DOCUMENT RESUME

ED 190 199 IR 008 667

AUTHOR Rosenberg, S.
TITLE An Intelligent Support System for Energy Resources in the United States.
INSTITUTION California Univ., Berkeley. Lawrence Berkeley Lab.
SPONS AGENCY Department of Energy, Washington, D.C.
PUB DATE Apr 80
CONTRACT W-7405-ENG-48
NOTE 10p.: For related document, see IR 008 664.

EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS *Artificial Intelligence; *Energy; *Management Information Systems; Natural Resources; *Semantics; Systems Development

ABSTRACT

Based on artificial intelligence research, the frame based system for reasoning described in this paper is one of the components of an intelligent decision support system for an information system on petroleum resources and use which is being designed by the Information Methodology Research Project as the first step in the development of a comprehensive intelligent information system for dealing with energy resources in the United States. By extending the notion of frames to include rule frames which can then be interpreted and applied, expertise of various kinds can be directly encoded into the frame representation. Frame based rules are useful in encoding constraints, performing actions, noticing complex situations, and deducing solutions. By varying the interpretation of a rule frame, the same competence knowledge can be used in performing each of these tasks. Rules are able to use frame based representation in finding other rules, avoiding most pattern-directed invocation. (Author/RAA)

Submitted for
ASIS Annual Conference
Anaheim, CA, October 1980

LBL 10618

AN INTELLIGENT SUPPORT SYSTEM FOR ENERGY RESOURCES IN THE UNITED STATES

S. Rosenberg

Information Methodology Research Project
Lawrence Berkeley Laboratory
Berkeley, California 94720

April 1980

Prepared for the U.S. Department of Energy,
Technical Information Center,
under Contract W-7405-ENG-48

AN INTELLIGENT SUPPORT SYSTEM FOR ENERGY RESOURCES IN THE UNITED STATES

S. Rosenberg
Information Methodology Research Project
Lawrence Berkeley Laboratory
University of California
Berkeley, CA 94720

ABSTRACT

This paper describes a frames based system for reasoning in a petroleum resources domain. By extending the notion of frames to include rule frames, which can then be interpreted and applied, expertise of various kinds can be directly encoded into the frame representation. Frame based rules are useful in encoding constraints, performing actions, noticing complex situations, and deducing solutions. By varying the interpretation of a rule frame, the same competence knowledge can be used in performing each of these tasks. Rules are able to use the frame based representation in finding other rules, avoiding most pattern-directed invocation.

Increasing amounts of information are now commonly recorded and stored in computers. There is often, however, no corresponding increase in the ability to access and use this information. In this paper I shall discuss one approach to this task which draws on research in Artificial Intelligence. Our goal is the construction of intelligent information systems for dealing with energy resources in the U.S. I shall first describe the concept of a "friendly" representation which can use the power of a frame representation in organizing information. The remainder of the paper focuses on the use of rules in such a database. We have found frame based rules to be a useful way to provide "intelligence" in an intelligent information system.

Our goals are ultimately quite practical; namely the transfer of A.I. "technology" into a real world domain. This imposes certain constraints on the design. Over 200 databases dealing with energy resources are already maintained by DOE. Our representation scheme must be able to use this existing knowledge (Rosenberg (1)). Real data in these databases is often "messy." Crucial information is sometimes missing. Of the data available, there are problems with validation, with information gaps, with variable definitions.

A "friendly" representation helps with these and other problems by taking the burden of performing routine, if sometimes complicated, functions, from the reasoning component or user. Some of these functions are quite simple, such as providing aggregated information, or default values. Some are complex, such as adjusting the representation by changing deduced consequences when erroneous facts are corrected. The net effect of such a friendly system is to allow a user to focus on higher level tasks, while leaving lower level information processing to the representation system. In effect, we propose that in many domains, semantic representations must function dynamically, drawing on interlaced procedural and world knowledge, to provide a solid basis for creating higher quality information.

We use FRL (Frame Representation Language) (Roberts and Goldstein (2)) as the basis for our representation. FRL provides a hierarchically organized, frames-based semantics with inheritance and procedural attachments among other features. This augmented notion of what a data object is allows us to create the types of "friendly" representation we need. FRL is based on Minsky's (3) notion of frames. FRL is a sophisticated, higher level language designed for the representation of knowledge in a variety of domains.

A frame can be thought of as a named collection of slots which form the semantic definition of a concept. A slot (= property) can be specified further through the use of associated user and system defined "facets." Useful system defined facets are: Value, which contains the value of that slot; Default, which specifies a default value; Require, which specifies procedural constraints on the values for that slot; and If-Added and If-Removed, which specify actions to be taken when a value is added or removed. Notice that many of these facets are procedural attachments which can perform calculations when required. Thus, if-added attachments can be used to encode and execute constraint relations between frames. For example, in the portion of a Frame for Iran shown below, the if-added attachment on the carryover slot automatically updates aggregation information on world supply. The default predicate on the production slot provides a typical value when this data is reported late.

```
IRAN   production   $default  (use-last-month-value)
       carryover    $if-added (add-to-world-supply)
```

FRL allows concepts (= frames) to be arranged in an inheritance hierarchy. Thus the frame system forms a tree structure. Generic information is stored higher up in the hierarchy and shared by frames lower down; specialized frames specify new distinguishing knowledge. The generic knowledge, including computational procedures, is inherited automatically.

Suppose an overseas supplier of crude oil, such as Iran, decreases supplies to U.S. sites. Figure 1, below, shows a portion of the petroleum flow network. In it, Iran and Dallas supply the site of Newark. We want to model this reduction in the shipment of oil. This change in production will in turn alter some of the constraints we have set in our database. We must then notice when meaningful patterns of information (such as reduced supplies at Newark) occur. Finally, we may want to find alternate sources.

To deal with these four problems of driving our model, encoding constraints, noticing the development of important situations, and deducing solutions, we have found it useful to extend the concept of frames to include rules. By creating a class of frames called rules, and varying our interpretation of these rules, we can do all four of these tasks.

All knowledge is represented as frames. Thus rules are represented as declarative knowledge in the frame tree. The only indication that such a frame is a rule consists in the value of the generic pointer (e.g., = AKO

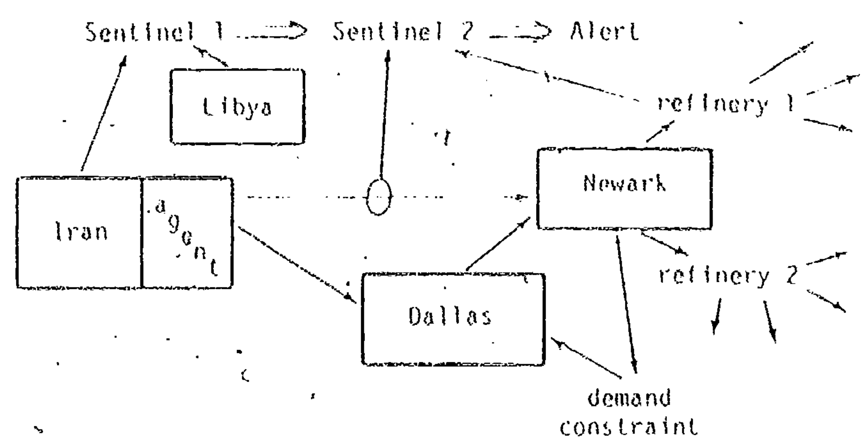

FIGURE 1

rule). Rule frames contain condition and action slots. To use rules, a rule frame is interpreted as a procedure, with the slot values controlling the interpretation. Thus a condition slot causes a condition to be tested; the action slot specifies the action to be performed, and so on.

Many of the relations between semantic entities in our model can be encoded as constraints. Changes in information can cause propagation of these constraints to occur (although in a much simpler form than Doyle (4) proposes.) Simple constraints can be encoded directly as procedural attachments to frames. A change in the information content of one frame triggers a simple predicate which then modifies the information available at another frame. (For example, the if-added procedure in the Iran frame.) More complex constraints are encoded as rule frames. Figure 2 shows the constraint of Figure 1 in more detail. Here, a generic constraint exists whose purpose is to see that the petroleum needed by a site is equal to the amounts its suppliers intend to provide. This constraint places a trigger in the generic site frame. There it monitors demand at all sites. If a particular site, such as Newark, changes its monthly needs, this trigger is inherited, and fires. The constraint then tries to adjust supply among the suppliers to Newark to correspond to demand.

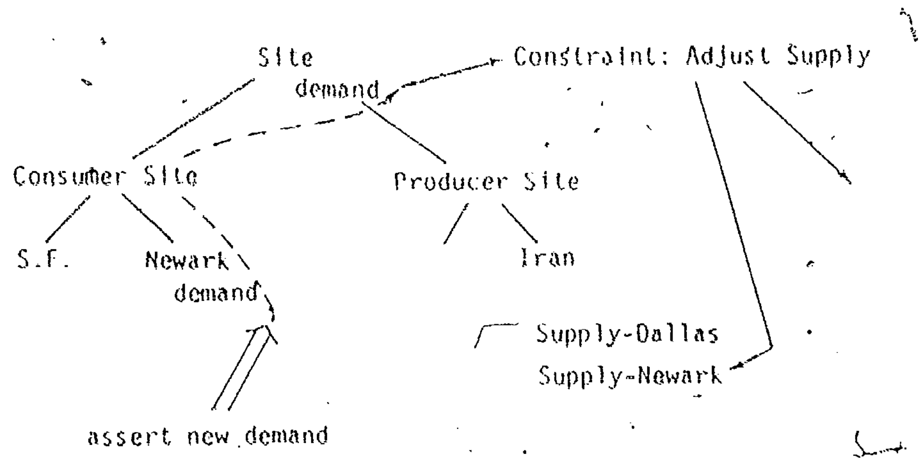

FIGURE 2

Actions such as "shipping" oil are performed by using sets of rules as agents. In Figure 1, an agent which ships oil is shown "attached" to Iran. By monitoring production and carryover in the Iran Frame the agent determines when to ship oil. At the right time the oil is allocated among the sites supplied by Iran. This action can in turn trigger a new agent. Using rules as agents provides a method for driving our model to simulate the changing state of our domain.

Given a database of changing information, we want to provide some capability to monitor important developments, and alert us when necessary. Many subtle problems can arise in providing such alerts. For example, small reductions in supply by various producers, together with changes in demand at several sites can result in a severe shortage at one particular site. Such dynamic noticing is done by treating rules as Sentinels (Rosenberg (5)) which leave active expectations in the data base.

Suppose we wish to be warned whenever a consuming site, such as Newark, will experience a severe shortfall in supply. Figure 3 shows in more detail the sentinels from Figure 1 which do this. By taking advantage of the semantic structure a frame hierarchy provides, we can create a sentinel which places a trigger in the generic site frame. This trigger will be inherited

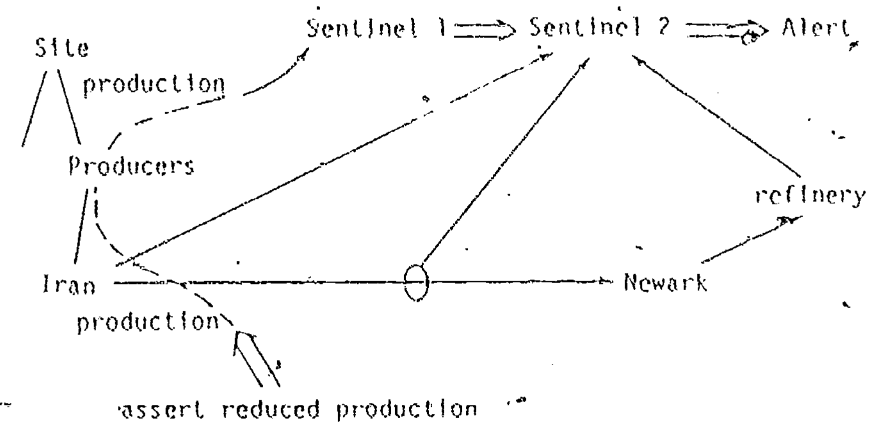

FIGURE 3

by all production site instances. Reductions beyond some local criterion expressed on the individual sites will trigger the sentinel. Thus the assertion of reduced production into the Iran frame causes the conditions of this Sentinel to succeed. A reduction of production ~~by~~ one producer however (or even several), does not necessarily mean a shortage of oil at a site. How large a proportion of its requirements are met by this producer? Can other suppliers make up the slack? Before a shortage warning can be issued, questions such as these must be considered. The first sentinel will examine the sites supplied by Iran to determine if any are excessively affected. If, for instance, Newark were solely dependent on Iran for oil, a reduction in Iranian production is sufficient evidence for a warning. If there are several suppliers to Newark, the best choice may be to monitor Newark's supply more closely. In this case, Sentinel1 creates another sentinel, Sentinel2, to monitor both shipments to Newark, and demand at Newark, directly. Iranian production is also monitored. If Sentinel2 notices a drastic supply imbalance at Newark, it will give an alert. If production returns to normal in Iran, this Sentinel will erase itself. Sentinels provide a flexible, powerful mechanism for encoding the noticing of meaningful patterns of information.

Once an alert has been given, we may need to answer questions such as those about possible alternative supplies. We are exploring the uses of frame based rules for such reasoning. While appropriate rules can be found by some variant of pattern matching (e.g., Planner, or production systems (Newell and Simon (6)) we take advantage of the organized semantic structure to have information directly trigger the appropriate rule(s). Figure 4 shows part of an adjust-supply frame. Sentinel2, on noticing an oil shortage in Newark, can assert Newark into the buyer slot of this frame. The addition of this information triggers a Rule, R1. By contrast, the assertion of a new supplier would trigger a different Rule, R2. R1 first collects all normal suppliers to Newark, and if these have sufficient stocks, adjusts supply using these stocks. Otherwise, it can try alternative methods of increasing supply. (Try supply-increase) will cause all rules which inherit from the supply-increase frame in Figure 4 to be evaluated. Although we do not know which specific rule might be relevant, we use the frame hierarchy representation to allow rules to call on other classes of rules known to be helpful in

```
adjust supply
     buyer          $value Newark ---> Rule1
     supplier
     new supplier $value Alaska ---> Rule2

Rule 1: Condition: buyer
        Action: (if((carryover(normal suppliers))
                           (demand buyer))
                      (adjust-supplies))
                (else (try supply-increase)))

Rule 2: Condition: New-supplier
        Action: Allocate new supplies

        Rule 3: Supply - Increase
           /           |           \
Rule 4: Spot   Rule 5: Borrow   Rule 6: Share
     Market      from other        shortage
                   sites
```

FIGURE 4

achieving their goals. In effect, the use of frames such as adjust supply, and a rule hierarchy, allows us to create small contexts of relevant rules and pertinent information. We have found that using a powerful representation semantics together with rules provides the flexibility and scope we need in creating an intelligent information system.

ACKNOWLEDGMENT

Prepared for the U. S. Department of Energy, Technical Information Center, under contract W7405-ENG-48.

REFERENCES

(1) Rosenberg, S., "A Knowledge Based System for Providing Intelligent Access to a Petroleum User Database," LBL-8720, University of California, Lawrence Berkeley Laboratory, January 1979.

(2) Roberts, R. B. and Goldstein, I. P., "The FRL Manual," AI Memo 409, MIT, June 1977.

(3) Minsky, M., "A Framework for Representing Knowledge," in P. H. Winston (ed.) The Phychology of Computer Vision, McGraw-Hill, NY, 1975.

(4) Doyle, Jon, "Truth Maintenance Systems for Problem Solving," AI TR-419, January 1978.

(5) Rosenberg, S., Reasoning in Incomplete Domains. Proceedings of the Sixth International Joint Conference on Artificial Intelligence, 1979.

(6) Newell, A. and Simon, H. A. "Human Problem Solving," Prentice-Hall, NJ, 1972.