

# Beyond Knowledge Representation: Commercial Uses for Legal Knowledge Bases

Surendra Dayal, Michael Harmer, Peter Johnson, David Mead  
SoftLaw Corporation, Canberra Australia

---

## Introduction

### ABSTRACT

At the 1991 Conference, SoftLaw presented a paper dealing with issues which arise in the modelling of legislation as English sentences and rules which a computer can process.<sup>1</sup> Using the techniques outlined in that article, knowledge bases may be constructed to model areas of the law, especially those concerned with public administration.<sup>2</sup> This paper illustrates the incorporation of such knowledge bases into a large scale application. This type of application may be used to drive the business of any organisation which primarily administers a large body of rules (legislative or otherwise).

Firstly, the paper gives a background description of the role played by ASSESS, a large scale application whose processing is based around legal knowledge bases.

Secondly, the system architecture of ASSESS is examined, focusing on:

- (i) the overall architecture of the application, and why that architecture was adopted,
- (ii) the structure of the knowledge based component of the application, and the reasons for that structure.

### BACKGROUND

SoftLaw developed ASSESS as a PC based decision support system for the New Zealand Accident Rehabilitation and Compensation Insurance Corporation (ACC). The system was built to process claims under the Accident Rehabilitation and Compensation Insurance (ARCI) Act 1992, and is linked to the ACC's mainframe database system. It is currently in production use and processing about 4800 claims for cover each working day. Every claim for compensation in New Zealand is processed through ASSESS.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1993 ACM 0-89791-606-9/93/0006/0167 \$1.50

The ACC administers a general accident insurance scheme, available as a right to all people in New Zealand. The ACC is New Zealand's insurer for worker's compensation and third party motor vehicle accident insurance. ACC pays long term compensation for the widows and dependants of persons who died as a result of an accident. Health care insurance is provided for injuries which arise from an accident. Due to the wide-ranging nature of its activities, the Corporation keeps records on most of New Zealand's 3 million residents.

The ASSESS system was implemented to coincide with the introduction of a new Act that substantially altered the system of compensation in New Zealand. Its implementation has enabled a re-definition of work procedures in the ACC. The ACC implemented a new PC and network environment at the same time as building and deploying the ASSESS system.

There are currently two main modules in ASSESS which model complex legislation and administrative processes:

- **Claim for Cover:** processes all claims for cover under the ARCI Act, including claims for coverage of medical fees and compensation,
- **Entitlements:** provides substantial assistance to client officers who process compensation payments.

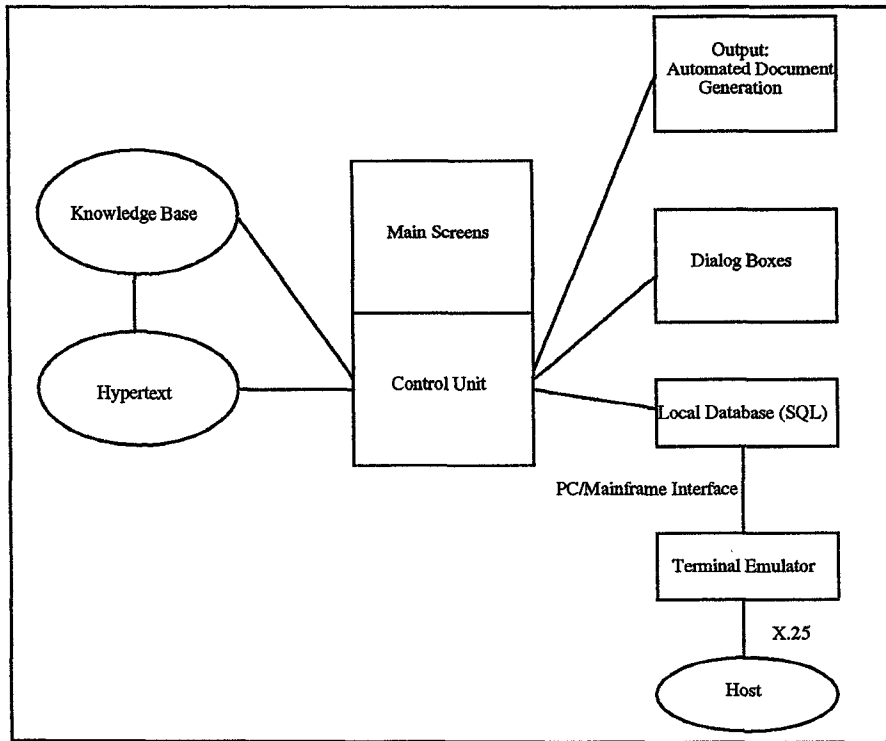
ASSESS is assuming the main automated processing role for the ACC, replacing paper-based systems, manual procedures and the dumb-terminal front-ends to the Corporation's databases. The mainframe databases remain part of a distributed processing environment, performing server functions for the ASSESS client, intelligent database interaction on command from ASSESS and batch processing for electronic disbursements.

---

## Architecture

### ASSESS SYSTEM ARCHITECTURE

Diagram 1 illustrates the architecture of the *Claim for Cover* and *Entitlements* modules of the ASSESS system. These modules rely extensively on a knowledge base to perform processing.



ASSESS System Architecture  
Diagram 1

## TOOLS USED

The *Claim for Cover* and *Entitlements* modules of the ASSESS system each consist of 15 program components, the major ones being described below.

## Expert System

The ASSESS knowledge bases are stored and processed via the *STATUTE* Inference Machine (SIM), which acts as a knowledge base server for the application. Knowledge bases are developed and maintained using the *STATUTE* Rulebase Workshop (developed by SoftLaw but not part of the ASSESS processing system).

## Textual Assistance

The *Hypertext* is used by the system in two ways. Firstly, it is accessed in conjunction with the expert system to provide source material and policy assistance for questions asked by the expert system<sup>3</sup>. Secondly, the hypertext forms a research system in its own right in the standard *STATUTE* Hypertext reader, providing access to the *Hypertext Database* which contains all the relevant legislation and policy (and soon the case law). Multiple windows can be opened on the hypertext database, which contains links between related sections and words. At present, the hypertext database contains about nine megabytes of text

and links. The text for the database is authored in Microsoft's *Word for Windows*. The text and the hyperlinks are managed by the *STATUTE* Hypertext Author (developed by SoftLaw but not a part of the ASSESS processing system).

## Graphical User Interface

The *Main Screens* program was developed using Asymetrix Corporation's Toolbook. This program is the front-end that users see when the *Claim for Cover* module is launched. It contains most of the screens for decision support functions (see diagram 3). The program also hosts the *Control Unit* which coordinates all inter-application communications.

The *Dialog Boxes* for user communication with ASSESS are written both in Microsoft Visual Basic and Gupta's SQLWindows. These dialog boxes provide data entry screens for most of the data which is stored on the Corporation's mainframe. Data validation typically occurs in code placed immediately behind the data entry fields.

## Document Generation

Microsoft's *Word for Windows* is used to generate and print all output from ASSESS. Document generation is driven from the knowledge base. A set of Word Basic macros interacts with the *Control Unit* to produce reports and letters which incorporate comprehensive publishing features such as bit-mapped graphics,

headers and footers and so on. Some of the content of the output is written by ACC staff involved in the development of the system, some parts are forms designed by professional form consultants, while other parts are reports automatically generated directly from the knowledge base.

### Database Facilities

The local *PC/LAN Database* is resident in Gupta's SQLBase. It is a relational model of a subset of the mainframe hierarchical database, with some extensions for the ASSESS system. It is accessed via an SQL Windows program, and is used to store two types of data:

- It caches the personal data that is received from the mainframe or is destined for the mainframe. This improves performance of the system, aids in robustness and provides security of data while resident on the PC.
- It stores large semi-static tables of data such as tables of employer codes that are accessed by the decision support system, but are too unwieldy to download from the mainframe while on-line.

The *PC/Mainframe Interface* program was developed in Microsoft C. It provides the communications functions to ASSESS. Using IBM's Ehllapi standard, it implements a standard transaction protocol as a vehicle for the interaction between ASSESS and the mainframe database. All interaction is via a set of dedicated ASSESS transactions. The protocol specifies workspace coding and de-coding. Data is transmitted in logical packets of no more than about 2 kilobytes in size. The communications program also provides error recovery services for such situations as line drop-out.

Eicon's *Access for Windows* provides the mainframe emulation capabilities to which the Ehllapi service is linked.

### Use of a Variety of Tools

The construction of ASSESS from a large variety of tools was one of the most contentious design issues.

The decision on whether to use a variety of high level prototyping tools such as Toolbook, Visual Basic and SQLWindows as compared to a low level language like C will depend on the circumstances of the particular application. Where time is a crucial factor, and the system is not clearly specified but will be subject to change, then the flexibility afforded by the prototyping tools will be valuable. However, where there is a more stable development environment, then a lower level language may be more appropriate. For example, the front end to a similar processing system which is currently under construction for the Australian Department of Veterans' Affairs has been constructed in Microsoft C.

In retrospect, one or two of the higher level tools such as SQLWindows or Toolbook might not have been used were it not

for the tight time frame in which the system had to be built. Ultimately *Claim for Cover* had to be designed and constructed in a period of four months. The system had to be operational by July 1 1992. The legislation on which it was to be based was in Committee stage until late March and had changed substantially from the original Bill.

However, most of the tools clearly offered essential services which could not be provided by hand-crafted software, and we continue to use those tools for application development:

- the embeddable inference engine is essential,
- the embeddable hypertext facilities provide essential access to the relevant primary and secondary material available through the system,
- the use of a word processor to manage output allows comprehensive document generation,
- a database is necessary for storing information accessed and changed by the application,
- a communications package is essential where a mainframe is used in order to maintain a link from the mainframe to the local PC.


## THE KNOWLEDGE BASES

### Introduction

The knowledge bases are central to the operation of ASSESS, and are comprised of three major rule-types:

- "Legal rules", most of which are the encoding of the Accident Rehabilitation and Compensation Insurance Act and related policy. This rulebase is constructed following the verbatim modelling techniques outlined in the last paper.<sup>4</sup>
- "Business rules" which are the organisational rules for processing claims as determined by the ACC.
- "Navigational rules" which instruct the ASSESS application on what to do next.

Conceptually we have divided the ASSESS knowledge bases into two components: the procedural rulebase and the substantive rulebase. The procedural rulebase contains navigational rules, and the substantive rulebase contains legal rules, with business rules found in both. The relationship is captured by Diagram 2:



Knowledge Base Architecture  
Diagram 2

It has been possible to model certain legal rules for some time now,<sup>5</sup> however as yet very few practical uses have been made of such models. We have found that the knowledge base structure shown above will allow practical uses to be made of legal knowledge bases in an administrative environment. Such a structure will therefore be a desirable characteristic of any processing system which attempts to model the following:

- a large body of rules (legislative or otherwise), and
- an administrative process designed to apply those rules in the real world.

In order to provide illustrations of the relationships described above, examples will be drawn from the knowledge-based component of the *Claim for Cover* module of ASSESS.

## The substantive rulebase

### Introduction

Representing law and procedures as rules is well suited to applications supporting administrative decision making. This is so because:

1. The law which defines the tasks to be carried out in the administrative environment, and the criteria upon which decisions are to be based, is usually framed in terms of *rules* - large chains of logical conditions which must be satisfied in order for a decision to be made.
2. The places in the legislation which call for discretion may be left to the officer making the decision. The main assistance offered by the application is in the reduction of logical complexity, and in finding the right rules to examine.
3. Where interpretation is required, to resolve ambiguity or procedural questions, the domain experts of the client organisation may be consulted, because it is the client's point of view which is being implemented.

### Legal Rules in *Claim for Cover*

The threshold of the ACC insurance system is the concept of *cover* for an injury. If an injury is covered under the Act, then all medical fees, treatment fees, related costs and compensation can be paid in respect of the injury. If the injury is not covered, then none of this money is available. Since it is a threshold question, every claim made to the ACC is tested to see whether the injury claimed is actually covered under the Act. All of these claims are processed using the *Claim for Cover* module in ASSESS. This module therefore determines whether each injury claimed in New Zealand is covered by the Act.

The rules for determining cover are contained in the legislation, and the *Claim for Cover* module is modelled very closely around the relevant provisions. The following provisions are modelled:

- section 3: definitions;
- section 4: personal injury;
- section 5: medical misadventure;
- section 6: work injuries;
- section 7: gradual process injuries;
- section 8: cover requirements;
- section 9: cover for injuries caused outside NZ;
- section 10: exceptions to granting cover;
- section 11: gradual process injuries before 1974;
- section 12: injuries to non-residents;
- section 13: when injuries outside NZ will be treated as injuries in NZ;
- section 63: proper claims for cover;
- section 65: work injury claims;
- section 100: the employers' fund;
- section 108: the motor vehicle account;
- section 113: the earners' account;
- section 120: the non-earners' account;
- section 121: the subsequent work injury account;
- section 122: the medical misadventure account.

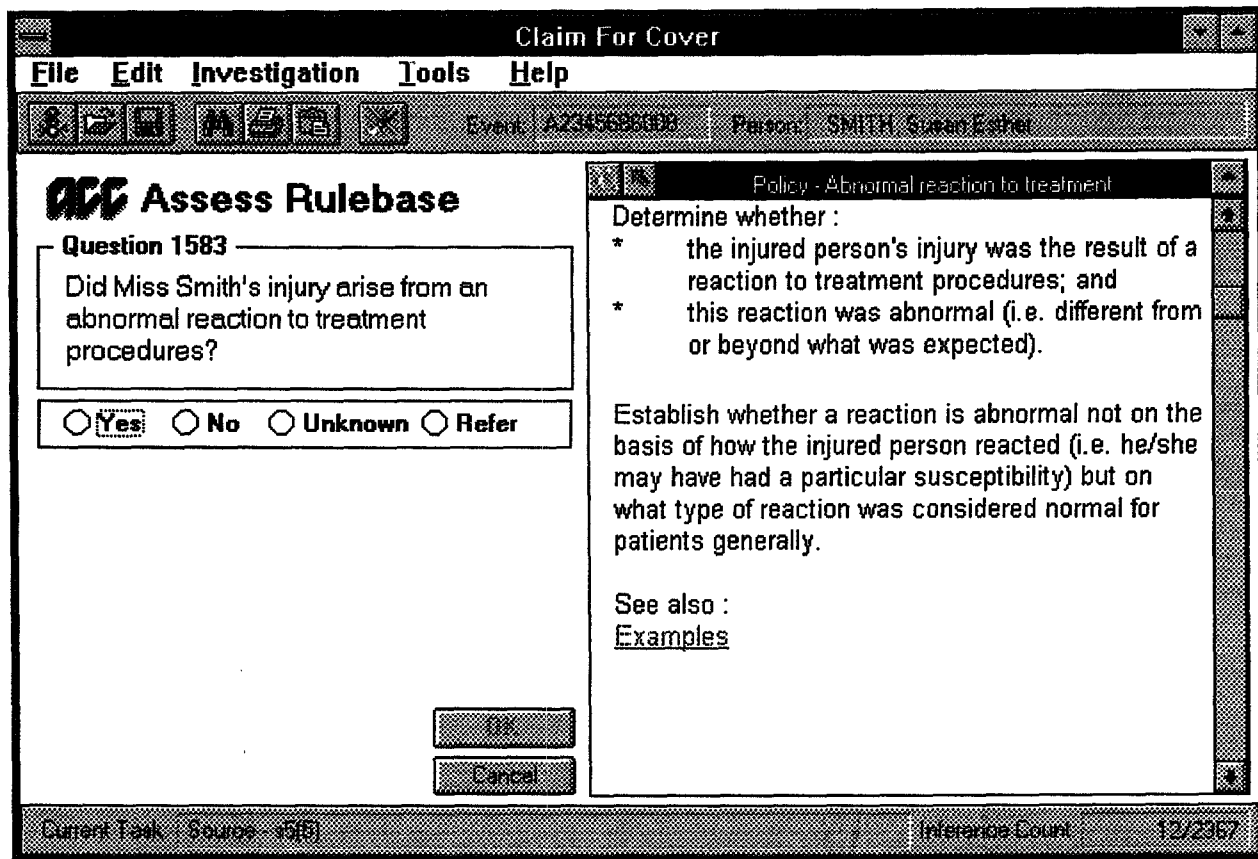
These provisions are substantial and complex, and every case requires that they be completely and exhaustively examined. The rulebase in *Claim for Cover* is composed of well over 3000 rules,<sup>6</sup> which model these provisions verbatim to comprehensively test cover and all of the other issues which must be determined at the same time as cover.

The other main component of *Claim for Cover* which is derived from the legislation is the interpretative decisions made by the top level staff in the ACC. These decisions are captured in two ways:

- In interpretative rules, which force users to adopt a certain approach to a legislative provision,
- In the hypertext assistance which is provided to the user whenever they are asked a question or given the opportunity to input data. This assistance provides guidelines on how to approach the legislative issues presented by the system (see diagram 3).<sup>7</sup>

### The need for Business Rules

Where a body of paper rules must be administered by a large organisation, in order to relate these abstract rules to the world of practical affairs, much more must be accounted for. The organisation must define how its own business will operate. The organisation must set up procedures so that cases are not dealt with in an *ad hoc* fashion, but are methodically and consistently handled. It is true that many organisations charged with the administration of rules fall well short of these goals. Nevertheless almost every organisation will have a set of rules which dictate its business. In some cases unarticulated, this set of rules will have a substantial effect on the way the legislative or paper rules are applied to the real world. When designing and constructing applications for large organisations, we have chosen to call these rules "business rules".<sup>8</sup>



The question and answer screen used by the *Claim for Cover* module of ASSESS. Many other screens are standard data collection screens with a button to access the policy assistance. Inferences occur in the background as a result of data entered on those screens.

Diagram 3

The main factors accounted for in the business rules are:

- **Information.** The collection of enough information to make decisions occupies much of the time spent by officers in large organisations. Furthermore, the collection of information is a two way process, involving communication from one party to another. Not only must the organisation process information that it collects from the real world, information must be fed back into the real world. For example, when someone fills out an application form incorrectly, or leaves mandatory information missing from the form, there must be some means of:

- (a) Communicating the lack of adequate information to the provider,
- (b) Collecting the remaining information,
- (c) Processing the remaining information once it is returned.

The rules governing the flow of information between clients and the organisation will not be spelt out in the legislation. Often they grow over time as it becomes apparent what is necessary. Occasionally, they are captured in procedures manuals or inter-departmental memos.

- **Time.** Administrative decisions are often made over a protracted period of time. Where a claim is made, or a request for information generated, time must be allowed for

the cogs and wheels to turn and responses to occur. There are usually rules governing how long it is appropriate to wait on responses, or how to prioritise tasks so that the essential ones are completed first while other less important ones are left until later. In many organisations, these rules do not find their way into print, but they exist and provide criteria on which to base decisions.

- **Authority.** In any large organisation, not everyone has the same tasks to fulfil. A hierarchy must exist to allow those with more ability and experience to handle the difficult cases, and those with less experience to deal with the cases they are capable of processing. Review is an important part of good decision making, and the accompanying responsibility must be allocated around the organisation. In practice, the rules governing *who* has authority to do *what* dictate much of the way the organisation operates.

It is possible for organisations to operate without regard to these primary business rules. However, they will not operate effectively. It is not possible for knowledge-based technology to improve the efficiency of large organisations without articulating these rules. While people may muddle along on a mixture of folk-lore and out-of-date paper procedures, an automated system requires precision and planning to operate at all, let alone effectively.

Following a verbatim modelling approach,<sup>9</sup> legal rules are actually much easier to model than the business rules which drive an organisation, since at least they are clearly articulated in documents. However, legislative rules are almost impossible to apply and embed within a large scale application unless they are modelled *together* with the business rules.<sup>10</sup> Without the business rules they have no context and hang suspended, waiting to find some portion of reality to cling to. This is why the substantive rulebases in ASSESS, and any large application modelling rules applied by a business, consists both of business rules and legislative rules.

### Business Rules in *Claim for Cover*

The complete rules for action at any stage in the administrative process are modelled as business rules in *Claim for Cover*. Most of these rules did not come directly from the legislation, but had to be elicited from high level ACC staff.

Most processing of claims for cover is carried out at one of the ACCs four registration centres.<sup>11</sup> There are a number of business tasks which must be performed on any claim for cover. These jobs, which will usually be performed at the processing centre which receives the claim, include:

- a check on whether the claim is a proper claim or not (on the right form, signed etc);
- if the claim is not proper, provision of a letter to the client rejecting the claim, in which case no further action is taken;
- collection of the medical fees number, in case some medical costs are claimed;
- a test on whether the claim is covered;
- if the claim falls within one of the categories of cover for which registration centre staff have no authority, referral of the claim to Head Office;
- if the claim is covered, determination of whether the injury is a work injury;
- if the claim is covered, determination of whether the client is nevertheless disentitled from receiving any payments;
- if the claim is covered, determination of the fund from which any money will come;
- if the claim is not covered, provision of a decline letter to the client and their medical practitioner;
- entry of all relevant details for all properly claimed cases, whether accepted, declined or disentitled, onto the corporate database;
- creation of appropriate reports to put on the file.

In some cases, the claim may not be able to be fully processed immediately, but must be delayed:

- there may be some missing information which must be collected before a final decision can be made;
- the officer may have authority to investigate the case, but no authority to confirm the decision made, in which case it will have to be referred to Head Office for confirmation;

- the officer may have no authority even to investigate the case in which case it will have to be referred to Head Office for investigation and decision.

Once the decision on cover has been made, the file is sent to one of the 35 Branch Offices located around New Zealand. If cover has been granted, this is sufficient for payment of medical fees to be authorised. The decision on cover and the details of the case will be registered on the corporate database by ASSESS. If the client wishes to claim some payment other than medical fees, he or she will have to submit an application for the payment to the Branch Office. These compensation applications can then be processed using the *Entitlements* module of ASSESS.

As can be seen from the administrative process described above, the business rules dictated by that process can be procedural or substantive in character. Hence, many of the business rules are in the procedural rulebase. The main business rules which are part of the substantive rulebase relate to:

- the authority required to make a decision on a given claim for cover, and
- the output to be generated by the system in different circumstances.

### The procedural rulebase

#### Business and Navigation

A procedural rulebase in a processing system such as ASSESS models both an administrative process and a sequence of necessary procedural steps for the application.

Many of the rules concerning which course of action to adopt will come from the business rules of the organisation. For example, in some situations it will be necessary for processing of a claim to be suspended while more information is collected, or in order to refer the claim on to a higher authority. All of these sorts of rules are captured as business rules, and form one portion of the procedural rulebase.

In addition to business requirements, much of any large application is devoted to determining exactly what step to perform next. For example, once one task is finished then another task might commence on another screen. All of these sorts of rules are captured as "navigational rules", and form the remainder of the procedural rulebase.

#### The Procedural Rulebase in *Claim for Cover*

What we have called the "procedural rulebase" in *Claim for Cover* determines the following:

- The correct path required to fully decide a given claim, and

- What step to take next. These steps include:
  - which part of the graphical user interface to enter,
  - which document to generate,
  - which portion of the substantive rulebase to investigate next.

For example, one of the core parts of the procedural rulebase for *Claim for Cover* reads as follows:

*Legislative Rule 1023-L from module Cover Control*

*The Claim for Cover has been fully investigated if:*

- The accident categories screen has been completed;
- and** The major culling screen has been completed
- and** Processing of the *Claim for Cover* need not cease until sufficient information has been provided to make a decision on gradual process
- and** The category of cover has been investigated
- and** The Date of Injury stage has been completed
- and** The location of the event has been investigated
- and** The Location stage has been completed
- and** The definition of 'personal injury' has been investigated
- and** The issue of whether the injured person's injury is a work injury has been fully investigated
- and** The check on whether it is necessary to request a work injury report has been completed
- and** The investigation of the Voluntary Cover Flag has been completed
- and** The Employment screen has been completed
- and** The Determination of Fund stage has been completed
- and** The disentitling provisions for treatment have been fully investigated

*If this is not the case, then it will be concluded that the Claim for Cover has not been fully investigated<sup>12</sup>*

## Use of a Procedural Rulebase

Currently, most major software systems in the administrative environment use application code to perform the same functionality provided by the use of a procedural rulebase. It is certainly *possible* to write application code to perform the same function as the procedural rulebase. However, especially where the application models an administrative process, there are distinct advantages afforded by harnessing the power of a good inference machine to control essentially procedural tasks:

- There is a lot of *if-then* logic in the flow of control of screens within the application. Whether dialogue box A or B is to be generated at a particular point will depend on a number of factors. The same is true of the administrative process, because the appropriate action to perform at any time will depend on specific criteria. This type of *if-then* logic is ideally suited to rulebase modelling - the rules determine the behaviour of screen units, and the correct path

to take. An example of this is the pre-conditions for sending out a request for work information in *Claim for Cover*.

*Legislative Rule 1702-L from module Cover Control*

*The work injury report stage has been completed if:*

It is not necessary to generate a letter requesting a work injury report

or

It is necessary to generate a letter requesting a work injury report

**and** The Payment Delayed letter has been generated

**and** The letter requesting a work injury report has been generated

*If all the options of this rule are disproved, then it will be concluded that the work injury report stage has not been completed<sup>13</sup>*

- The procedural rulebase allows the procedural code to be very transparent. When captured in the knowledge base, complex rules are easily readable and are gathered together in one place. This allows the flow of control to be tested and debugged reliably.
- Modelling the procedural behaviour of the application as rules allows use to be made of the powerful navigational tools which are built into an inference engine<sup>14</sup>. Most complex processes will sometimes require the retracing of one's steps. A significant functional specification that ASSESS had to meet was the ability to backtrack from one point to any previous point. Since backtracking is handled automatically by a good inference machine, a user may easily retrace his/her steps to the point at which an error was made and re-enter data at the appropriate screen.
- The use of a rulebase to model procedural aspects of the application allowed these facets to be developed quickly, reliably and by a business analyst. Since many of the procedural rules will in fact come from business analysis, this will often be desirable, allowing the business analyst to encode those parts of the navigation which come primarily from the business rules.
- Event-driven Graphical User Interface environments frustrate the use of traditional programming techniques used by application programmers to code control flow. We found that using an embedded inference engine for this task solved all of the problems associated with attempting to force an event-driven operating system environment to behave like a procedure driven system.

---

## Conclusions

ASSESS is a good example of a major production system which uses legal knowledge bases in an administrative environment. At

a fundamental level, it demonstrates that such large applications are both useful and possible.

## Benefits Offered by ASSESS

Large organisations can benefit from the implementation of the type of technological platform which ASSESS represents in the following ways:

- Since the rulebase strictly adheres to the legislation and management decisions of the ACC, there is greatly increased consistency in decision-making and in adherence to policy and administrative guidelines.
- The application provides a platform to ensure business rules and procedures are adhered to by ACC officers, resulting in substantial quality assurance and risk management benefits.
- With the centralisation of knowledge in the application, there is a greatly increased capacity to effect changes in policy and procedures in a uniform and effective manner.
- Because the knowledge base may assume bulk processing responsibility, there is an increased capacity to effect structural changes within the organisation, and to increase the structural efficiency of the organisation.
- With the logical complexity of the legislation handled by the knowledge base, the quality of decision making and client service can be improved to best practice standards.
- The system improves the capacity of the organisation to distribute processing between PC's and mainframes.

## System Architecture

Large processing applications such as ASSESS will always require the combination of a variety of tools if they are to provide comprehensive decision support. At the very least, independent tools will facilitate mainframe interaction, document generation and user-friendliness, without the developer having to build such functionality from the ground up.

The use of procedural rulebases to provide the flow of control over complex administrative processes was undoubtedly successful, and is a method which we will continue to use where appropriate in other modules. Harnessing the powerful navigational capabilities of an inference engine is safe and profoundly assists the more stable and rapid development of the rest of the system. Ultimately, this method simply uses rules for the purpose for which they are designed: the modelling and execution of *if-then* conditions.

---

## Footnotes

<sup>1</sup> Johnson and Mead "Legislative Knowledge Base Systems for Public Administration - Some Practical Issues" Proceedings

---

of the Third International Conference on Artificial Intelligence and Law, 1991, pp.108 - 117.

- <sup>2</sup> For other examples see Svensson, "Tessec: an Expert System for Social Security Legislation, in Krach, De Vey Mestdagh and Svensson (eds) Legal knowledge based systems: An interview of criteria for validation and practical use", Koninklijke Vermande BV, Lelystad (1990), pp.87 - 92, Bench-Capon, Robinson, Routen and Sergot, "Logic Programming for Large Scale Applications in Law: A Formalisation of Supplementary Benefit Legislation", Proceedings of the First International Conference on Artificial Intelligence and Law, 1987, pp.190 - 198, Sergot, Cory, Hammond, Kowalski, Kriwaczek and Sadri, 'Formalisation of the British Nationality Act' (1986) 2 Yearbook of Law, Computers and Technology, 40.
- <sup>3</sup> See diagram 3. The text on the right is provided using an embeddable *STATUTE* Hypertext object.
- <sup>4</sup> See Johnson and Mead "Legislative Knowledge Base Systems for Public Administration - Some Practical Issues" Proceedings of the Third International Conference on Artificial Intelligence and Law, 1991, pp.108 - 117.
- <sup>5</sup> For example Sergot, Cory, Hammond, Kowalski, Kriwaczek and Sadri, 'Formalisation of the British Nationality Act' (1986) 2 Yearbook of Law, Computers and Technology, 40.
- <sup>6</sup> The Entitlements rulebase is of a similar magnitude.
- <sup>7</sup> The hypertext assistance provided to the user was also subject to an independent audit by Pete Marwick (KPMG). See diagram 3.
- <sup>8</sup> See also part 2 of Bench-Capon (Ed.) **Knowledge Base Systems and Legal Applications** Academic Press 1991.
- <sup>9</sup> See Johnson and Mead "Legislative Knowledge Base Systems for Public Administration - Some Practical Issues" Proceedings of the Third International Conference on Artificial Intelligence and Law, 1991, pp.108 - 117.
- <sup>10</sup> See chapter 16 of Bench-Capon (Ed.) **Knowledge Base Systems and Legal Applications** Academic Press 1991.
- <sup>11</sup> Located around New Zealand, in Auckland, Hamilton, Wellington and Dunedin.
- <sup>12</sup> Source code from cover.drb, the rulebase used by claim for cover. (drb = Development RuleBase).
- <sup>13</sup> Source code from cover.drb.
- <sup>14</sup> For other desirable and essential facilities in modern knowledge base shells, see Stylianou, Madey and Smith, "Selection Criteria for Expert System Shells" (1992) 35 Communications of the ACM, Vol 10, pp.31 - 48.