ADJUDIPRO

J. P. Little and Mark Gingrich, United HealthCare Corporation

United Health Care (UHC) is a diversified managed care company, serving virtually every segment of the health-care market. Founded in 1974, UHC has become an industry front runner by offering a broad portfolio of products and services to help clients, both providers and purchasers of health care, successfully manage their health-care costs. UHC's highly integrated range of products and services includes a network of owned and managed health plans as well as specialty companies that individually address key and fast-growing cost areas of health care.

UHC's unique business mix provides an ideal research and development environment for developing, testing, and refining health-care–management products and service. Recognizing the power of knowledge-based systems, UHC's senior management hired JP Little to investigate and apply expert systems and other advanced computing technologies to the health-care industry. Success was top priority for the selection of an initial application. There were many possibilities, but the domain of claim adjudication was selected as the best choice.

Problem

Over the years, a complex coding scheme has emerged between physicians and the insurance industry for explaining the medical services

provided to patients. The American Medical Association has enumerated most medical procedures and services, along with variations, in a book called the Physicians' Current Procedural Terminology (CPT) (Kirschner et al.), which is published annually. CPT is written by physicians, and they agree to code according to the guidelines stated within it. Because of the complex coding, many physicians, whether intentionally or accidentally, submit incorrect claims. For example, many times a group of procedures will be billed separately when a single comprehensive CPT code actually exists for the entire group. This practice is known as *unbundling* and costs insurance companies large sums of money because the comprehensive code usually pays less than the sum of the separate procedures.

At UHC, a complex information system known as COSMOS processes medical claims submitted by physicians and hospitals. COSMOS suspends (puts on review) many claims because of high claimed amounts and procedure-payment complexities. Departments of medical analysts must manually adjudicate these suspended claims by first classifying or categorizing the claim and then applying appropriate coding and payment guidelines (rules).

The medical analysts are typically registered nurses, have several years of clinical experience, and are hard to replace. It generally takes three to six months of training before the medical analysts achieve full productivity. As UHC grows, the work load of the medical analysts is continually increasing; new plans appear annually. Despite the increased work load, contractual agreements require expedient claim processing. To compound the problem, management pressures exist to minimize additional staffing, maximize cost savings, and standardize adjudication policies across the plans.

Additional automation of the adjudication decision process seemed necessary, but COSMOS is a large table-based COBOL system. The complexity of adding more tables for additional associations (rules) or recoding the inference process of COSMOS did not seem viable solutions for a system already so complex that it was difficult to maintain. The integration of a rule-based inference mechanism was deemed a better approach (because it was more flexible and allowed for a higher level of abstraction). Coding and payment guidelines fit well into rules, an medical procedures fit nicely into class hierarchies. COSMOS remains the vehicle for placing claims on review.

Finally, the strategic objectives of all management information system (MIS) activities at UHC include (1) managing the medical loss ratio (the cost of processing claims versus premiums received); (2) containing selling, general, and administrative expense (SG&A); and (3) selling systems externally (when it makes sense). ADJUDIPRO was de-

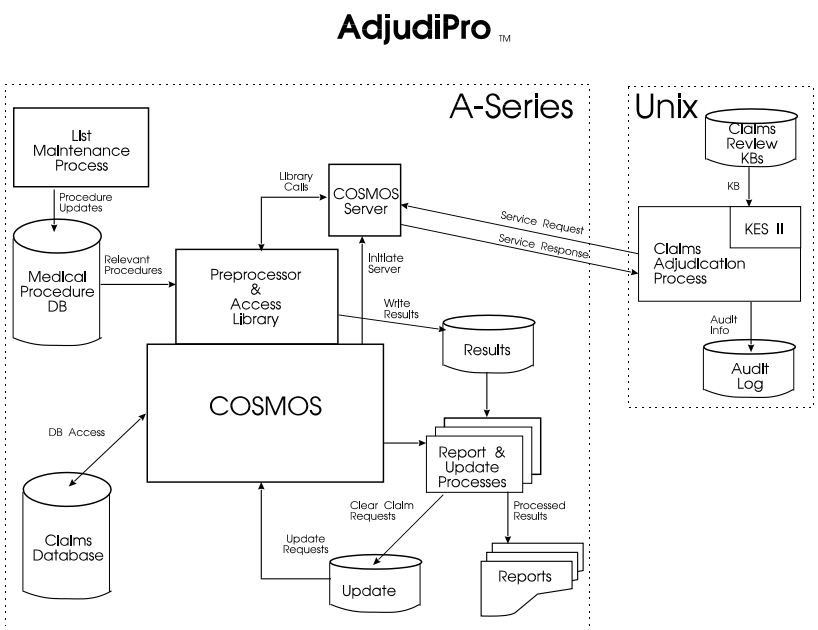

Figure 1. Current System Design.

veloped with all three objectives in mind, and the first two have been met.

System Design and Implementation

ADJUDIPRO (until recently known as the CARE system) is an intelligent augmentation to COSMOS running on a Unisys A17L mainframe. This system runs as a batch process during off-peak hours, autonomously clearing claims that the system understands 100 percent and provides medical analysts with a series of reports for those claims that still require manual intervention.

The system is logically composed of four main components: a preprocessor for collecting claims and relevant medical information, an embedded expert system component for analyzing the information, a system access component for retrieving historical claim information, and a report generator for presenting the recommendations of the expert system. Figure 1 depicts the current system design.

The principal component of ADJUDIPRO is the expert system component that is shown in figure 1 as the claims adjudication process. This component comprises the expert system shell KES-II (Template Software

of Herndon, Virginia) and C code for accessing the other components of the system. A library interface to the knowledge bases provides the C code with the ability to service the information needs of the knowledge bases.

KES-II is a C-based expert system shell that runs on all relevant UHC hardware platforms (that is, personal computers [PCs], RS/6000s, a Unisys A17L, and an IBM ES-9121), with a graphic user interface supported in the PC and RS/6000 environment. The fact that knowledge bases are 100-percent portable across these platforms made KES-II the best choice. This choice enabled knowledge base development and testing on PCs and UNIX systems, and knowledge base deployment occurred on the A-Series mainframe and UNIX systems.

The ADJUDIPRO knowledge bases primarily use a backward-chaining approach to analyze the claim information. The highest-level goal is to obtain adjudication advice. At subgoal levels, ADJUDIPRO attempts to identify situations that cannot be adjudicated before attempting to deny or accept claims. These situations and important information (that is, historical claims) are logged for the reporting component. If claims are fully or partially adjudicated, reasons for decisions and payment corrections are logged for the reporting component.

The rules of the ADJUDIPRO knowledge bases fall into five categories: *ignore rules* for identifying information ignored by medical analysts, *medical analyst rules* for identifying situations that must be handled by a medical analyst, *denial rules*, *accept rules*, and *payment rules*.

Besides rules, the problem domain called for extensive use of the classification capabilities of KES-II. CPT structures medical procedure codes in a hierarchical fashion by systems of the body (for example, cardiovascular system, musculoskeletal system), which fits neatly into a KES-II class hierarchy. Also, the medical analysts' visual and mental views of the claim structure are modeled with KES-II classes (lines, claim reviews, history claims, and so on).

The system currently shares 80 classes and 10 demons between 8 different review knowledge bases. We use the demons primarily for initialization and termination processing for a claim, but the encapsulated business and medical knowledge is represented with rules. The largest knowledge base contains 168 rules (multiple surgical review), and the smallest contains 32 (radiology review). The total current number of rules between the 8 knowledge bases is 461.

ADJUDIPRO has evolved into an interesting system with roughly 26,125 lines of code: 18,000 procedural (ALGOL, C, COBOL) and 8,125 rule-based lines. Most of the procedural code was written to integrate ADJUDIPRO into the production database world and to deliver the information to the medical analysts through reports. This code is critical to the

success of the system. As a stand-alone prototype, ADJUDIPRO received little attention, but as an integrated piece of the information system, it has become a required tool of the medical analysts. In today's environment, AI systems cannot survive as separate entities. If one wants to build a system that is going to have a significant corporate impact, this system must interact with corporate data.

The procedural code has formed a platform for future expert system (as well as client-server) applications at UHC. Until recently, the entire system ran on the A17L; however, because of performance considerations, the system was reimplemented as a client-server application with the computation-intensive expert system component residing on an IBM RS/6000 Model 550 running AIX. Without the client-server architecture, the operations management would not allow ADJUDIPRO to run against more than 3 of UHC's 19 health plans. This architecture reduced A-series processing requirements significantly. Initial estimates indicated the following A17 reductions: central processing use, 99 percent; input-output, 69 percent; elapsed time, 86 percent. This reduction was more significant than anticipated and is an exciting testament to the virtues of client-server.

The server process services the requests (for example, next claim, history, results) of the expert system component using the facilities of a utilities library. Logically, this library is composed of the preprocessor and system access components of ADJUDIPRO.

The server uses the preprocessor to extract the raw data for multiple surgical claims on review, appends the declarative knowledge of the medical analysts, and formats the information into KES-II syntax. Figure 2 shows a sample claim. The initial intent of the preprocessor was only to format claims data for use by the expert system; however, it quickly became apparent that vital information for claims adjudication appeared on a series of cheat sheet lists that medical analysts use daily. Also, the medical analysts identified many attributes of procedures in CPT as necessary to their decision process.

This declarative knowledge is now automated in the preprocessor. Offline, a maintenance process takes a set of lists containing procedures and their attribute values and updates a medical procedure database that is indexed by medical plan ID, procedure ID, effective date, and attribute. The preprocessor uses the database to supply the expert system component with the appropriate procedure information for each claim. Responsibility for list maintenance is currently being turned over to the users. As needed, new attributes (lists) are added to the system, requiring only modification to the knowledge base.

The server process uses the system access component when the expert system requests additional information while inferencing. De-

pending on the request, a response might return anything from a Boolean value to several database records. Often, more than one request is required to complete adjudication of a claim, but by only accessing data on an as-needed basis, excessive data access and processing time are spared.

The system access component also logs the expert system's results. The report-generation routines use these results to create seven reports for the medical analysts. These reports replace a single old report containing a long list of claims on review (the Age Report). One of the new reports, the Knock-Out Report, lists the claims ADJUDIPRO cannot yet handle (those with only 0- to 20-percent understanding), and the Clear Claim Report lists the claims that ADJUDIPRO can solve completely (determine payment and clear the review). The remaining new reports take the medical analysts 30 to 80 percent of the way toward adjudicating a claim. Also, the reports display additional information requested during the inference process, including relevant historical information. This information saves the medical analysts valuable lookup time because they need this information to manually process claims.

Once ADJUDIPRO can solve a particular type of claim completely and once medical analysts have performed extensive review and testing through the Clear Claim Report, a reporting routine writes claims of this type to an update file that COSMOS uses to autonomously clear claims. In other words, completely understood claims are removed from the medical analysts' claim inventory (claims on review).

Both the system access and reporting components were developed and are maintained with XGEN, a COBOL 4GL tool. This approach enabled the reuse of COSMOS components, also written with XGEN, simplifying development and maintenance. The group obtained an XGEN developer from the MIS group to write these two components.

The application of expert system technology to medical claims processing is still a new concept; however, the real effectiveness and innovation of ADJUDIPRO is the tight system integration to COSMOS, enabling quick and selective access to historical claim information during the inference process. Approximately 57 percent of the claims that ADJUDIPRO processes require historical information; however, if all historical information were provided for every claim, COSMOS and ADJUDIPRO performance would be degraded severely by excessive database access and information processing. The key is selective access.

The tight integration to the database system—combined with the declarative knowledge (cheat sheet information), the flexible reporting, the knowledge representations, and inference capabilities—has proven effective in helping us respond to enhancement requests. Several new reviews will be released during the second

```
AUDIT_NUMBER = “83114439”.
AUDIT_EXT = “00”.
PROVIDER_NUMBER = “0601487”. GROUP_PLAN = 01.
DATE_RECEIVED = “910717”.
DATE_PAID = “000000”.
HMO_ID = MSP.
INITIALS = “ “.
MEMBER_NUMBER = “5910000008415100”. assertclass MEDICINE = p01.
assertclass N_S_CRANIECTOMY OR CRANIOTOMY = p02. assertclass I_R_REPAIR_INTERMEDI-
ATE = p03. assertclass LINE = l01,l02,l03.
assertclass REVIEW = r1,r2,r3,r4. LINE:l01>proc=MEDICINE:p01. LINE:l01>modifier=” “.
LINE:l01>reason_code=32. LINE:l01>dos=”910309”.
LINE:l01>site=04.
LINE:l01>units=01.
LINE:l01>claimed= 135.00.
LINE:l01>copay=  0.00.
LINE:l01>eligible=  84.00.
LINE:l01>pcr=  0.00.
LINE:l01>fee_max=  84.00.
LINE:l01>icda=1.
LINE:l02>proc=N_S_CRANIECTOMY OR CRANIOTOMY:p02. LINE:l02>modifier=” “.
LINE:l02>reason_code=32. LINE:l02>dos=”910309”.
LINE:l02>site=04.
LINE:l02>units=01.
LINE:l02>claimed= 3350.00.
LINE:l02>copay=  0.00.
LINE:l02>eligible= 2001.00.
LINE:l02>pcr=  0.00.
LINE:l02>fee_max= 2001.00.
LINE:l02>icda=1.
LINE:l03>proc=I_R_REPAIR_INTERMEDIATE:p03. LINE:l03>modifier=” “.
LINE:l03>reason_code=67.
LINE:l03>dos=”910309”.
LINE:l03>site=04.
LINE:l03>units=01.
LINE:l03>claimed= 745.00.
LINE:l03>copay=  0.00.
LINE:l03>eligible=  0.00.
LINE:l03>pcr=  0.00.
LINE:l03>fee_max=  0.00.
LINE:l03>icda=1.
MEDICINE:p01>proc_code=”90620”.
N_S_CRANIECTOMY OR CRANIOTOMY:p02>proc_code=”61312”. N_S_CRANIECTOMY OR
CRANIOTOMY:p02>reviewable_code=true. N_S_CRANIECTOMY OR
CRANIOTOMY:p02>surgical_assistant=allowed.
I_R_REPAIR_INTERMEDIATE:p03>proc_code=”12037”.
I_R_REPAIR_INTERMEDIATE:p03>no_cut_reason=repair. I_R_REPAIR_INTERMEDIATE:p03>re-
viewable_code=true. I_R_REPAIR_INTERMEDIATE:p03>base_code=”12031”. I_R_REPAIR_INTER-
MEDIATE:p03>base_code_assoc=only_one. REVIEW:r1>number=38.
REVIEW:r2>number=22.
REVIEW:r3>number=75.
REVIEW:r4>number=35.
REVIEW:r4>lines=LINE:l02,LINE:l03.
%
```

Figure 2. Sample Claim, Formatted by Preprocessor for Use by Knowledge Base.

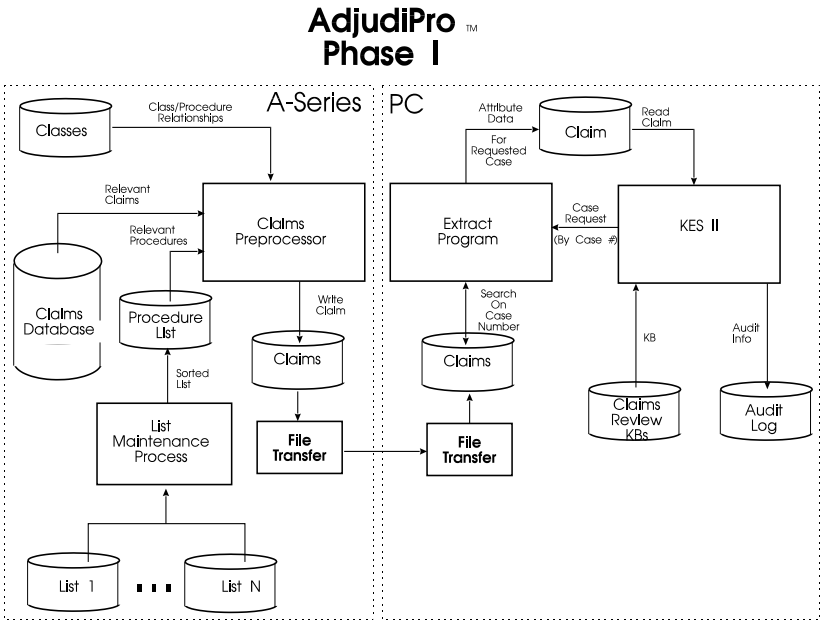

Figure 3. Phase 1.

and third quarters of 1992, bringing the total reviews to eight. This system is considered innovative in the insurance industry, and UHC filed a patent application in October 1991 to protect ADJUDIPRO.

Application Development and Deployment

This system was the first foray by UHC into AI. Although the primary objective was certainly to produce a viable system, perhaps as important was technology acceptance and the construction of a solid expert system development platform. JP Little started in January 1990 with this mission.

As previously described, there was definitely a need for further automation of the adjudication process: improving the productivity of experienced professionals, reducing the training cycle, dealing with an increasing work load, minimizing additional staffing, standardizing guidelines, and so on. The initial prototype, however, could not address the entire domain of claim adjudication. Little selected one review known as the multiple surgical review. These claims involve more than one surgery on the same date of service as well as special contractual circumstances between insurance payers and health-care

providers. This problem domain was selected because of the financial importance of paying multiple surgical claims correctly and in a timely manner. Pederson (1989) was useful for identifying the problem and narrowing the scope of the domain.

The development and deployment of the system was broken into three phases. Phase 1, figure 3, was the initial prototype phase. Phase 2, figure 4, was the first phase deployed in production. Finally, the current phase, with client-server and autonomous clearing, enabled deployment across all the plans of UHC.

Phase 1

In March 1990, Little obtained the services of Mark Gingrich of the Unisys Applied Technology Group through the Artificial Intelligence Apprenticeship Program. Together, they studied CPT, policy guidelines, and other medical information to become familiar with the domain. MEDICA (formerly PHP of Minnesota and Share/Minnesota) was selected as the initial health plan for the prototype because of its size (approximately 40 percent of total claim volume) and the availability of medical analysts.

Knowledge-acquisition sessions took place with MEDICA medical analysts. These sessions took several forms. First, observation sessions and unstructured questioning were used to obtain a better yet general understanding of the problem. The medical analysts were good at explanation through examples, where they would walk through their decision process. In time, dependency-oriented techniques from the KES-II training class (Template Software 1989) were applied. Basically, with *dependency-oriented acquisition*, one attempts to identify the solutions (or types of solutions), the information used to reach the solutions, and the dependencies between the two.

Over the first few months, the knowledge-acquisition process did not go smoothly. The elicitation of knowledge is a slow, arduous task when entering a new domain. It takes a patient and introspective domain expert (that is, medical analyst) to express his/her decision process, possibly many times. It also takes the availability of an expert. Up front, promises of time were given; however, when push came to shove, ADJUDIPRO understandably took second priority. Then, Jean Toftely, a medical analyst supervisor, was identified as the expert for the prototype. With both a manager's perspective and many years of adjudication experience, Jean's availability provided an incredible boost to development productivity.

With the acquired knowledge, an initial prototype was constructed for a subset of the multiple surgical review called the "no-brainers."

These cases required a somewhat shallow level of medical knowledge, yet their volume was large enough to warrant autonomous clearing. It was assumed that the knowledge depth would grow over time. The preprocessor was developed to extract and format pending claims from COSMOS, and KES-II ran interactively on a personal computer for validation purposes.

The validation process went much more quickly than expected. To summarize, in 4 testing sessions (of 2 to 4 hours) over a 6-week period, the expert system progressed from 0 to nearly 100-percent correctness. The users were surprised at how quickly we could react to problems and create solutions.

Phase 2

Development of the production phase commenced prior to the completion of phase 1. Full deployment on the A17 seemed the most viable solution because having all processes on one machine made integration somewhat less complicated. A C programmer was added in July 1990 to develop the KES-II support code that made up ADJUDIPRO's expert system component.

One important conclusion from the prototype was that to enhance intelligent processing, access to historical claim information was necessary. Also, if access was provided to the expert system, why not report the requested information to the medical analysts; besides, this information would be useful for validation. An XGEN-COBOL programmer came on board in November 1990 to build the database-access and database-reporting mechanisms.

Once the reports were available, the medical analysts preferred working from a known format (we emulated the appearance of the Age Report) to verify the results of the expert system. The reports greatly increased the speed of the testing-correction process and enabled the implementation of formal verification procedures. To this end, the medical analysts worked from development reports for several weeks before signing off on the transmittal to production.

ADJUDIPRO was deployed in production in March 1991. Use of the ADJUDIPRO reports completely replaced the existing Age Report in June 1991. The key aspect of the deployment was the medical analysts' confidence in the accuracy of the reports.

With this confidence came a wave of requests for new features and reviews. An operational process evolved for handling the request. First, all requests were filtered through Toftely. Next, the knowledge engineers made the appropriate changes to the knowledge bases and did some rudimentary testing. After testing, the system was handed to the C

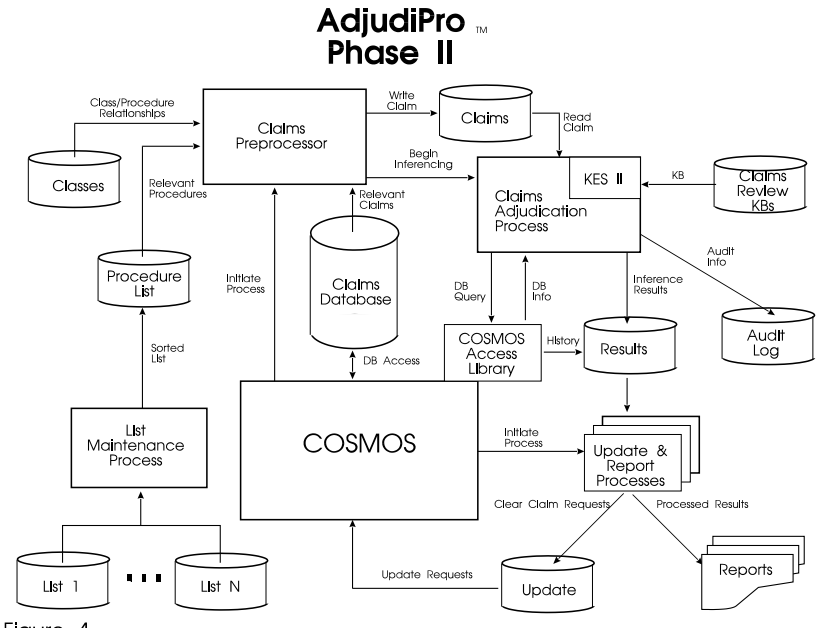

Figure 4. Phase 2.

programmers for KES-II support code modification. They reacted to new classes, types, attributes, and so on, and produced a functional embedded system. The C programmers then handed off the system to the reporting and database programmers, who reacted to new database and reporting requirements. Finally, after testing, Toftely signed off on the transmittal to production. This process is still essentially used to introduce new capabilities.

A knowledge engineer was added in August 1991 to assume maintenance and new development responsibilities for the knowledge bases.

Phase 3

Within the wave of requests, two stood out as priorities: autonomous clearing of claims and propagation of ADJUDIPRO across all UHC plans. The ability to autonomously clear claims required total buy-in by members of the MIS group because they would develop the code; that is, MIS members had responsibility for COSMOS, so they should be responsible for updating its databases. This link meant the integration of ADJUDIPRO into COSMOS was complete, and the users were the driving force that placed its implementation as a top priority. It was a great feeling.

The second major request dictated the migration to client-server.

The UHC A17L is a large mainframe; however, COSMOS and other processes keep the machine busy. Operations keeps a close eye on processor-intensive applications, and they would not allow a new implementation that would more than double ADJUDIPRO's processor requirements. Client-server not only allowed for the processing of all plans but dramatically increased the speed of development and testing by providing online access to COSMOS data.

With all three phases, the total cost of development for ADJUDIPRO thus far has been $450,000. This figure includes salary, contract services, hardware and software, the RS/6000, and mainframe software (a C compiler, KES-II, and TCP/IP). From a labor perspective, we have invested 5.8 person-years thus far in the product.

Maintenance

Bug fixing and new features are incorporated into production as previously described. The procedural code of the system has become a flexible platform where minimal changes are required in reaction to knowledge base changes.

Knowledge bases are developed and maintained using the KES-II graphic user interface on both PCs and UNIX machines. The ability to run both offline (phase 1) and online (client-server) has provided a flexible testing environment.

Domain knowledge changes over time but has been limited to annual changes to CPT and some contract changes. Such changes between 1990 and 1992 required few knowledge base changes. Use of classes and abstraction has helped ensure that we do as little maintenance as possible. The medical analysts maintain certain auxiliary information (the cheat sheets), which changes annually with CPT.

Although maintenance has not been a problem thus far, it will be an issue in the future. Currently, all maintenance is handled by the original developers. Eventually, this responsibility will be handed to another group. The lack of automated maintenance facilities is definitely a system weakness and is an area under investigation. Some possibilities include seamless association of COSMOS database information with knowledge base attributes, natural language techniques for eliciting information from CPT, and database mining using case-based reasoning and neural networks.

Use and Payoff

The client-server version of ADJUDIPRO was deployed in March 1992. Until this time, not all the plans of UHC could use ADJUDIPRO. The user

community and the Strategic Development Group are in the process of quantifying the total benefits, largely in part because they now fully understand what we can deliver and want to see more. This milestone, in and of itself, was probably the most significant in the political landscape of the project. Enthusiastic user acceptance of this new technology has created a ground swell of requests for enhancements and new applications. The users are driving the success of the project.

At this time, the most apparent quantifiable benefits of ADJUDIPRO are the cost and time savings for claims it can autonomously clear. For UHC, cost savings translate directly into improving the medical loss ratio, and personnel time savings improve SG&A. The projected cost savings for these claims are $1.6 million for 1992 and $3.4 million for 1993. A portion of the savings cannot be considered new cost savings to the company because medical analysts would have cleared these claims manually; however, there is still a time savings of 0.74 full-time equivalents (that is, 3/4 of a person) that allows the medical analysts to adjudicate other claims or perform more research. A conservative return-on-investment analysis based solely on cost savings projected over 5 years indicates a 36.3-percent return on investment.

This quantifiable portion, although not trivial, is only a fraction of the total value that will be realized over the next three to five years. No figures account for the many reviews currently in their development or testing phases, nor do the figures account for the many benefits that have not been fully quantified. These benefits primarily include the following:

Work load partitioning: The work load of the medical analysts has been partitioned to a finer level of understanding. The result is better control over the claim inventory, which is a major issue in managed health care.

Reduced online use: For claims that require checking historical claims (57 percent), ADJUDIPRO has reduced online access to the database by 83 percent. Some of the resource reduction is gobbled up during third shift when ADJUDIPRO is run, but central processing unit cycles are cheaper at that time.

Training: One of the UHC plans has used ADJUDIPRO reports to train new medical analysts and has noted that these reports greatly facilitated training.

Standardization: Prior to ADJUDIPRO, the medical analysts worked from notes in their CPT books and lists on their walls. Based on the assumption that humans make errors, this situation could be construed as a less than optimal environment for error-free work. As previously described, this information is now codified and used as input for a claim.

Architecture: The expertise of ADJUDIPRO will continue to grow as medical analysts request the automation of additional reviews. The general-purpose architecture that has evolved will provide a rapid development and deployment environment for future reviews.

Institutionalization: UHC medical and business entities are using ADJUDIPRO as a strategic tool for attacking new reviews (that is, new cost savings). Toftely was given responsibility for ADJUDIPRO as if it were another health plan that she oversees and will be accountable for cost savings.

Finally, we are in the process of positioning ADJUDIPRO as a service to sell to other managed care or indemnity companies. There are several technical and marketing challenges ahead to make it happen, but the process has begun.

Conclusion

ADJUDIPRO has opened the door for expert systems and AI technologies at UHC. Although skeptical at first, the users have embraced the technology and expect great, yet realistic results of ADJUDIPRO. New ideas for savings abound as the system and the technology gain momentum at UHC. ADJUDIPRO and advanced technologies have a bright future at UHC.

Acknowledgments

The authors of this chapter would like to acknowledge the invaluable efforts of many talented people: UHC's medical analysts; Advanced Technology's Strategic Development Team, which includes Joe Ampaabeng, Roleigh Martin, Scott McMahon, and David Williams; and Unisys's Applied Technology Group, most notably Nick Luzeski and Ed Sorensen.

References

Kirschner, C.; Coy, J.; Edwards, N.; Leoni, G.; McNamara O'Heron, M.; Pollack, A.; Ryan, C.; and Willard, D. 1992. *Physicians' Current Procedural Terminology.* Chicago: American Medical Association.

Pederson, K. 1989. *Expert Systems Programming: Practical Techniques for Rule-Based Systems.* New York: Wiley.

Template Software. 1989. Training Manual: Knowledge Engineering System. Herndon, Va.: Template Software.