PROCEEDINGS

of the

SEVENTH INTERNATIONAL JOINT CONFERENCE ON ARTIFICIAL INTELLIGENCE

IJCAI·81

24-28 August 1981
University of British Columbia
Vancouver, B.C., Canada

Sponsored by
The International Joint Conferences on Artificial Intelligence

Co-sponsored by
The Canadian Society for Computational Studies of Intelligence

The American Association for Artificial Intelligence

VOLUME I

USING ACTIVE CONNECTION GRAPHS
FOR
REASONING WITH RECURSIVE RULES

Donald P. McKay and Stuart C. Shapiro

Department of Computer Science
State University of New York at Buffalo
Amherst, New York 14226

Abstract

Recursive rules, such as "Your parents' ancestors are your ancestors", although very useful for theorem proving, natural language understanding, question-answering and information retrieval systems, present problems for many such systems, either causing infinite loops or requiring that arbitrarily many copies of them be made. SNIP, the SNePS Inference Package, can use recursive rules without either of these problems. A recursive rule causes a cycle to be built in an active connection graph. Each pass of data through the cycle results in another answer. Cycling stops as soon as either the desired answer is produced, no more answers can be produced or resource bounds are exceeded.

----

1. Introduction

Recursive rules, such as "Your parents' ancestors are your ancestors", occur naturally in inference systems used for theorem proving, question-answering, natural language understanding and information retrieval. Transitive relations, e.g. $\forall(x,y,z)$ [ANCESTOR(x,y) & ANCESTOR(y,z) -> ANCESTOR(x,z)], inheritance rules, e.g. $\forall(x,y,p)$[ISA(x,y) & HAS(y,p) -> HAS(x,p)], circular definitions, e.g. "a limb is a leg or arm" and "a leg is a limb", and equivalences, e.g. $\forall(x,y,z)$ [RECEIVE(x,y,z) <=> PTRANS(z,y,z,x)], are all occurrences of recursive rules. Yet, recursive rules present problems for system implementors. Inference systems which use a "naive chaining" algorithm can go into an infinite loop, like a left-to-right top-down parser given a left recursive grammar [6, 11, 15, 23]. Some systems fail to use a recursive rule more than once, i.e. are incomplete [8, 20]. Other systems build tree like data structures containing branches the length of which depend on the number of times the recursive rule is to be applied [4, 21]. Since some of these build the structure before using it, the correct length of these branches is problematic. Some systems eliminate recursive rules by deriving and adding to the data base all implications of the recursive rules in a special pass before normal inference is done [16]. Another

measure taken to avoid the problem of circular definitions and equivalences in one system was to use "a depth first expansion policy and to limit the total depth of the expansion" [2]. This is essentially the same solution proposed by Black [1] and by Simmons and Chester [22].

Not too surprisingly, recursive rules cause problems for many programming languages developed for artificial intelligence (AI) research. AI languages, such as MicroPLANNER [23], FUZZY [11], and PROLOG [15], have differing approaches to a basic problem: there exist well formed statements in the language which cause infinite loops in the language interpreter when some theorem, procedure or clause is used. Using the terminology from MicroPLANNER, these languages are sensitive to the order of assertion, or equivalently the order of retrieval of theorems, and to the application of a theorem as a subgoal of itself. Both interact to make a collection of theorems incomplete, i.e. statements which are logically implied by its data base of assertions and theorems are not derivable because the system gets into an infinite loop.

A typical example would be the ANCESTOR example mentioned previously. In a MicroPLANNER like syntax, such a statement can be represented as the consequent theorem:
(CONSEQUENT (ANCESTOR ?X ?Y) (Z)
    (GOAL (ANCESTOR $X ?Z))(GOAL (ANCESTOR $Z
    $Y)))
The use of this theorem for solving some goal introduces a problem. If there is no other way in which to deduce instances of ANCESTOR, either by finding assertions in the data base or through application of some other rule, or if the order of application of theorems picks this theorem first regardless of any other available theorems then the above theorem reuses itself without making any progress towards finding a solution, i.e. use of the theorem causes the interpreter to enter an infinite loop.

MicroPLANNER provides a primitive, THUNIQUE, which can be used to check whether a theorem has previously been entered with the current bindings and does solve the infinite regress problem for recursive theorems. However, the user must explicitly include the appropriate statement, so the possibility exists that the user may not in fact notice that a theorem will be used recursively. This could happen when circular definitions or equivalences are inadvertently

introduced into a collection of theorems. FUZZY suffers from a similar problem and compounds it by not providing an operator equivalent to THUNIQUE. Two points should be noted. First, apparently the developers of FUZZY did not need to represent recursive procedures (LeFaivre, personal communication) and second, THUNIQUE can be simulated in FUZZY. Pure PROLOG also does not explicitly contain a THUNIQUE primitive -- using recursive rules properly is a problem with the procedural semantics of some implementations of PROLOG, but it is not a problem of the declarative semantics. Some implementations of PROLOG include an equivalent primitive. Since a primary mode of definition is recursive definition by listing clauses, this is a potential source of problems for users of PROLOG.

SNIP [12, 18, 19] was designed to use rules stored in a fully indexed data base. When a question is asked, the system retrieves relevant rules and builds an <u>active connection graph</u> which attempts to derive the answer from the rules and other information stored in the data base. Since a semantic network is used to represent all declarative information available in the system, we differ from the basic assumption of several data base question-answering systems [5, 14, 16] by not making a distinction between "extensional" and "intensional" data bases, (i.e. non-rules and rules are stored in the same data base), nor do we distinguish "base" from "defined" relations. Specific instances of ANCESTOR may be stored as well as rules defining ANCESTOR. In addition, the inference system described here does not restrict the left hand side of rules to contain only one literal which is a derived relation [5], does not need to recognize cycles in a graph [5, 9, 14] and does not require that there be at least one exit from a cycle [14].

The active connection graph may be viewed as an AND/OR problem reduction graph in which the root node represents the original question and rules are problem reduction operators. Partly influenced by Kaplan's producer-consumer model [7], the system is designed so that if a node representing some problem is about to create a node for a subproblem and there is another node already representing that subproblem or some more general instance of it, the parent node can reuse the extant node and avoid solving the same problem again. In addition, if the extant node is a more specific instance of the proposed subproblem then the results produced by the extant node are immediately made available and the extant node cancelled. The method employed handles recursive rules with no additional mechanism and, as will be seen below, the size of the resulting graph does not depend on the number of times a recursive rule will be used.

This paper describes how SNIP handles recursive rules. Aspects of the system not relevant to this issue are abbreviated or omitted. In particular, neither the details of the match routine which retrieves formulas unifiable with a given formula [17], the representation of logical connectives and formulas in SNePS [18] nor the implementation of SNIP is fully described.

2. Predicate Connection Graphs

Others have described <u>predicate connection graphs</u> [4, 8, 10] or clause interconnectivity graphs [21] which have been used in resolution theorem proving systems and question answering systems.

Our view of predicate connection graphs depends on a match algorithm which retrieves <u>all</u> formulas unifiable with a given formula. The details of the algorithm appear elsewhere [17]. Here, we relate it to the unification algorithm [3]. A match function is given as input a formula S, called the source, which it uses to find all formulas unifiable with S. The result of the match is a list of triples, $\langle T, \underline{t}, \underline{s} \rangle$, where T is a retrieved formula called the <u>target</u>, and $\underline{t}$ and $\underline{s}$ are substitutions called the <u>target binding</u> and <u>source binding</u> respectively. Essentially $\underline{t}$ and $\underline{s}$ are factored versions of the most general unifier (mgu) of S and T, and might have been computed (but are really not) in the following way. [Note: This description is accurate for first order predicate calculus without function symbols. Inclusion of function symbols adds a complication which is not treated in this paper, but is in a forthcoming paper.] Let $\underline{R}_S$ be the substitution $\{t_1/v_1, \ldots, t_n/v_n\}$ where $v_1, \ldots, v_n$ are all the variables in S, and if v_i occurs in T, t_i is a variable used nowhere else, but if v_i does not occur in T, t_i is v_i (the pair v_i/v_i is not normally allowed in substitutions since it is superfluous, but it will make our algorithms easier to describe). Let $\underline{R}_T$ be the substitution $\{v_1/v_1, \ldots, v_k/v_k\}$, where $v_1, \ldots, v_k$ are all the variables in T. Note that $T\underline{R}_T=T$ and that $T\underline{R}_T$ and $S\underline{R}_S$ have no variables in common. Now let Θ be the mgu of $T\underline{R}_T$ and $S\underline{R}_S$ such that for each pair v_1/v_2 in Θ where v_1 is a variable, v_1 occurs in T. Finally, $\underline{s}=\underline{R}_S\backslash\Theta$ and $\underline{t}=\underline{R}_T\backslash\Theta$, where $\underline{a}\backslash\underline{b}$ denotes the <u>application</u> of $\underline{b}$ to $\underline{a}$ -- the substitution derived from $\underline{a}$ by replacing each term t in $\underline{a}$ by $t\underline{b}$. For example, if $S=P(x,a,y)$ and $T=P(b,y,x)$, where a and b are constants and x and y are variables, then $\underline{R}_S=\{u/x, v/y\}$, $\underline{R}_T=\{x/x, y/y\}$, $\Theta=\{b/u, a/y, x/v\}$, $\underline{s}=\{b/x, x/y\}$ and $\underline{t}=\{x/x, a/y\}$. Note that $S\underline{s}=T\underline{t}=P(b,a,x)$, the variables in the variable position of the substitution pairs of $\underline{s}$ are all and only the variables in S, the variables of $\underline{t}$ are all and only the variables in T, all terms in $\underline{s}$ came from T, and the non-variables in $\underline{t}$ came from S. It is important to note that factoring the mgu in this way looses no information.

A predicate connection graph (pcg) is a collection of statements in predicate calculus with unifiable literals linked together by <u>edges</u> labelled with the most general unifying substitution (mgu) of the literals. In systems which use pcg's, the inference algorithms may impose constraints on which literals may have an edge between them. For example, systems which use resolution as the only rule of inference [4, 10, 21] require that the predicate calculus statements be represented in clause form and that only complementary literals be joined by an edge, i.e.

369

a literal L is linked to a literal ~L. In a system which does not represent statements in clause form, e.g. [8], and which uses the standard connectives of predicate calculus, the edges usually link an instance of a literal in the antecedent of some statement with a unifiable instance of the same literal in the consequent of some other statement. In such a system which uses both backward chaining and forward chaining, an edge between $P(x)$ and $P(y)$ asserts that to show $P(x)$ use the statement in which $P(y)$ appears and that if $P(y)$ for some y is ever deduced then the result can be used to further satisfy the statement in which $P(x)$ appears.

The match operation specifies a pcg of a slightly different form. Instead of labelling the edge with the mgu, a directed edge linking a source node (S) to a target node (T) labelled with the target binding ($\underline{t}$) and the source binding ($\underline{s}$) is used. Figure 1 shows a pcg consisting of five rules, labelled R1 through R5. While the rules considered in the remainder of this paper are of the form $A_1 \& \dots \& A_j \rightarrow C$ where $j \geq 1$ and all variables are universally quantifed (i.e. Horn clauses), SNIP is not so limited (see [19]). The edges are labelled with the pair ($\underline{t},\underline{s}$) where $\underline{t}$ is the target binding and $\underline{s}$ is the source binding. The source node is the literal at the tail of the edge and the target node is at the head of the edge. For example, the edge labelled "a" in Figure 1 has $P(a,y,x)$ as the source literal and $P(x,y,z)$ as the target literal. SNIP does not explicitly store the pcg but uses the match function described above to compute the edges on demand. The remainder of this paper is concerned with pcg's of this last form.

3. Active Connection Graphs

An active connection graph (acg) is a connection graph in which edges link literals and are labelled with a target binding and source binding. These graphs are active because instances of literals flow from one formula to another formula via the edges. Using the producer-consumer
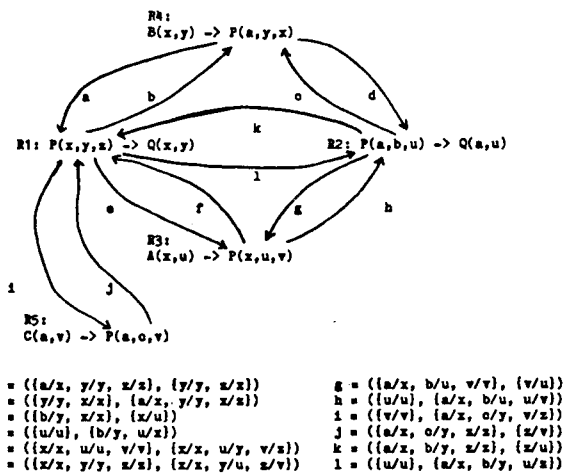
analogy, a rule instance can be considered a producer of instances of its consequents and a consumer of instances of its antecedents. Furthermore, the acg contains instances of the rules in the pcg, the edges in the acg point in the opposite direction to the corresponding edges in the pcg and the bindings play an active role (see below). The target binding filters the flow of instances of literals. The source binding translates between variable contexts.

Suppose a consequent reasoning system has been asked to deduce all instances of Q using the pcg of Figure 1. It can use R1 and R2 to deduce instances of Q if appropriate instances of P can be deduced. Thus, rules R3-R5 can be used. A full acg for this scenario is presented in Figure 2. Rectangles enclose formulas. The partitions contain literals instances. Antecedents appear on the left of the double line, consequents to the right. The rectangle at the top of Figure 2 represents the request to deduce all instances of Q and as such has an empty consequent part. In this example, each acg rule labelled A_i is an instance of the pcg rule labelled R_i. These labels are arbitrary and the reader should not infer anything about the construction of the acg based on the labels alone. The remainder of this section expands this simple notion of active connection graphs using the acg of Figure 2 as an example.

Target bindings and source bindings operate on the bindings flowing through the acg. A target binding is a filter which only lets through those bindings which have the binding of the filter as a subset. For example, in the active connection graph of Figure 2, if G2 produced the bindings $\{a/x, b/y, c/z\}$ and $\{a/x, d/y, e/z\}$, only $\{a/x, b/y, c/z\}$ would be allowed to pass through the filter $\langle a/x, b/y \rangle$ to A2. The source binding is used to switch variable contexts. A2 contains only the variable u while G2 produces bindings with the

Figure 1
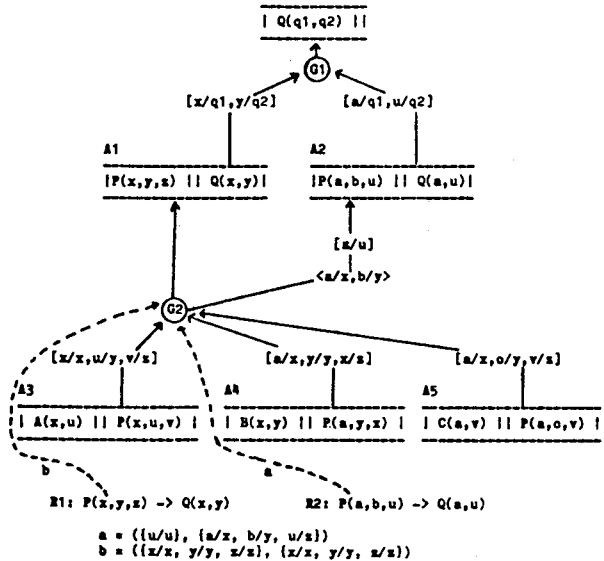PCG using target bindings and source bindings.

R4:
$B(x,y) \rightarrow P(a,y,x)$

R1: $P(x,y,z) \rightarrow Q(x,y)$ R2: $P(a,b,u) \rightarrow Q(a,u)$

R3:
$A(x,u) \rightarrow P(x,u,v)$

R5:
$C(a,v) \rightarrow P(a,o,v)$

a = (($a/x, y/y, z/z$), ($y/y, z/z$))
b = (($y/y, z/z$), ($a/x, y/y, z/z$))
c = (($b/y, z/x$), (x/u))
d = ((u/u), ($b/y, u/z$))
e = (($z/x, u/u, v/v$), ($z/x, u/y, v/z$))
f = (($z/x, y/y, z/z$), ($z/x, y/u, z/z$))
g = (($a/x, b/u, v/v$), (v/u))
h = ((u/u), ($a/x, b/u, u/v$))
i = ((v/v), ($a/x, o/y, v/z$))
j = (($a/x, o/y, z/z$), (z/v))
k = (($a/x, b/y, z/z$), (z/u))
l = ((u/u), ($a/x, b/y, u/z$))

Figure 2
Full ACG to deduce instances of Q.

$Q(q1,q2)$

G1

$[x/q1,y/q2]$ $[a/q1,u/q2]$

A1   A2
$|P(x,y,z)\ ||\ Q(x,y)|$ $|P(a,b,u)\ ||\ Q(a,u)|$

$[z/u]$

$\langle a/x,b/y \rangle$

G2

$[z/x,u/y,v/z]$ $[a/x,y/y,z/z]$ $[a/x,o/y,v/z]$

A3   A4   A5
$|\ A(x,u)\ ||\ P(x,u,v)\ |$ $|\ B(x,y)\ ||\ P(a,y,x)\ |$ $|\ C(a,v)\ ||\ P(a,o,v)\ |$

b   a

R1: $P(x,y,z) \rightarrow Q(x,y)$ R2: $P(a,b,u) \rightarrow Q(a,u)$

a = ((u/u), ($a/x, b/y, u/z$))
b = (($z/x, y/y, z/z$), ($z/x, y/y, z/z$))

370

variables x, y and z, so in order for G2 to send bindings to A2, the variable z must be mapped to the variable u. Continuing with the example above, since the binding {a/x, b/y, c/z} passed through the filter it next encounters the switch [z/u]. The switch uses the binding application operation defined in Section 2 to generate a new binding. In this case, {z/u} \ {a/x, b/y, c/z} yields {c/u}, which is an appropriate binding in the context of A2.

To deduce all instances of Q requires back chaining through the pcg until only ground literals are found or all possible rules are tried. Initially, a request is created which contains a literal. In Figure 2, Q(q1,q2) is that request. The next step is to create a <u>goal node</u> for the literal. The goal node matches its literal with the pcg to find all literals which unify with it. If there are ground instances then the source bindings of those matches are answers and the goal node produces them immediately. Other matches can be literals which are antecedents or consequents of rules in the pcg. For every literal in the consequent of some rule, a new rule instance is added to the acg using the target bindings, and the instance is connected through a switch containing the source binding to the goal node. The same process of creating goal nodes is applied to each of the antecedents of the new rule instances created in the previous step. This process is repeated until either no more rules apply or only ground instances are found. However, a new goal node need not be created if an existing one will suffice. The remainder of this section describes how to find adequate extant goal nodes without doing extra matching and how to use results previously generated by a goal node.

Suppose a goal node is about to be created for some literal in a rule instance of an acg. Which other literals are likely to have goal nodes which should be checked? Namely, those other literals in the pcg unifiable with the literal which are in the antecedent of some rule of the pcg and which have already had goal nodes created for them. Thus, when a new goal node performs its match as described above, the matched antecedent literals are marked with a pointer to the newly created goal node. These pointers, which link antecedent pcg literals to goal nodes in the acg, are called <u>goal pointers</u>. Just like pcg edges, each goal pointer is marked with a target binding and source binding. When a new goal node is about to be created, if its literal has no goal pointer then no existing goal node will be useful. If it does, it is possible that one of the goal nodes pointed to could be used instead of the proposed goal node.

Let's consider as an example the process of building the acg of Figure 2 using the pcg of Figure 1. Recall that the request was for all instances of Q. A goal node is created for Q(q1,q2) and a match performed. The results of the match of Q(q1,q2) are the tuples <Q(a,u), {u/u}, {a/q1, u/q2}> and <Q(x,y), {x/x,y/y}, {x/q1, y/q2}>. Both Q(a,u) and Q(x,y) are consequents of rules, so rule instances are created for them in the acg connected through the appropriate switches

to G1. There are no goal pointers created because Q(q1,q2) does not match any literal in an antecedent of a rule. Continuing to expand the acg using the antecedents of A1 and A2 requires picking one of them to expand first. Whichever order is picked, the same acg is constructed.

Suppose A1 is picked for expansion. First, the literal P(x,y,z) is checked for a goal pointer. There are none, so a new goal node is built for P(x,y,z). The result of the match are the tuples <P(a,b,u), {u/u}, {a/x, b/y, u/z}>, <P(x,u,v), {x/x, u/u, v/v}, {x/x, u/y, v/z}>, <P(a,y,x), {y/y, x/x}, {a/x, y/y, x/z}>, <P(a,c,v), {v/v}, {a/x, c/y, v/z}> and <P(x,y,z), {x/x, y/y, z/z}, {x/x, y/y, z/z}>. P(a,b,u) is an antecedent of R2, so P(a,b,u) gets a goal pointer to the goal node for P(x,y,z). Also, P(x,y,z) is given a goal pointer to its own goal node. We draw a goal pointer as a dashed line labelled with the target binding and source binding of the match. Figure 3 shows the acg and part of the pcg with goal pointers after this step. The remaining target literals are each consequents of some rule and the rules are added to the acg as before. Next, a goal node is to be created for P(a,b,u) of A2. But the pcg literal P(a,b,u) in R2 has a goal pointer. Thus, some instance of the old goal literal, P(x,y,z), unifies with the literal which has the goal pointer but the binding of the current acg rule instance is not necessarily compatible with the binding of the old goal literal.

In the current example, the old target binding from the goal pointer is identical to the binding associated with A2. However, G2 is not identical to the proposed goal for P(a,b,u). Rule instance A2 is interested in a subset of all instances of P(x,y,z), namely those instances which have a/x and b/y. It is important to note that G2 will produce all that the proposed goal node for P(a,b,u) would produce and more. Also, above G2 all instances are in terms of the variables of R1. Instead of creating a new goal node for P(a,b,u) and the edges associated with it, G2 is reused. The results from G2 must be filtered by <a/x, b/y> and variable contexts switched by [z/u]. The filter is computed from the application of the current binding to the old source binding. This assures that the filter


Figure 3
Active connection graph with goal pointers.

contains variables which the old goal node produces. The switch is obtained by considering the binding pairs of the old source binding for which the term is a variable as a "variable map". Discarding those binding pairs for which the term is a constant and inverting the remaining pairs yields the appropriate binding for the switch. For example, discarding pairs with constant terms and then inverting {a/x, b/y, u/z} yields {z/u}. Figure 2 shows the acg after the filter and switch have been built.

If A2 were expanded before A1 then a goal node would have been created for it and goal pointers established from the antecedents of R1 and R2 to the goal node. To obtain the same acg requires that the goal node created for P(a,b,u) in A2 be superseded by the goal node eventually created for P(x,y,z). When a goal node is superseded, it is erased from the acg and all the consumers of the superseded goal nodes become consumers of the new, more general goal node with an appropriate filter and switch between the superseding node and the old consumers.

We stated above that rule instances in the acg consume instances of antecedents and produce instances of consequents but have totally ignored what constitutes such an instance. There are basically two alternatives. An instance of a literal can be either a separate literal or a reference to a literal and a binding which when applied to the literal would yield an instance of it. We prefer the second alternative because the inference algorithms need not produce extra literals and because the match operation mentioned above returns such information. Furthermore, information about literals is superfluous in communication between rule instances in the acg because the literals are known to unify (that's why there is an edge in the pcg) and therefore, only bindings need be communicated along the edges of the active connection graph.

Finally, we mentioned above that the acg contains rule instances but the previous examples use the pcg rules directly. In general, rule instances in the acg are pairs -- a rule from the pcg and an associated binding. The rule instances in the examples all have an identity binding, i.e. a binding in which each variable of a binding pair is bound to itself. The binding is used to restrict the rule instances to generating just those literal instances which are requested. In other words, it is used as an internal filter in the rule instances in order to keep the inference more focussed.

In summary, an acg contains instances of rules which are linked to each other through filters, switches and goal nodes. The separation of the target and source bindings allow rule instances to work in their own variable environment, relying on the source binding of the pcg to enable switching of variable contexts and on the target bindings to allow more general producers to be used by less general consumers. The goal nodes are the production site of instances of literals and so are useful to more than one consumer. Finally, goal nodes are indexed on antecedent literals using goal pointers.

4. Recursive rules cause cycles

A set of recursive rules is of the form A_1 &...& A_j -> B_1, B_1 &...& B_k -> ... -> C, with C unifiable with at least one of the antecedents, A_1 say. In an acg, this means that the goal node for C can be used instead of creating a new goal node for A_1. As an example, consider the pcg of Figure 4 which contains the recursive ANCESTOR rule. The two rules represent the predicate calculus statements ∀(x,y,z)[ANCESTOR(x,y) & ANCESTOR(y,z) -> ANCESTOR(x,z)] and ∀(x,y)[PARENT(x,y) -> ANCESTOR(x,y)]. The remainder of the entries represent the ground literals PARENT(Bill, John), PARENT(John, Mary), ANCESTOR(Bill, Bob), ANCESTOR(John, Mary) and ANCESTOR(Mary, Sarah). Since the ground literals are never a source node in a match, the edges from the ground literals to literals in rules are omitted.

Consider a request, ANCESTOR(q1,q2), for all instances of the ANCESTOR relation. Inference proceeds by creating a goal node for the request, creating rule instances for all rules in the pcg which have the literal as a consequent and creating a goal pointer for each literal in the antecedent of some rule. Figure 5 shows the acg at this point. Here, instances which are produced are shown in terms of the variables of the goal literal

a = ({x/x, z/z}, {x/x, z/y})
b = ({x/x, y/y}, {x/x, y/z})
c = ({x/x, z/z}, {x/y, z/z})
d = ({y/y, z/z}, {y/x, z/z})
e = ({x/x, y/y}, {x/z, y/y})
f = ({x/x, y/y}, {x/x, y/y})
g = ({y/y, z/z}, {y/z, z/y})
h = ({z/z, y/y}, {x/y, y/z})

i = ({}, {Bill/x, Bob/y})
j = ({}, {John/x, Mary/y})
k = ({}, {Mary/x, Sarah/y})
l = ({}, {Mary/y, Sarah/z})
m = ({}, {John/y, Mary/z})
n = ({}, {Bill/y, Bob/z})
o = ({}, {Bill/x, John/y})
p = ({}, {John/x, Mary/y})

Figure 4
Sample data base for Section 4.

next to the goal nodes which collect them. When such an instance is produced, it flows through the acg until either it encounters the top level request or it is consumed by a goal node which has previously produced it. Thus, the bindings {Bill/q1, Bob/q2}, {John/q1, Mary/q2} and {Mary/q1, Sarah/q2}, representing the ground literals ANCESTOR(Bill, Bob), ANCESTOR(John, Mary) and ANCESTOR(Mary, Sarah) respectively, are stored by G1 and consumed by the top level request. Picking one rule instance, A1 or A2, to expand first has no significant impact because if A2 were expanded first then any results derived from A2 would be stored in the goal node just as these are. The order of the creation of the goal nodes in A1 is insignificant for this example because both reuse G1.

Consider ANCESTOR(x,y) as the first goal literal. ANCESTOR(x,y) has the goal pointer ({x/x, y/y}, {x/q1, y/q2}), and the binding associated with it in A1 is {x/x, y/y}. Thus, the goal pointer's old target binding is identical to the rule instance binding and G1 can be used. This is accomplished by giving G1 another consumer, the ANCESTOR(x,y) antecedent, connected through the switch [q1/x, q2/y]. The bindings from G1 are produced for this new consumer but further processing by A1 is not possible because the other antecedent of A1 has not yet consumed any instances of ANCESTOR(y,z).

Goal nodes act as data collectors [13, 19]. A data collector stores all literals it has consumed and never produces a literal it has previously produced. When an old goal node is given an additional consumer, all literals previously produced are immediately available to the new consumer. Also, the new consumer receives any new literals produced by the goal node. The fact that data collectors never produce the same literal twice protects SNIP from getting into an infinite loop by prohibiting the passing of the same literal around a cycle in the acg.

Next, consider trying to create a goal node for ANCESTOR(y,z). Again, G1 is found using the goal pointer and as before the old target binding and the binding of the rule instance are identical.

A new consumer is added to G1 through a switch of [y/q1, z/q2] and the three instances stored by G1 are produced for this new consumer. Since A1 has now received the bindings {John/x, Mary/y} for one antecedent and and {Mary/y, Sarah/z} for the other, it can produce {John/x, Sarah/z} and pass it to G1. Again, G1 produces {John/q1, Sarah/q2} to all consumers because it has not produced it before. Now, A1 cannot produce any further instances.

Finally, a goal node is created for PARENT(x,y), call it G2. Since there are no goal pointers for PARENT(x,y), a match is performed resulting in the ground instances of {Bill/x, John/y} and {John/x, Mary/y}. Figure 6 shows the acg after G2 has matched these instances. These are stored by G2 and produced to all consumers. A2 now has its antecedent satisfied in the bindings above. Since these instances are next passed to G1 and {Bill/q1, John/q2} has not been produced before, it is produced to the top level request and the other consumers. After passing through the switches, the appropriate binding arrives in A1 which produces ANCESTOR(Bill, Mary). Again, G1 has not produced {Bill/q1, Mary/q2} previously, so it passes it on to all consumers. A1 in turn produces {Bill/x, Sarah/z} to G1. Since G1 has not previously produced {Bill/q1, Sarah/q2}, it produces this binding to all consumers. Now, no further results can be produced by A1 and since there are no other rules left to expand, inference terminates.

This example demonstrates that recursive rules can be productively used in an acg and not cause an infinite loop. The reason no infinite loop is encountered is that no further results could be produced and that the acg contains a cycle as opposed to continually trying to use the recursive ANCESTOR rule. This example also demonstrates the accessing of data stored by a goal node previous to adding a new consumer. This is a property of the goal node acting as a data collector. Finally, this example suggests how rule instances produce bindings, waiting until sufficient instances have been consumed in the appropriate binding. SNIP
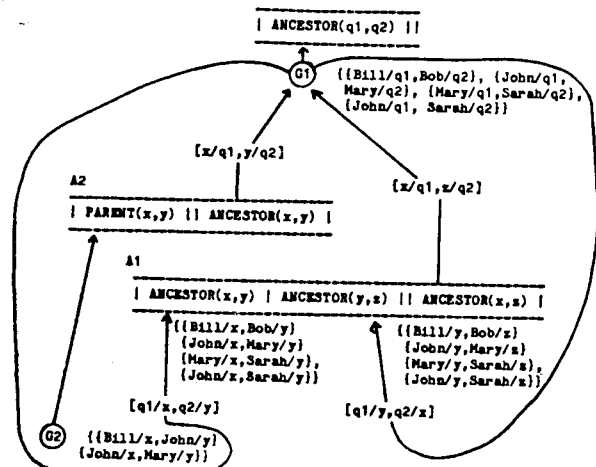

Figure 5
Partial acg for request of ANCESTOR(q1,q2).

Figure 6
ACG after creation of goal node for PARENT(x,y).

allows other logical connectives so that other schemes for rule instances to produce instances of their consequents are required (see [18]). Note, this kind of structure is built for indirectly recursive rules as well as the directly recursive rule in the above example, the resulting active connection graph is a directed graph and neither a tree nor a directed acyclic graph, and once it is built its size is constant.

5. Summary

In SNIP, recursive rules cause cycles to be built in an active connection graph. The key features of active connection graphs which allow recursive rules to be handled are: 1) goal nodes are data collectors; 2) data collectors never produce the same answer more than once; 3) the data collector may report to more than one consumer; 4) a new consumer may be assigned to a data collector at any time -- it will immediately be given all previously collected data; 5) variable contexts are localized, switches change contexts dynamically as data flows around the graph; 6) filters allow more general producers to be used by less general consumers; 7) goal nodes are indexed on antecedent literals which were matched by the goal literal.

References

1. Black, F. A deductive question-answering system. In Semantic Information Processing, Minsky, M. (ed.), MIT Press, Cambridge, 1968.
2. Bobrow, D.G., Winograd, T. et al. Experience with KRL-0 one cycle of a knowledge representation language. Proc. IJCAI-77, 1977, 213-222.
3. Chang, C.-L. and Lee, R.C.-T. Symbolic Logic and Mechanical Theorem Proving, Academic Press, New York, 1973.
4. Chang, C.-L. and Slagle, J.R. Using rewriting rules for connection graphs to prove theorems. Artificial Intelligence, Vol. 12(2), August, 1979, 159-180.
5. Chang, C.-L. On evaluation of queries containing derived relations in a relational data base. In Formal Bases for Data Bases, Gallaire, H., Minker, J. and Nicolas, J. (eds.), Plenum, New York, 1980.
6. Fikes, R.E. and Hendrix G.G. The deduction component. In Understanding Spoken Language, Walker, D. (ed.), Elsevier North-Holland, 1978, 355-374.
7. Kaplan, R.M. A multi-processing approach to natural language understanding. Proc. NCC, AFIPS Press, Montvale, NJ, 1973, 435-440.
8. Klahr, P. Planning techniques for rule selection in deductive question-answering. In Pattern Directed Inference Systems, Waterman, D.A. and Hayes-Roth, R. (eds.), Academic Press, 1978, 223-239.
9. Kellogg, C. and Travis, L. Reasoning with data in a deductively augmented data management system. To appear in Advances in Data Base Theory - Volume 1, Gallaire and Minker (eds.).
10. Kowalski, R. A proof procedure using connection graphs. JACM, Vol 22(4), October, 1975, 572-595.
11. LeFaivre, R. FUZZY Reference Manual. Computer Science Department, Rutgers University, 1977.
12. Martins, J.P., McKay, D.P. and Shapiro, S.C. Bi-directional inference. Department of Computer Science, Technical Report 174, SUNY/Buffalo, March, 1981.
13. McKay, D.P. and Shapiro, S.C. MULTI- A LISP based multiprocessing system. Proc. 1980 LISP Conference, Stanford University, 1980.
14. Naqvi, S.A. and Henschen, L.J. Performing inferences over recursive data bases. Proc. NCAI, Stanford University, 1980.
15. Pereira, L.M., Coelho, H. and Pereira, F. User's guide to DECsystem 10 PROLOG (Provisional Version). Divisao de Informatica, Lab. Nac. de Engenharia Civil, Lisbon, Portugal, 1978.
16. Reiter, R. On structuring a first order data base. Proc. Second National Conference, Canadian Society for Computational Studies of Intelligence, 1978, 90-99.
17. Shapiro, S.C. Representing and locating deduction rules in a semantic network. SIGART Newsletter, No 63, June, 1977, 14-18.
18. Shapiro, S.C. The SNePS semantic network processing system. In Associative Networks, Findler, N.V. (ed.), Academic Press, 1979, 179-203.
19. Shapiro, S.C. and McKay, D.P. Inference with recursive rules. Proc. NCAI, Stanford University, 1980.
20. Shortliffe, E.H. Computer Based Medical Consultations: MYCIN. American Elsevier, New York, 1976.
21. Sickel, S. A search technique for clause interconnectivity graphs. IEEE Transactions on Computers Vol. C-25, 8, August, 1976, 823-835.
22. Simmons, R.F. and Chester, D. Inference in quantified semantic networks. Proc. IJCAI-77, 1977, 267-273.
23. Sussman, G.J., Winograd, T. and Charniak, E. Micro-Planner Reference Manual. AI Memo No. 203A, MIT AI Laboratory, December, 1971.