are generally not sophisticat-

ed computer users; they are accustomed to being guided through the tasks they perform on the computer and do not, in general, learn to use additional commands or functions outside the scope of their normal activities, nor do they wish to learn the underlying structure of the applications they use. Therefore, CALTREC must be simple, easy to use, self-explanatory, and look as much like existing PROFS applications as possible. PROFS is organized as a hierarchy of menus from which the various functions available are selected by using program function keys; so, CALTREC was added as yet another option on a PROFS menu. To further buffer the user from all the steps required to bring up the CALTREC application and to protect the user from possible error conditions, a program was written that is invoked when the PROFS travel expense claim option is selected. It verifies whether the user has sufficient disk and memory space to run the system, then displays screens that offer tutorial information and provide the option to actually invoke the expert system. Once the expert system session is complete, control is returned to the invoking program, which deletes work files that are created during the course of the consultation and returns control to the PROFS system.

The fourth area of effort involved making the CALTREC consultation itself as friendly as possible, which required extensive work on the screens presented by the expert system. Although building screen interfaces is really beyond the scope of implementing the knowledge required for the CALTREC application, effective screen interfaces are essential to the success of the application. No matter how well the system represents the knowledge involved in creating travel expense claims, CALTREC would not be useful (or used) if its audience could not understand how to use it or what it was talking about. The screen-generation facility of the ESE shell was used to provide a user interface that is consistent with other PROFS applications and that provides a great deal of explanation, in familiar language, of what is expected from the user. Now, users feel that they do not have to do anything special to run CALTREC. In fact, most users are completely oblivious to the fact that CALTREC is an expert system or that such things as expert systems even exist; CALTREC is just another tool that helps them do their jobs.

Of all the tasks involved in the CALTREC project, the fifth task, implementing print support for travel expense claims, was by far the most difficult and time consuming. The travel expense claim form is detailed and complex; it must be printed on a laser printer. Clearly, if CALTREC is to be of any practical use, travel expense claims should be printable on whatever laser printer happens to be available at any user's location. However, the various state departments have many different kinds of laser printers; there is no standard for statewide printer use.

Even among HWDC user departments, no standard printer exists. In addition, different brands of printers, even different models made by the same vendor, have unique protocols for receiving, formatting, and printing data. These protocols vary widely, such that there is no standard way to interface with the multitude of printers. This situation meant that a unique interface had to be written to support the printing of CALTREC claims on each type of printer, which would not be easy.

The prototype version of CALTREC interfaced to a printer that happened to be available at HWDC, which, as luck would have it, was the only one of its kind in the state. In the early releases of the system, all output was printed on this one printer and then routed through interdepartmental mail to the user's site. As the project progressed, individual interfaces were built to support what were believed to be the most widely used laser printers, which turned out to be various models made by IBM and XEROX. Then, software was acquired that supports a common interface for IBM and XEROX printers. This software was of great interest to the knowledge-based system group, but it felt that learning the ins and outs of print software was straying too far from its knowledge engineering domain, and by this time, they were too busy to devote the time. Luckily, one of HWDC's technical support groups agreed to attempt, as an exploratory project, to produce a common print interface for CALTREC that allowed printing on both IBM and XEROX laser printers, with the idea that what it learned could be applied to other applications as well. This project was a success, and now, all CALTREC forms are being printed by one interface. The system supports printing on about six IBM and XEROX laser printers, some of which are inexpensive. Potential CALTREC users are required to either acquire one of the supported printers or have their claims printed at a central location and routed to them through interoffice mail.

The final area of effort outside the knowledge engineering domain was a consequence of the use of a technology that was new to the mainframe computing arena at the time. Much time was devoted to the trials and tribulations of using early releases of software, which is simply not, in general, as robust and well debugged as it might be later. The team was forced to devote a great deal of time to resolving problems with the ESE shell and its particular implementation at HWDC, which tended to frequently terminate abnormally and perform inconsistently. These problems were severe enough to delay the deployment of CALTREC and undermine the confidence in expert system technology of at least one department that participated in the original 8-week experiment. IBM worked hard to resolve the problems but had some difficulty responding quickly with fixes. A great deal of the onus for troubleshooting and problem determination was placed on the knowl-

edge-based system group, and there were occasions when applying a fix would cause a whole new crop of problems to arise. This situation was compounded by the fact that both members of the knowledge-based system section were new to HWDC and had to learn the organizational protocols for getting the support they needed to get fixes applied and tested. All in all, 10 months elapsed before all needed fixes were applied and in production, during the course of which over 3 person-months of effort were expended.

The dissatisfaction with the unreliability of the ESE shell, along with the realization that CALTREC's functions were severely limited by the shell's relatively unsophisticated knowledge representation techniques and problem-solving paradigm, led the knowledge-based system group to search for an alternative shell. In October 1989, the KBMS (knowledge base management system) shell was acquired, and the knowledge-based system team immediately set out to reimplement CALTREC in KBMS. This effort was not without its own set of problems; KBMS turned out to be no less bug free than ESE.

In addition to working on the CALTREC project, the HWDC knowledge-based system section has devoted time to various other activities. Among its major goals is to stimulate interest in the use of expert systems by the state and to promote the use of CALTREC. To this end, members of the knowledge-based system section are continually giving presentations, making themselves available to almost anyone who expresses an interest in what they are doing. Their audiences have ranged from informal groups of 1 or 2 people to a group of 400 at a government conference. The most common presentation introduces the CALTREC application to potential users who generally have no interest in expert systems but are involved in processing travel expense claims. A second presentation is directed toward management and potential application developers; using CALTREC as an example, the presentation is an introduction to expert systems and demonstrates why they are important. The knowledge-based system section also presents occasional tutorials that describe some aspect of expert systems in depth. Finally, the section members have consulted with application developers and management on the suitability of potential applications to be implemented using expert systems. Several applications have been identified, and recently, the knowledge-based system section began working with novice developers to facilitate all aspects of building knowledge-based system applications.

Results

Nearly 3 years have passed since the initial 8-week expert system experiment began. Except for occasional minor additions, the CALTREC knowl-

edge base has been complete since early 1989, although only in the last few months have print interfaces been available to allow the printing of CALTREC claims on the major printers used outside HWDC. The ESE version of the CALTREC system has been available to the entire HWDC PROFS user community since August 1989, and about 400 individuals have used it. With the advent of KBMS and the plan to reimplement CALTREC in KBMS, SCO has become involved with the project. As the control agency with ultimate approval authority over travel expense claims, SCO has the greatest interest in the proper functioning of CALTREC. SCO has trained 2 people in the use of KBMS, 1 of whom is now working full time on the conversion project. When the conversion is complete, responsibility for, and control of, the application will be transferred to SCO. This sort of cooperative effort between 2 unrelated departments to produce software that will be distributed statewide is unprecedented in state history. Thus, expert system technology has not only enabled an entirely new set of applications for the state of California but has inspired new ways of building and distributing applications.

The response to CALTREC has been overwhelming, much larger than expected. Word of CALTREC's existence spread from person to person and through an article in an employee newsletter, resulting in many requests for demonstrations of the system. In the beginning, demonstrations of CALTREC were given with the objective of generating interest in expert systems and inspiring members of the audiences to build their own expert system applications. However, to the surprise of the knowledge-based system section, the audiences often became so excited about CALTREC itself that they paid little attention to the topic of expert systems; instead, they clamored for access to CALTREC, even though the knowledge-based system section did not consider the application to be ready for release. In one case, a large department that is not one of HWDC's customers has literally been making a nuisance of itself by demanding access to the system.

Initially, CALTREC was released to about a dozen users at HWDC. These users were administrative staff members, who had been generating travel expense claims for the other members of the department, and managers, who had been passing along the details of their trips to administrative staff members for processing. Interestingly, at first, administrative staff members resisted CALTREC. They were confident that they already knew all the rules of processing travel expense claims and felt that going through the expert system consultation took longer than simply typing the forms. However, they grudgingly used CALTREC. Then, what has come to be a common occurrence began to take place. The knowledge-based system section would cheerfully be notified that CALTREC made a mistake: Such-and-such an expense was or was not sup-

posed to be allowed, and CALTREC was not or was allowing it. The knowledge-based system section would then scramble through its documentation, run tests, and call the expert, only to discover that CALTREC was correct and that the person who identified the so-called mistake was simply not fully aware of the rules. The first few occurrences of this scenario were met by the knowledge-based system section with great excitement (and relief), but it didn't take long before users just accepted the results given by CALTREC. In fact, CALTREC is trusted so highly now that in the event of a question, the burden of proof tends to lie on the accuser; CALTREC's word is taken for granted as correct. However, every time a new group of users is given access to CALTREC, this chain of events occurs again (except that now the anxiety level of the knowledge-based system section is much lower). Recently, a woman who had been doing travel expense claims by hand for a large office for 20 years was surprised to be shown by CALTREC that she had been doing them wrong.

The real benefit of CALTREC is for accounting staff members who must audit travel expense claims. Incorporated into the CALTREC system is an innovation called the "To Do List," which appears as a second page of the claim. This list details all signatures, receipts, and other documentation required for the expenses being claimed to be approved for reimbursement; the lack of such documentation is a major cause for delays in travel expense claim reimbursement. Because of the "To Do List," the claimant now knows exactly what to attach to the claim and is much more likely to submit a complete claim. Likewise, the auditor can now simply verify that the items listed on the "To Do List" are complete without having to remember the rules involved. Auditors no longer have to check for addition and other basic errors and can concentrate on more sophisticated issues. The audit staff at HWDC quickly began to operate under the assumption that CALTREC is always right and subjects CALTREC claims to rather cursory audits. The auditors are also enthusiastic proponents of the system. Other auditors in other departments have also been enthusiastic about CALTREC and often become convinced that they cannot survive another second without CALTREC.

CALTREC's expert, Diane Hachey, has been essential to the acceptance and success of CALTREC. Hachey is highly respected by both the accounting staffs and the travel expense claim audit division at SCO and is in continual contact with accounting staffs in all departments. She also regularly teaches classes to accounting staffs on the travel expense claim process. Hachey is a classic example of the overworked expert. She has a great deal of difficulty getting what she perceives to be her real work done because she is constantly being hounded by people with routine questions, and she immediately saw the potential benefits

of an automated travel expense claim system. She uses every opportunity to inform people about CALTREC, and they listen to her: She conveys her excitement about the project with great conviction, and because she is the acknowledged state expert on travel expense claims, her statements bring a degree of credibility that would never be possible otherwise.

Among Hachey's early CALTREC converts was Jeff Braun, manager of the SCO travel expense claim audit division. Braun had reason to be skeptical about CALTREC. There had been attempts in the past to automate travel expense claims, but these products were little more than data-entry devices and, therefore, did not affect the error rates being experienced. Furthermore, it was perceived that the use of these systems actually made the auditing process more difficult because the forms generated tended to be inconsistent with the forms produced manually. Hachey talked Braun into taking a look at CALTREC; he quickly realized that this product was something that could help him tremendously. He is so enthusiastic about CALTREC that he lent his support to the suggestion that the CALTREC application be turned over to SCO and was instrumental in making this idea become a reality. Braun has stated that CALTREC output will be the only form of automated travel expense claim that his office will accept. He expects that his staff members will be able to eliminate most audit procedures for CALTREC claims that are currently being done on manual travel expense claims. Again, this situation will speed the reimbursement process and amount to a great savings of time and effort for Braun's staff.

The knowledge-based system section discovered that the benefits of the CALTREC system are more immediately obvious and understandable to those who must audit travel expense claims than to many of those who submit the claims. If the submitter is a secretary who submits claims for a group of people or is someone who travels frequently, it might actually take longer to run through a CALTREC session than to type the claim by hand. Even in the face of the dismal error statistics with travel expense claims, most frequent claim submitters believe that they already know all the travel expense claim rules, and it can be difficult convincing them that using CALTREC is of any advantage, although, as previously noted, most users of CALTREC discover that their knowledge of the travel expense claim rules is incomplete. The claimant who simply gives his(her) travel data to someone else who then creates a travel expense claim might or might not be willing to use CALTREC, depending on how effective the person actually submitting the claims is as well as how willing the claimant is to try new things. In addition, some people seem to believe that using a terminal is beneath them and will persist in having a secretary produce their travel

expense claims. However, CALTREC should be of great use for the average infrequent claimant who must complete his(her) own travel expense claim and is painfully aware that s/he does not know the ins and outs of doing it correctly. This person generally spends a great deal of time trying to find out how to correctly fill out the form and then, because of inevitable errors in the claim, is subjected to many delays before s/he is actually reimbursed. CALTREC will significantly speed the travel expense claim process for this claimant.

Not only is CALTREC viewed as a new way to do an old task better, but it has stimulated some ideas about how to improve the entire travel expense claim process. The CALTREC system was implemented with the objective of simply replacing a process that was being done manually, that is, producing a travel expense claim document, which is a physical piece of paper. CALTREC streamlines the process by formatting and printing the claim and storing personal and claim information on a database to avoid inputting this information more than once. Such information could be used to generate statistics that could be useful in a number of ways, and accounting staffs are interested in this possibility. Also, claims stored electronically could be passed around electronically, reducing the paper flow among destinations. A second new idea inspired by CALTREC is to automatically generate account coding that is assigned by accounting departments to individual expense items on claims. These data are currently being assigned by hand, then input into accounting systems. If CALTREC could produce the account codes, error rates in this area would also be reduced. A number of questions and problems need to be solved before these two ideas can come to fruition, but the fact that CALTREC can not only do its own job well but also inspire new, better ways of doing other facets of the travel expense claim process is encouraging indeed.

Conclusions

CALTREC was implemented in a large, centralized data processing environment (as opposed to using smaller systems) by data processing people for some important reasons. First, incorporating expert system technology into the existing data processing environment is probably the most effective way for state administrative government to make use of expert systems. This approach assures the widest possible user audience and, therefore, the largest possible impact and also affords a great opportunity to leverage existing applications with expert system technology. Second, a major benefit of implementing on a centralized system is that even though there might be frequent updates to the application,

every user will be assured of using a current, up-to-date version of the knowledge base. This ability would be unlikely if the application were to run on smaller, more numerous systems. Whenever rules tend to be volatile, the need exists for a centralized control strategy; this characteristic is expected to be common of state expert system applications. (Incidentally, the level of sophistication of these expert systems in AI terms will probably be low, which is not necessarily a problem because these applications have tremendous potential for benefit and relatively low risk and could not have been implemented using traditional tools).

Thus, in this kind of environment, data processing skills are as important (or more so) as knowledge engineering skills. Building the CALTREC knowledge base required only about a quarter of the total time required for the project. Much of the other work that was required must be classified as ordinary data processing activities as opposed to knowledge engineering as such. All this extra effort was required, as previously described, to offer the application in a form that is understandable and palatable to its intended users. The application would have virtually been useless had it not been integrated into the existing user platforms and had it not supported printing on a wide range of printers, no matter how elegant or effective the knowledge base was. In addition, the reality is that the skills required to accomplish these tasks are outside the scope of what is considered to be knowledge engineering work. Thus, given this scenario, the involvement of the data processing staff in the expert system project is essential to its success. However, data processing people are valuable for more than their technical skills because these people are the ones who have been building and maintaining the organization's automated systems; they should be able to offer real insights into opportunities for effectively using expert system technology. Finally, these people are professionals at managing automated system projects: They know what it takes to get a project done and can apply the same skills to knowledge-based system projects.

Despite its problems and pitfalls, the CALTREC experiment unequivocably proved that there is a place in California state government for expert systems. CALTREC successfully automated an application that could not be automated in the past using traditional techniques and is paving the way for other expert system applications: State departments have begun using the technology now that it has been proven by CALTREC. The horizon for expert system use by the state of California looks virtually limitless.