Journal of Automated Reasoning **15:** 237–265, 1995.

The Anatomy of Vampire

Implementing Bottom-up Procedures with Code Trees

ANDREI VORONKOV*
Computing Science Department, Uppsala University, Sweden

(Received: 24 August 1994; in final form 12 February 1995)

Abstract. We present an implementation technique for a class of bottom-up logic procedures. The technique is based on *code trees*. It is intended to speed up most important and costly operations, such as subsumption and resolution. As an example, we consider the forward subsumption problem which is the bottleneck of many systems implementing first-order logic.

To efficiently implement subsumption, we specialize subsumption algorithms at run time, using the *abstract subsumption machine*. The abstract subsumption machine makes subsumption-check using sequences of instructions that are similar to the WAM instructions. It gives an efficient implementation of the "clause at a time" subsumption problem. To implement subsumption on the "set at a time" basis, we combine sequences of instructions in *code trees*.

We show that this technique yields a new way of indexing clauses. Some experimental results are given.

The code trees technique may be used in various procedures, including binary resolution, hyper-resolution, UR-resolution, the inverse method, paramodulation and rewriting, OLDT-resolution, SLD-AL-resolution, bottom-up evaluation of logic programs, and disjunctive logic programs.

Key words: bottom-up algorithms, indexing, subsumption, otters, vampires.

1. Introduction

Most automatic theorem-proving systems can be generally divided in two parts by their mode of evaluation (or clause generation). *Top-down* systems start from a goal, reducing it to subgoals until all subgoals become axioms. Such systems usually exploit just one branch of the search tree. *Bottom-up* systems usually deal with a database of clauses and generate new clauses by applying inference rules to clauses in the database. There are also systems that can combine top-down with bottom-up.

Systems that keep track of previously generated clauses are said to *retain information* in (Wos *et al.*, 1991). All bottom-up systems retain information. If a top-down system retains information, it is usually characterized as a combination of top-down and bottom-up.

The systems retaining information show a superior performance when dealing with problems involving large search spaces. This has been widely recognized

* Supported by Swedish TFR grant No. 93–409.

in the deductive database community. There are two main reasons why systems that retain information become more efficient on hard problems:

1. Systems that retain no information must often do the same job on different branches of the search tree.
2. In systems that retain information one can implement major operations on the "set at a time" basis, as opposed to "tuple a time" or "clause at a time" procedures used by, e.g., Prolog.

Procedures retaining information have been traditionally used in automated deduction, since its very beginning. Recently, the need for such procedures has been recognized in logic programming and deductive databases. In logic programming, it has been noted that tabulation can both speed evaluation and give more declarative treatment of some subgoals. In deductive databases, where the databases can be very large, the old "tuple at a time" methods are not adequate any more. Also, deductive databases do not allow the query answering process to be controlled by a user. In this case recursive specifications can cause nontermination, unless (some) intermediate goals are retained. In logic programming and deductive databases the systems that make use of previously generated clauses (or goals) are usually considered as bottom-up procedures. Throughout this article we shall use "bottom-up" to denote procedures that retain clauses.*

There are also systems that combine bottom-up and top-down evaluation. For instance, various modifications of magic sets transformations require the semi-naive bottom-up evaluation of the transformed logic program (Bancilhon *et al.*, 1986; Beeri and Ramakrishnan, 1987). SLD-AL resolution (Vieille, 1989) and OLDT-resolution (Tamaki and Sato, 1986; Warren, 1992) are further examples of systems that combine the bottom-up and top-down evaluations. In automatic theorem proving, model elimination, which is a top-down procedure, has recently been augmented with lemmas (Astrakhan and Stickel, 1992; Goller *et al.*, 1994), which allow one to use previously proved subgoals.

This article deals with an implementation technique for bottom-up systems, which can be used to speed up many important algorithms.

Bottom-up algorithms have been used through the years in automated deduction. They are used by many automatic theorem-proving procedures, e.g., binary resolution (Robinson, 1965b), hyperresolution (Robinson, 1965a) and the inverse method (Maslov, 1967; Voronkov, 1992). In order to efficiently implement such procedures, indexing techniques have been developed (Overbeek, 1975; McCharen *et al.*, 1976a; McCharen *et al.*, 1976b; Stickel, 1989; McCune, 1988; McCune, 1992). Indexing allows main operations (like resolution and subsumption) to be implemented on the "set at a time" basis. As noted in (Wos, 1992), the implementation of discrimination trees in the well-known theorem prover Otter resulted in a much faster performance.

* The bottom-up methods of search are also sometimes called *indirect* methods, as opposed to *direct* top-down methods. (Lincoln Wallen, private communication).

Most bottom-up systems are based on the following loop that resembles the "closure algorithm" from (Lusk, 1992):

Let D be the initial database of clauses.
`loop`:
1. Apply given inference rules to D (or a part of D). Let ΔD be the new clauses.
2. If $D, \Delta D$ satisfy a termination criterion, then terminate.
3. Apply a retention test to ΔD.
4. Change D, based on ΔD.
5. Add ΔD to D.

`end of loop`.

For this class of procedures, we call clauses in D *kept* clauses, and clauses in ΔD *new* clauses.

Let us consider these steps in more detail.

1. Inference rules are usually resolution rules applied to clauses from D. In SLD-AL resolution, inference rules and clauses are more complicated. In fact, it supports two kinds of clauses: the first represents queries, and the second represents solutions to the queries.
2. There are many kinds of termination criterion. Usually, the procedure terminates when a goal is found. The termination criterion can also be based on the *saturation*: the procedure terminates when no new clauses can be generated.
3. The retention test for ΔD usually includes *forward subsumption*: remove from ΔD every clause subsumed by a clause in D. In principle, there might be other criteria for discarding new goals (McCune, 1990). When the number of generated clauses is big compared with the number of kept clauses, forward subsumption becomes the most costly operation.
4. The changes in D are induced by new clauses from ΔD. Two algorithms may be applied at this stage. *Back subsumption* consists of removal from D every clause subsumed by a clause from ΔD. If a new clause includes equalities, it can be used for *back demodulation*, i.e., rewriting applied to clauses from D.

This loop is used in most implementations of binary resolution and hyper-resolution (Robinson, 1965b; Robinson, 1965a), the inverse method (Maslov, 1967; Voronkov, 1992), implementations based on magic sets (Bancilhon *et al.*, 1986; Beeri *et al.*, 1987), tabled computations of logic programs (Tamaki and Sato, 1986; Warren, 1992; Sudarshan and Ramakrishnan, 1993), various deductive databases procedures, for example, query answering and integrity checking (Lloyd and Topor, 1985; Lloyd *et al.*, 1986; Decker, 1986; Das, 1992), bottom-up static analysis of logic programs (Codish and Demoen, 1993), constraint logic programming, parsing (Tambe and Rosenbloom, 1994), and production systems

(Maeda *et al.*, 1994). Top-down procedures that use lemmaizing, like SLD-AL resolution (Vieille, 1989) or some implementations of model elimination (Astrakhan and Stickel, 1992), also use a similar scheme.

Certain specific features of such systems pose a challenge to people who implement them:

1. There is usually a small number of operations that must repeatedly be applied to clauses in the database (for example, resolution).
2. The database may be very large and dynamically changing.
3. The number of calls of these operations is so big that it is practically impossible to implement them on the "clause at a time" basis.

Below we shall consider particular examples that illustrate the number of operation calls for some hard problems in automated deduction.

The main problems faced by bottom-up procedures are based on the necessity to efficiently handle large dynamically changing databases of clauses. Implementations of Prolog meet very different problems. Goals in Prolog are executed on the "goal at a time" basis, which requires very fast processing of the current goal. WAM, which has become a standard for implementing Prolog, is hence designed with the aim of the efficient execution of one (current) subgoal or a chain of subgoals at a time. Such "goal at a time" processing is inadequate for bottom-up algorithms which must often handle millions of clauses.*

In this paper we introduce an implementation technique called *code trees* which is intended for speeding these basic procedures. We shall show its use on the forward subsumption problem. The idea of code trees is very general and may as well be applied to various concepts of resolution. It can also be used for the implementation of equality reasoning procedures, like paramodulation (Robinson and Wos, 1969) and rewriting.

Our technique is based on the following two main ideas. The first idea is to compile specialized subsumption procedures for kept clauses. It can be considered as a run time specialization of a general subsumption algorithm. This technique has much in common with the technique of WAM-based Prolog implementations. We call it *the abstract subsumption machine.* It gives a very efficient subsumption algorithm for the problem of subsuming a clause by a clause. A specific feature of the use of this approach is that the compilation must be performed at run time.

When the set D of kept clauses is large, the implementation of subsumption on the "clause at a time" basis becomes inefficient. The second idea of this paper is to perform subsumption on the "set at a time" basis via code trees that merge several specialized algorithms into one.

* In our experiments we tried benchmarks where around 900,000,000 clauses were generated, and the number of subsumption-checks was about 20,000,000,000.

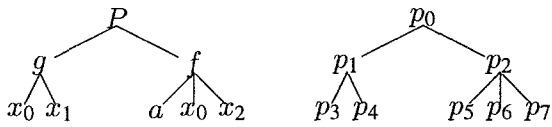

Fig. 1. The clause $P(g(x_0, x_1), f(a, x_0, x_2))$ and its set of positions $\{p_0, \ldots, p_7\}$.

2. Subsumption: The Abstract Subsumption Machine

We assume acquaintance with the basic notions of terms and substitutions. A *clause* is a list of terms. Terms in a clause are also called *literals*. *Variables* will be denoted by the lower case letters x, y, z, u, v, w, maybe with indices. Let c_1, c_2 be two clauses. Clause c_1 *subsumes* clause c_2 iff there is a substitution θ such that $c_1\theta$ is a subset of c_2*. For example, let $c_1 = \{P(x, f(y)), P(z, f(f(z)))\}$ and $c_2 = \{P(y, f(f(y)))\}$. Then c_1 subsumes c_2 with the substitution $[x/y, y/f(y), z/y]$ and c_2 subsumes c_1 with the substitution $[y/z]$.

A clause is a *unit clause* iff it consists of exactly one literal. Clauses with more than one literal will be called *multi-literal* clauses. In most logic programming applications, the bottom-up procedures use only unit clauses. Multi-literal clauses are used in disjunctive logic programs, theorem proving procedures and static analysis of logic programs. **In this section, we only consider unit clauses.** Multi-literal clauses will be considered in Section 4.

The algorithms considered in this paper deal with *positions* in clauses. A position in a clause c corresponds to an occurrence of a subterm in c. A tree representation of the (unit) clause $P(g(x_0, x_1), f(a, x_0, x_2))$ and the corresponding set of positions are given in Figure 1. This clause will be our running example until the end of Section 3.

In order to present our algorithms, we need to introduce a few operations on positions. Let p be a position, corresponding to t_i in a subterm $f(t_1, \ldots, t_n)$. Then $right(p)$ is the position to the right of p, i.e. the one that corresponds to t_{i+1}. If such position does not exist, $right(p)$ is undefined. For example, on Figure 1 we have $right(p_1) = p_2$, $right(p_5) = p_6$ and $right(p_0)$ is undefined.

If the position p corresponds to a term t then $down(p)$ denotes the position of the first argument of t. If t has no arguments (i.e. t is a variable or a constant) then $down(p)$ is undefined. For example, on Figure 1 we have $down(p_0) = p_1$, $down(p_1) = p_3$ and $down(p_6)$ is undefined.

$var(p)$ is true iff the term at the position p is a variable. $functor(p)$ denotes the principal functor of the term in the position p. For simplicity, we assume that the arity of each function symbol is fixed. For example, terms $g(x)$ and

* In other words every literal in $c_1\theta$ is a literal in c_2.

$g(x, y)$ cannot be used together.* *defined*(p) is true iff p is a valid position in the clause (operations *down*, *right* may lead to invalid positions). For p, q positions, $p == q$ is true iff terms at positions p, q are equal. $\neq\neq$ is the negation of $==$. For example, on Figure 1 $var(p_0)$ is false, $var(p_7)$ is true, $functor(p_0) = P$, $p_3 == p_6$ is true and $p_3 == p_1$ is false.

For unit clauses c_1, c_2, the subsumption problem reduces to the *matching problem*: find a substitution θ such that $c_1\theta = c_2$. An example of the matching algorithm *match* is given in Figure 2. We do not claim this algorithm to be the most efficient. One can often do better, depending on the nature of the problem and the concrete representation of clauses.

```
Algorithm match(c1, c2)
  Let p, q be positions
  Let subst, post be initially empty stacks of pairs of positions
begin
  push(⟨c1, c2⟩, post);
  while nonempty(post) {
    ⟨p, q⟩ := pop(post);
    if var(p)
    then push(⟨p, q⟩, subst)
    else if var(q) or functor(p) ≠ functor(q)
    then exit with fail
    else {
      if defined(right(p)) then push(⟨right(p), right(q)⟩, post);
      if defined(down(p)) then push(⟨down(p), down(q)⟩, post)
      }
    }
  forall cliques c = ⟨p0, q0⟩, ..., ⟨pn, qn⟩ ∈ subst
  forall i ∈ {1, ..., n} {
    if q0 == qi
    then remove ⟨pi, qi⟩ from subst
    else exit with fail
    }
  exit with success
end
```

Fig. 2. The general matching algorithm for clauses c_1, c_2.

* In logic programming language such use of functors is allowed. In this case we have to consider pairs g/n consisting of a function symbol and its arity. We do not treat this case, for the sake of simplicity.

In this algorithm the stack *subst* represents the matching substitution. The stack *post* is used to keep the pairs of positions whose matching is postponed. p, q are used to denote the positions which are currently under consideration.

For a stack S and an element e, $push(e, S)$ denotes the operation of adding e at the beginning of S. To add e to the end of the stack S, we write $add(e, S)$. $pop(S)$ removes the first element from the stack S and returns this element as the result. $nonempty(S)$ is true if the stack S is not empty.

Let $S = \langle p_0, q_0 \rangle, \ldots, \langle p_n, q_n \rangle$, where p_i are positions. Then $tail(S)$ is $\langle p_1, q_1 \rangle$, $\ldots, \langle p_n, q_n \rangle$. A *clique* in S is any maximal subsequence $\langle p_{i_0}, q_{i_1} \rangle, \ldots, \langle p_{i_k}, q_{i_k} \rangle$ of S such that $k \geqslant 1$ and $p_{i_j} == p_{i_0}$ for all $j \in \{0, \ldots, k\}$.

The algorithm *match* tries to construct a substitution θ such that $c_1\theta = c_2$. First of all, *match* traverses the c_1 and c_2, putting into *subst* pairs $\langle v, t \rangle$, where v is a variable in c_1, t a term in c_2. Then it compares terms in c_2 corresponding to different occurrences of the same variable in c_1. There were two reasons for us to delay the comparison. The first reason is that the term comparison is potentially very costly operation. The second reason is that this form of the algorithm makes more *code sharing* considered below in Section 3.

THEOREM 1. *The algorithm match is sound and complete. More precisely:*

1. *If c_1 does not subsume c_2, then $match(c_1, c_2)$ terminates with failure.*
2. *If c_1 subsumes c_2, then $match(c_1, c_2)$ terminates with success. In this case after termination subst contains the matching substitution.*

Proof. Omitted.

The set of clauses in bottom up procedures is dynamically changing. Thus, it is not possible to pre-compile the subsumption algorithm. However, one can try to compile subsumption algorithms for kept clauses c at run time, in order to execute subsumption more efficiently. The compilation requires some time, but it can improve the performance in general, if the subsumption algorithm between c and other clauses will be called several times. Since clauses c_1, c_2 in $match(c_1, c_2)$ play different roles, there are two ways to specialize the subsumption algorithm *match*. If c_1 is given and c_2 is unknown, we obtain a specialized algorithm $smatch_{c_1}(c_2)$. It is used for the subsumption of new clauses by c_1 (forward subsumption). In the other case, if c_2 is given and c_1 is unknown, we obtain a specialized algorithm $smatch_{c_2}(c_1)$. It is used to subsume c_1 by new clauses (back subsumption). The two specializations are quite different. For simplicity and because lack of space, we shall consider only forward subsumption.

We shall represent the specialized matching algorithms as a sequence of instructions of an abstract machine. All possible instructions form "the building blocks" of matching algorithms. For reasons that will be clear in Section 3 it is preferable that the instruction set be as simple as possible. We shall show an example how to produce these instruction sets for the matching problem.

```
Algorithm match_{P(g(x0,x1),f(a,x0,x2))}(c)
  Let q be a position
  Let subst, post be initially empty arrays of positions
begin
  q := c;                                          Initialize
  if var(q) or functor(q) ≠ P then exit with fail; Check P
  q := down(q);                                    Down
  if var(q) or functor(q) ≠ g then exit with fail; Check g
  post[0] := right(q);                             Push 0
  q := down(q);                                    Down
  subst[0] := q;                                   Put 0
  q := right(q);                                   Right
  subst[1] := q;                                   Put 1
  q := post[0];                                    Pop 0
  if var(q) or functor(q) ≠ f then exit with fail; Check f
  q := down(q);                                    Down
  if var(q) or functor(q) ≠ a then exit with fail; Check a
  q := right(q);                                   Right
  subst[2] := q;                                   Put 2
  q := right(q);                                   Right
  subst[3] := q;                                   Put 3
  if subst[0] ≠≠ subst[2] exit with fail;          Compare 0 2
  exit with success                                Success
end
```

Fig. 3. The specialized matching algorithm for the clause $P(g(x_0, x_1), f(a, x_0, x_2))$. In the right column there are instructions of the abstract subsumption machine. Unnecessary instructions are shown in roman letter, like Right.

Let $c_1 = P(g(x_0, x_1), f(a, x_0, x_2))$. When we partially evaluate *match* with respect to c_1 as the first argument, we obtain the specialized algorithm shown in Figure 3.*

Each line of this specialized algorithm is an instruction of our *abstract subsumption machine*. There are nine different kinds of instruction (*Initialize*, *Check*, *Down*, *Right*, *Push*, *Pop*, *Put*, *Compare*, Success). They are displayed in the right

* We also perform an obvious optimization. Instead of the sequence of instructions

$push(E, post)$,
$q := pop(post)$,

where E is an expression, we use the equivalent $q := E$.

column of Figure 3. Some of instructions have parameters (*Check*, *Push*, *Pop*, *Put* and *Compare*).

One can imagine the abstract subsumption machine as a kind of Turing machine. q plays the role of the head of the machine. Instead of moving along the tape, the head q moves on the set of positions in the clause. It uses two arrays: *post* to record positions to be considered later and *subst* for keeping tracks of terms substituted for variables. The meaning of the set of instructions is the following:

Initialize	Set q to the initial position (the term c).
Check P	Check that $functor(q)$ is P.
Down	Go down the current position (i.e., from a term to its arguments).
Right	Go to the right (i.e., from an argument in a term to its next argument).
Push n	Push the position to the right off the nth position in *post*.
Pop n	Set q to the nth position in *post*.
Put n	Put the current position on the nth position in *subst* (as the substitution for the nth variable).
Compare m n	Compare terms at positions m, n in *subst* (both positions correspond to the substitution for the same variable).
Success	Exit with success.

If a *Compare* or a *Check* instruction fails, then the whole algorithm fails. We shall prove in Theorem 2 below that these nine types of instruction are enough to specialize $match(c_1, c_2)$ for any c_1. The proof will use a specializing algorithm.

The specializing algorithm *matchspec* accepts a term c_1 as the input and constructs the instruction sequence for the subsumption of c_2 by c_1. The algorithm *matchspec* is shown in Figure 4. The following theorem shows the soundness and correctness *matchspec* algorithm.

THEOREM 2. *For all unit clauses* c_1, c_2, $matchspec(c_1)(c_2)$ *is equivalent to* $match(c_1, c_2)$; *that is one of them terminates with success iff the other one terminates with success.*

Proof (Sketch). By induction on the size of c_1 one can show the following. Let mat_{c_1} be obtained by the unfolding of *match* using c_1 as the first argument. Replace in mat_{c_1} all sequences of instructions of the form

$push(E, post)$,
$q := pop(post)$,

where E is an arbitrary expression by the equivalent $q := E$. It is easy to see that we obtained exactly $matchspec_{c_1}$.

```
Algorithm matchspec(c)
 Let p be a position
 Let subst, post be initially empty stacks of pairs of the form ⟨position,number⟩
 Let instr be an initially empty instruction sequence.
 Let var_count, post_count := 0
begin
 add(Initialize, instr);
 p := c;
 loop
   if var(p)
   then {
     Add Put var_count to instr;
     Add ⟨p, var_count⟩ to subst ;
     var_count := var_count + 1
     }
   else Add Check functor(p) to instr
   if defined(down(p))
   then {
     if defined(right(p))
     then {
       Add Push post_count to instr;
       push(⟨right(p), post_count⟩, post);
       post_count := post_count + 1
       }
     Add Down to instr;
     p := down(p)
     }
   else if defined(right(p))
   then {
     Add Right to instr;
     p := right(p)
     }
   else if nonempty(post)
   then {
     ⟨p, k⟩ = pop(post);
     Add Pop k to instr
     }
   else exit loop
 end of loop
 forall cliques c = ⟨p_0, k_0⟩, ..., ⟨p_n, k_n⟩ ∈ subst
 forall i ∈ {1, ..., n} add Compare k_0 k_i to instr
end
```

Fig. 4. The specializing algorithm *matchspec*(*c*).

In our early experiments we used the abstract subsumption machine in two ways. The first way is to create, for each kept clause c, a structure representing the sequence of instructions of $matchspec(c)$. Then, instead of performing subsumption on two terms c_1, c_2, we used an *interpreter* of the abstract machine. The interpretation could give from no speedup (ground terms) to speedup of the order of 3–4 (clauses with about 10 occurrences of variables). The second way was to *compile* the abstract machine into the native code of our computer. On clauses with 10 occurrences of variables the speedup of the compilation approach was of the order of 10. This latter is a considerable improvement of subsumption performed on the "clause at a time" basis. However, the main idea behind the use of the abstract machine is much more powerful than that. The idea is to combine instruction sequences in *code trees.*

3. Code Trees

It has been widely recognized in the deductive database research that the logic programming "goal at a time"* mode of evaluation is not adequate for handling large amounts of clauses. It is true both for rules producing new clauses and for subsumption algorithms. In this section we demonstrate how to perform subsumption on the "set at a time" basis using *code trees.*

The idea of code trees is simple but extremely powerful. In the preceding section, we created a sequence of the abstract subsumption machine instructions in order to represent a specialized subsumption algorithm $matchspec(c)$ induced by a given clause c. Code trees create similar instructions in order to represent a specialized subsumption algorithm $Match_D$ induced by a given *set* of clauses D. Code trees can be considered as the abstract subsumption machine extended by several new instructions. We gave it a different name to stress that the instructions of $Match_D$ do not form a linear structure as the instructions of $matchspec(c)$. Rather, they are structured as *trees.*** Thus, we *dynamically create an algorithm for the forward subsumption problem induced by a dynamically changing database of clauses.*

Assume that we have a (large) database D of kept clauses and a new clause d. We have to check whether d is subsumed by a clause in D. There are at least two ways to perform subsumption by D on the "set at a time" basis. The first idea is to exploit the common structure of clauses in D in order to select a (small, one hopes) subset of potential candidates for subsumption. This idea has given rise to the use of various indexing schemes, such as discrimination trees (McCune, 1992) or path indexing (Stickel, 1989). Indexing plays the role of a *filter* for the selection of potentially useful clauses. The second idea is to combine *algorithms*

* In deductive databases it is also called "tuple at a time," since deductive databases consider relations as sets of tuples.

** In the case of multi-literal clauses, the specialized code is even more complicated than trees. For example, there are instructions which cause backtracking.

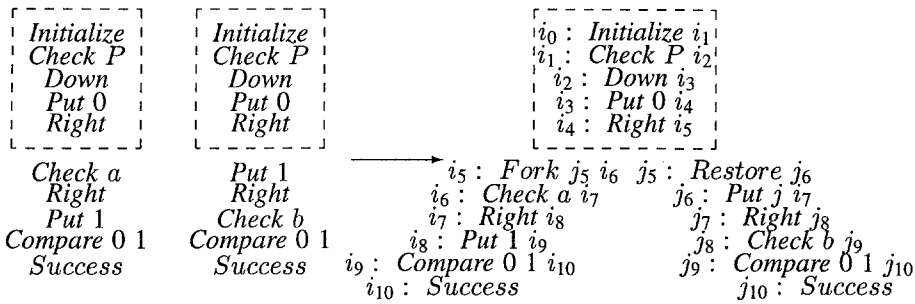

Fig. 5. Two instruction sequences and the corresponding code tree. Shared parts of code are put in a dashed box.

that perform subsumption by D in one algorithm (the code tree). In Section 5 we briefly compare the two approaches.

Let us start with an example. Consider two instruction sets, which represent forward subsumption algorithms for clauses $P(x, a, x)$ and $P(x, x, b)$. They are shown in Figure 5 on the left.

The first five instructions of the two sequences coincide. It means that two algorithms do the same job in the beginning. In order to not repeat the same job we can perform subsumption as follows. First of all, we execute the common five instructions. If one of these common instruction fails, then the algorithm fails. Then we execute the rest of the first instruction sequence. If it fails, we backtrack to the point where the two sequences diverge and execute the rest of the second instruction sequence. For such an algorithm to be sound we have to ensure that all state changes made after the backtrack point be restored after backtracking. To this end we add two additional instructions: *Fork* and *Restore*. *Fork* memorizes values of variables which may have changed after having executed the rest of the first instruction sequence. *Restore* restores the values. We call such a pair *Fork-Restore* a *fork*.* The new set of instructions representing subsumption for this set of two clauses is shown on Figure 5 on the right. The instructions form a tree. Since the set of clauses D can be changing dynamically, each instruction, except for *Success*, must have an additional argument representing the next instruction. This argument may change when we add new clauses to the database D or remove clauses from D. Such sets of instructions represent the subsumption algorithm for the *set* of clauses D. We called them code trees for the following reasons:

1. They can be considered as a code for the subsumption algorithm;
2. They have a tree-like structure;

* We have organized the abstract machine instructions in such a way that the only value that needs to be restored is the position of the head, i.e., the value of q. We achieved this by implementing stacks *subst*, *post* via arrays which resulted in having additional arguments to the *Push*, *Pop*, and *Put* instructions.

3. They can grow and shrink dynamically: new paths can be added and existing paths can be pruned;

The use of the code tree is the following. With each new clause d to be added to the database D, we *integrate* the corresponding instruction sequence $matchspec(d)$ into the code tree T. The integration process consists of finding the longest path in the tree that coincides with $matchspec(d)$ (disregarding forks). Then we add a new fork together with the rest of $matchspec(d)$ at the node where the path must diverge with $matchspec(d)$. We denote the insertion operation $Insert(I, T)$.

We also use the procedure of *removing* the path $matchspec(d)$ from the tree. We need to remove it for efficiency reasons when d is removed from the database (for example by back subsumption). The removal consists of finding the path that corresponds to $matchspec(d)$ and removing the part of this path starting from the last fork together with the fork. We denote the removal operation $Remove(I, T)$.

When we add/remove instruction sequences, we have to also change some arguments of instructions in the code tree, as can be seen on Figure 5. Assume that the tree consists of the instructions for the first clause $P(x, a, x)$, and we merge the instructions for the second clause $P(x, x, b)$ into it. Then the argument of the *Right* instruction must be changed from the address of the *Check* a instruction to the address i_5 of $Fork$.

Now we shall explain the semantics of the instruction set used in code trees similar to the one given on Figure 3. First of all we note that all previously used instructions have an additional argument representing the next instruction. After having executed an instruction, we have to pass the control to this next instruction. For example, *Right* n now means

$q := right(q)$;
`goto` n

The other instructions are changed in the same way. To implement new instructions $Fork$ and *Restore*, we introduce two new stacks *back* and *heads* representing backtrack points and head positions that need to be restored. Each $Fork$ instruction creates a backtrack point and memorizes the current position of the head q. Each *Restore* instruction restores the position of q. Instructions *Check* and *Compare* which may fail now try to backtrack to the previously stored backtrack point:

```
if var(q) or functor(q) ≠ f
then if empty(back)
     then exit with fail
     else goto pop(back)
else goto i_k                    Check f i_k
...
```

$push(q, heads)$;
$push(i_j, back)$;
`goto` i_k; $\quad$ *Fork* i_j i_k
...
$q := pop(heads)$;
`goto` i_k; $\quad$ *Restore* i_k
...

We defined the semantics of the code tree instructions. To prove that we use code trees correctly, we introduce some formal definitions. We call *extended instructions* all previously introduced instruction having the additional argument representing the next instruction. A *code tree* is any finite set of extended instructions. We call arguments representing instructions *addresses* of these instructions $(i_0, \ldots, j_{10})$ in Figure 5. We shall often identify instructions in a code tree with their addresses. An *extended path* in the code tree is any sequence of instructions defined by induction in the following way:

1. Any single instruction is an extended path.
2. Let $i_1, \ldots, i_n$ be an extended path.
   (a) If i_n has the instruction i as its last argument, then $i_1, \ldots, i_n, i$ is an extended path.
   (b) If i_n is a *Fork* j k instruction, then $i_1, \ldots, i_n, j$ is an extended path.

Paths in the code tree are obtained from extended paths by deleting all *Fork* and *Restore* instructions and omitting the last argument of each remaining instruction, except for *Success*. For example, the extended path

i_1: *Initialize* i_6
i_6: *Fork* i_7 i_8 $\quad$ gives the path $\quad$ *Initialize*
i_8: *Restore* i_{12} $\quad$ *Check P*
i_{12}: *Check P* i_{16}

A *completed (extended) path* is an (extended) path whose last instruction is *Success*. An instruction j is *reachable* from i in the code tree T iff there is an extended path $i, \ldots, j$ in T. A code tree T is *consistent* iff it satisfies the following properties:

1. There is exactly one *Initialize* instruction in T.
2. Every instruction $i \in T$ is reachable from the *Initialize* instruction.
3. For every instruction $i \in T$ that is not a *Success* instruction, there is a *Success* instruction reachable from I.
4. For every instruction $i \in T$, the last argument of i is not a *Restore* instruction.
5. For every *Fork* i j instruction, i is a *Restore* instruction.
6. There are no loops in T; i.e., the lengths of extended paths in T are restricted.

We shall only consider consistent code trees.

A code tree T *corresponds* to a set of clauses $c_1, \ldots, c_n$ iff the set of all paths in T is exactly $\{match_{c_1}, \ldots, match_{c_n}\}$. Now we can formally define the *execution* of the code trees T with respect to a clause d as the sequence of extended instructions, starting from the *Initialize* instruction, which would have been executed by T on d according to the given semantics of the extended instructions.

THEOREM 3. *Let the consistent code tree T correspond to the set of clauses C, and let d be a clause. Then the execution of T on d terminates with success iff d is subsumed by a clause in C.*

Proof (Sketch).

1. Assume that the execution succeeds. Consider the sequence I of instructions executed by T. If a backtrack happened in I (i.e., two instructions $Fork\ i\ j$ and i: $Restore\ k$ were executed), then the part of I between this two instructions can be safely removed. After removal of all backtrack points from I, we obtain an extended path J. Let J' be the path obtained from J by removing all $Fork$ and $Restore$ instructions. It is easy to see that the executions of J and J' give the same result, i.e., both of them succeed. We know that $J' = matchspec(c)$ for some $c \in C$. Applying Theorems 1 and 2, we obtain that d is subsumed by c.
2. Assume that the execution fails. Since any instruction $i \in T$ is reachable from the *Initialize*, there is at least one failing *Check* or *Compare* instruction on any path. Applying the same arguments as above, we obtain that no $c \in C$ subsumes d.

A consistent code tree T is in the *tree form* iff

1. The last arguments of any two instructions in T are different.
2. The first arguments m of any two different $Fork\ m\ n$ instructions are different.

THEOREM 4. *Let the code tree T be in the tree form with n instructions. Then, for every clause c, execution of T on c calls at most n instructions.*

Proof. Trivial.

This theorem formally states the importance of having fewer instructions in a code tree. It is thus important to have more *shared instructions* in the tree, i.e., instructions that lie on many paths. The need for more sharing explains some features of the design of code trees. In many cases we could speed single instruction sequences for subsumption by changing the order of instructions or by introducing new instructions. But it would create code trees with less sharing and thus decrease the efficiency of the "set at a time" subsumption.

The following theorem shows that we achieved a maximal sharing on some instructions.

THEOREM 5. *Let* $I_1 = matchspec(c_1)$ *and* $I_2 = matchspec(c_2)$. *Let the first* n *instructions of* I_1, I_2 *coincide. Let the* $n+1th$ *instruction of* I_1, I_2 *be* $Push\ m, Push\ n$. *Then* $m = n$. *The same holds for Pop and Put instructions.*

Proof. By routine inspection of *matchspec*.

THEOREM 6. *Let* T *be a consistent code tree in the tree form. Let* $I = matchspec(c)$. *The* $Insert(I,T)$ *and* $Remove(I,T)$ *are consistent code trees in the tree form. Let, in addition,* T *correspond to a set of clauses* C. *Then*

1. *If no clause in* C *is a variant of* c, *then* $Insert(I,T)$ *corresponds to* $C \cup c$, *otherwise* $Insert(I,T)$ *corresponds to* C;
2. $Remove(I,T)$ *corresponds to* C *without all variants of* c;

Proof. Straightforward, but tedious.

Note that by using code trees we obtained an interesting technique of indexing in presence of variables. It has been noted in (Sudarshan, Ramakrishnan, 1993) that the non-ground subsumption-checking is difficult: "In the general case, subsumption-checking is a costly operation, and we are not aware of efficient subsumption-checking techniques for the case of arbitrary non-ground facts...." Variables create a problem for indexing because in most logic programming and theorem proving implementations they are implemented as addresses in memory. Any two occurrences of variables in two different clauses are two *different* addresses. Our technique discriminates variables in the order of their occurrences in clauses. Thus, it *treats variables occurring in different clauses in a uniform way.* We do not know any implementation of a logic programming or theorem-proving system that makes use of this idea. In (Graf, 1994) a technique is proposed that tries to find some dependencies in path indexes in a uniform way. Ideas of that paper have something in common with our paper.

There are some quite obvious optimizations of code trees. For example, one can use hashing instead of forking in the tree. When the set of function/predicate symbols is fixed in advance, one can even implement parts of code trees via arrays. Another possibility is to introduce explicit backtrack points in the trees, instead of putting them in a stack at run time. It makes the insertion and removal operations on trees more costly, but it can speed the evaluation in general. In deductive database applications, where tuples are usually implemented via records, one can consider an alternative to *Down* and *Right* instructions. It is also possible to access arguments of functions in a nonstandard order, depending on their types. When the main memory is insufficient to keep the whole code trees, we can organize them in such a way that they will load/save their parts automatically and keep information about subtree sizes etc. In general when using code trees, one can try to get the best both of the particular properties of data and the properties of the computer.

Let us analyze at which circumstances code trees are most efficient. The first necessary condition for creating code trees is to have a simple abstract machine

with a small number of instructions for executing a single problem. The second condition is that the structure of code admit sharing. The third condition is that the execution of instructions on a single problem should either produce no side-effects, or to produce side-effects which are easily recoverable. In the case of matching, the only side effect was the change of the position of the head q, which we restored upon backtracking. Fortunately, all major procedures used in automated deduction satisfy these three properties in most cases.

4. Multi-Literal Clauses

In this section we consider multi-literal clauses. The subsumption-check for multi-literal clauses is much harder than that for unit clauses. It is known that the subsumption-check for multi-literal clauses is NP-complete (see, e.g., (Garey, Johnson, 1979)). The usual indexing methods used in theorem-proving programs do not work well with multi-literal clauses. For example, in Otter the clause $P(x), P(f(f(x)))$ will be selected as a potential candidate to subsume *any* clause in which the predicate symbol P occurs, since discrimination trees and path indexes in Otter make indexing on literals, not clauses.

Assume that we need to make the subsumption-check for clauses c, d. Then every literal in c must subsume a literal in d, and the matching substitutions must be compatible for all literals in c. Now every literal in d can be selected as a potential candidate for a literal in c.

The multi-literal subsumption can also be specialized. All commands used in the abstract subsumption machine for unit clause are needed for multi-literal clauses. In addition, we need two new instructions. The first instruction, *First* sets the first literal in d as the current potential candidate to be subsumed by a literal in c. The second instruction *Next* resets this candidate to be the next literal. Both instructions create new backtrack points: *Next* can be executed as many times as the number of literals in d, that is unknown at the time when we compile the specialized procedure $subsume_c$. Thus, both commands have an additional argument that denotes the backtrack point, namely, the (address of) the *Next* instruction. Since c can also be multi-literal, we need to remember the current candidates for each literal in c. To this end we introduce an array Q instead of the head q. Now q will denote the number of the current element in Q and $Q[q]$ can be considered as the head of the abstract machine. The semantics of the *Initialize* instruction changes. The new instructions of our machine have the following semantics:

$q := -1;$
`goto` j *Initialize* j
...
$q := q + 1;$
$Q[q] := d;$
$Push(i, back)$

```
goto j                                   First i j
...
i: if defined(right(Q[q]))
then
{
  Q[q] := right(Q[q]);
  Push(i, back)
}
else
{
  q := q − 1;
  if empty(back)
  then exit with fail
  else goto pop(back)
}
goto j                                   Next i j
...
```

More informally, the *First* i j command sets the (position of the) first literal of d as the candidate to be subsumed by the current literal in c (literal number q). *First* also sets an additional backtrack point i, corresponding to a *Next* instruction. *Next* i j is executed upon backtracking. It tries to reset the candidate for the literal number q to the next literal in d. If this next literal exists, *Next* i j puts the address i of itself as the new backtrack point. Otherwise, it backtracks, resetting q to $q - 1$, i.e., trying to reconsider the matching for the previous literal.

This modification of the subsumption abstract machine is friendly to the code trees idea. The reason is that we still have a small number of instructions with easily recoverable side-effects and good sharing properties.

The specializing algorithm for multi-literal clauses is very similar to the one for unit clauses. It compiles all literals one by one and adds two instructions *First* and *Next* before the code for every literal. The merge of instruction sequences into code trees is also very similar to the one for unit clauses. For example, the instruction sequences for clauses $P(a, x), P(x, b)$ and $P(a, y), P(f(y), y)$ and the corresponding code tree are given in Figure 6. As one can see, the code tree for this example shares the code for the first literals in both clauses and a part of the code for the second literal.

The code tree in this example is, strictly speaking, not a tree any more. The execution of the code tree is not the traversal of a relevant part of the tree. Now some parts of the code tree may be executed several times upon backtracking caused by the execution of *Next* instructions. It means that Theorem 4 about the number of steps called by the execution of a code tree is not true for multi-literal clauses. All other statements from the preceding sections remain true for multi-

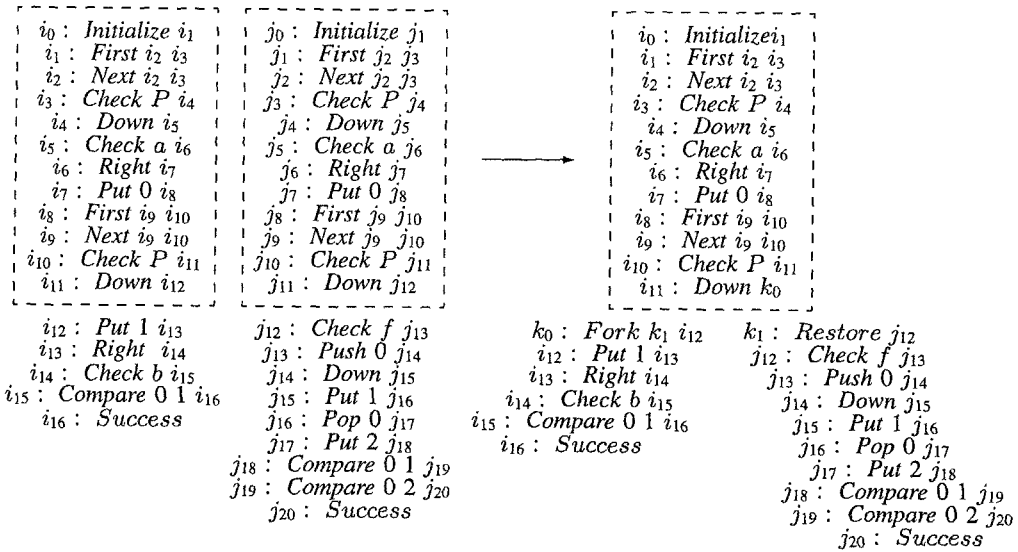

Fig. 6. Instruction sequences for multi-literal clauses $P(a, x)$, $P(x, b)$ and $P(a, y)$, $P(f(y), y)$ and the corresponding code tree.

literal clauses. However in the multi-literal case they are not obvious, since code trees now may have infinite paths caused by repeated applications of the *Next* instruction.

Instructions, extended instructions, and *code trees* are defined as in the preceding section, but may now include *First* i j and *Next* i j instructions. An *extended path* is now defined in the following way:

1. Any single instruction is an extended path.
2. Let $i_1, \ldots, i_n$ be an extended path.
   (a) If i_n has the instruction i as its last argument, then $i_1, \ldots, i_n, i$ is an extended path.
   (b) If i_n is a *Fork* j k, or a *First* j k or a *Next* j k instruction, then $i_1, \ldots, i_n, j$ is an extended path.

An extended k-path is an extended path in which the number of consecutive *Next* instructions is $\leqslant k$.

As before, a *completed* extended path is an extended path whose last instruction is *Success*. A code tree T is *consistent* iff it satisfies the following properties:

1. There is exactly one *Initialize* instruction in T.
2. Every instruction $i \in T$ is reachable from the *Initialize* instruction.
3. For every instruction $i \in T$ which is not a *Success* instruction, there is a *Success* instruction reachable from I.
4. For every instruction $i \in T$, the last argument of i is not a *Restore* or a *Next* instruction.
5. For every *Fork* i j instruction, i is a *Restore* instruction.
6. For every *First* i j instruction, i is a *Next* instruction.
7. For every *Next* i j instruction, i is (the address of) this instruction.
8. All loops in T are caused by repeated applications of the *Next* instruction. More precisely, for every $k \geqslant 1$, the lengths of extended k-paths in T are restricted.

Let us denote by *subsume* and *subsumespec* the subsumption algorithm and the specialization algorithm for multi-literal clauses, which are constructed based on *match* and *matchspec*.

A code tree T *corresponds* to a set of clauses C iff the set of all 1-paths is exactly $\{subsumespec(c) \mid c \in C\}$.

THEOREM 7. *Let the consistent code tree T correspond to the set of clauses C and d be a clause. Then the execution of T on d terminates with success iff d is subsumed by a clause in C.*

Proof (Sketch).

($\Rightarrow$) Assume that the execution of T terminates with success. Consider the sequence of instructions $I = i_1, \ldots, i_n$ which have been followed during the execution. Let J be obtained from I by removal of all backtrack points in the following way:

1. If there is an occurrence of two instructions *Fork* i j and i: *Restore* k, remove both instructions and all instructions in between.
2. If there is an occurrence of two instructions *First* i j and i: *Next* i k, remove all instructions between the two.

It is easy to see that the execution of J will have the same effect as the execution of I. Let P be the extended 1-path obtained from J in the following way: remove all multiple occurrences of the same *Next* instruction and add *Next* after each *First* not followed by *Next*. Thus, there is a clause $c \in C$ such that $subsumespec(c) = P$. Consider the substitution θ obtained by the execution of I. Let f: $N \to N$ be a partial function of natural numbers defined as follows: if the kth *First* instruction in J is immediately followed by n *Next* instructions, let $f(k) = n+1$. Using the properties of the *matchspec* algorithm and properties of *First*, *Next*, one can prove $c_k\theta = d_{f(k)}$, where c_j (d_j) denotes the jth literal of the clause c (d). Thus, c subsumes d.

($\Leftarrow$) Similarly. Let $c \in C$ subsumes d such that there is a substitution θ with $c_i\theta = d_{f(i)}$. Consider the extended path in T defined as previously. One can check that either the execution of T on d will follow this path and terminate with success, or there is a path terminated with success before.

It is not possible any more to prove the theorem on the number of steps in the code tree execution, since instructions *Next* can cause a backtracking that causes the exponential upper bound on the number of steps. It is not surprising, since subsumption on multi-literal clauses is NP-complete. The exponent arises when we have (failed) *Next* instructions inside other *Next* instructions. Nevertheless, code trees show superior behavior over other approaches on multi-literal clauses for the following reason: such deep backtrack points may be merged together in code trees, thus avoiding re-executing them many times.

Theorem 5 about sharing of code in code trees remains valid for multi-literal code trees. Insertion/removal of code in the code tree for the multi-literal case are similar to those of the trees for matching. As in the unit case, both operations applied to consistent code trees give consistent code trees and Theorem 6 also holds in the multi-literal case.

5. Comparison with Indexing

A number of approaches have been proposed to improve the performance of subsumption and other algorithms used in bottom-up procedures. Some of them are reviewed in (McCune, 1992). Most of these techniques use indexing to filter out the relevant clauses and to perform the algorithms "set at a time". There are two main types of indexing schemes which are used in modern theorem provers – path indexing (Stickel, 1989) and discrimination trees (McCune, 1992). In this section we briefly compare code trees with indexing schemes.

Discrimination trees are trees that encode structure of terms occurring in clauses. Paths in the discrimination trees correspond to the terms. An example of a set of terms and its discrimination tree is shown on Figure 7 (the example is taken over from (McCune, 1992)).

The main difference in the design of code trees and indexing schemes is the following. Indexing is used to filter out relevant clauses for a procedure. It may often be not enough for efficient algorithms. Thus, indexing schemes are usually provided with additional information. For example, in Otter leaves of the discrimination tree contain not only references to clauses, but also the lists of variables in clauses. It sometimes helps to complete an operation, e.g., subsumption, without actually doing subsumption. The traversal (or, better, execution) of code trees is a complete algorithm. Code trees use a different idea: instead of indexing on data they dynamically create the *code* for a procedure. In the case of forward subsumption, where the structure of specialized procedures is very close to the structure of data, code trees and discrimination trees are very similar. However the code trees for back subsumption and unification are very different from those of forward subsumption. Experience shows that discrimination trees do not perform well for back subsumption. One of the reasons lies in the fact that the code for back subsumption does not so closely resemble the structure of the kept clause as does the code for forward subsumption. In the case of unification, the specialized instructions for a clause already form a tree with forks created by if-then-else statements.

It is also very important that code trees can be translated into the native code of a computer. We made experiments with specialized algorithms compiled into the native code. The real compilation can give another speedup factor of 3–4.

The idea of indexing in the case of subsumption can be summarized as follows: find structural properties P_1, P_2 of data such that whenever c_1 satisfies P_1 and t_2 satisfies P_2, then c_1 subsumes (or does not subsume) c_2. Then keep information about these P_1, P_2 in special data structures, for example discrimination trees. We propose to exploit shared code, which is not the same as indexing. We do

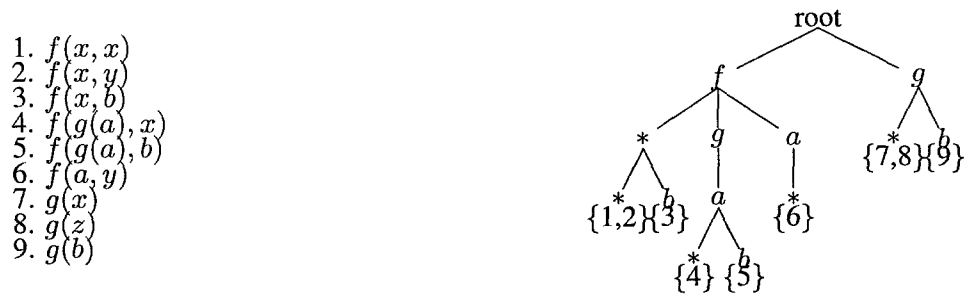

Fig. 7. An example of a set of terms and its discrimination tree.

not claim that the use of code trees is always superior over indexing. They are just another technique to implement a class of procedures.

We have already seen how the use of code instead of data helped us find a new way of indexing on variables. The use of code trees also resulted in a new solution for subsumption on multi-literal clauses. It shows that the decision to manipulate code instead of data helps one find new design decisions. *Code trees provide an important methodology for designing bottom-up procedures. The consistent use of this methodology helps to find new solutions to old problems.* We believe that code trees give a different viewpoint on problems and their solutions, compared with the idea of indexing data.

Indexing methods may require less memory than code trees, when several procedures participate in deduction. In theorem-proving applications, some problems use resolution, both kinds of subsumption, rewriting and paramodulation. The full use of code trees would require five different trees to be constructed. But all these procedures may use the same indexing scheme. At the same time, if one would try to provide indexing schemes with more information helpful for an algorithm, it may result in a huge memory consumption, as was shown in (Graf, 1994) for path indexing.

6. Experiments

Here we give a statistics of two examples proved in our theorem-proving system *Vampire* on a Hewlett Packard 735. We did not try to set the best methods/switches to solve the two problems. On the contrary, we tried to set options in such a way so that to increase the number of forward subsumptions required.

One of the hardest problems dealing with *unit clauses* was the condensed detachment 4 problem from (Wos *et al.*, 1991). The initial clauses are the following:

1. $P(x_1), \neg P(x_2), \neg P(i(x_2, x_1))$.
2. $\neg P(i(i(i(a, i(b, c)), c), i(b, a)))$.
3. $P(i(x_1, i(i(x_2, i(x_3, x_1)), i(x_3, x_2))))$.

The problem has been tried by positive hyperresolution with the restriction 20 on the maximal weight of clauses.* On this problem positive hyperresolution generates only unit clauses. A typical example of a clause used in this proof is

$$P(i(i(i(i(i(x_1, x_2), i(x_2, x_3)), i(x_4, i(x_2, x_3))), x_5), i(x_6, x_5))).$$

Subsumption is very hard for such kinds of problem for the following two reasons:

* Here the weight of a clause is the number of occurrences of symbols in the clause. For example, the weight of $P(f(x)), Q(a)$ is 5.

TABLE I.

Problem/options	Kept clauses	Subsumed clauses	Clause-to-clause subsumptions	Forward subsumption time	Subsumptions per second	Database subsumptions per second
Condensed detachment (weight = 20)	30,509	497,141	1,581,472,615	20.14	78,523,963	26,199
Condensed detachment (weight = 24)	163,509	534,456	17,763,995,378	81.37	218,311,360	8,577
Steamroller (weight = 17)	12,161	872,061	2,861,295,196	355.53	8,059,986	2,491
Steamroller (weight = 20)	11,012	803,982	2,173,049,289	539.44	4,028,342	1,510
Steamroller (weight = 22)	26,185	2,074,607	9,695,480,345	2422.34	4,003,088	867
Steamroller (no restrictions)	41,743	3,048,078	17,133,045,644	4287.16	3,996,362	711

TABLE II.

	Typical clause	Kept	Forward subsumed	Forward subsumption time	Ratio
	Condensed detachment (maximal weight = 20)				
Vampire	$P(i(i(i(i(i(x_1,x_2),i(x_2,x_3)),$	30,509	497,141	20.14	2.36–2.53
Otter	$i(x_4,i(x_2,x_3))),x_5),i(x_6,x_5)))$	28,390	482,119	46.04	

1. The number of variables in clauses is quite big (usually 7 to 10 in this example).
2. There is only one predicate symbol and only one function symbol. Thus, subsumption never fails on the comparison of function symbols.

On *multi-literal clauses,* we tried the "steamroller" problem from the Pelletier list (Pelletier, 1986).

The problem has been tried by binary resolution (no restrictions or optimizations). To increase the database size and the number of forward subsumptions, we switched back subsumption off. This problem have been tried with different weight restrictions. Typical clauses used in the found proof are the following:

5. $R(x_1, f(x_1)) \vee \neg P_5(x_1)$.
14. $R(x_1,x_2), R(x_1,x_3), \neg Q_0(x_3), \neg R(x_2,x_4),$
$\neg Q_0(x_4), \neg S(x_2,x_1), \neg P_0(x_2), \neg P_0(x_1)$.
1277. $\neg P_0(c_5), \neg P_0(c_4), \neg Q_0(x_1), \neg R(c_4,x_1), \neg Q_0(x_2), R(c_5,x_2), R(c_5,c_4)$.

There are no deep terms involved in the proof search, as in the first problem. But subsumption-checks for this problem are even more difficult for the following reason. This problem generates only a very small number of relatively short clauses. Most of clauses have from 7 to 11 literals. (The algorithm *subsume* on clauses with m and n literals must perform m^n matches in the worst case.)

We used a counter that counted the number of subsumption-checks which would have been executed if subsumption had been implemented on the clause-to-clause basis. The results are given in Table I. (For the weight = 24 case a proof was not obtained because of insufficient memory. The figures in the table are given up to the moment when Vampire ran out of memory after using about 62 megabytes.)

As one can see (Table I), Vampire achieved a speed of more than 200,000,000 subsumptions per second on unit clauses, where code sharing is very high. The speed is still very high (about 4,000,000 subsumptions per second) on multi-literal clauses with up to 15 literals.

TABLE III.

	Typical clause	Kept	Forward subsumed	Forward subsumption time	Ratio
	Steamroller (maximal weight = 17)				
Vampire	$\neg Q0(x0), \neg Q0(x1), \neg P2(g), \neg P0(g),$	12,161	872,061	404	17.3–23.1
Otter	$\neg S(g, g), R(g, x1), \neg R(g, x0)$	9,091	56,579	453	
	Steamroller (maximal weight = 20)				
Vampire	$\neg Q0(x0), \neg Q0(x1), \neg Q1(x2), \neg P4(x2), \neg P3(x1),$	11,012	803,982	539	17.7–18.8
Otter	$\neg P0(x3), \neg P0(x1), \neg S(x1, x3), \neg R(x1, x0)$	10,384	57,752	687	
	Steamroller (maximal weight = 22)				
Vampire	$\neg Q0(x0), \neg Q0(x1), \neg Q1(x2), \neg P4(x2), \neg P3(x1), \neg P0(x3),$	26,185	2,074,607	2422	15.4–35.3
Otter	$\neg P0(x1), \neg P1(x3), \neg S(x1, x3), \neg R(x1, x0)$	11,397	58,567	1,052	
	Steamroller (no restrictions)				
Vampire	$\neg Q0(x0), \neg Q0(x1), \neg Q1(x2), \neg P5(x3), \neg P3(x2), \neg P2(x2),$	41,743	3,048,078	4287	49.2–128.9
Otter	$\neg P0(f6(x3)), \neg P0(x2), \neg S(f(x3), x2), R(x2, x1), \neg R(f6(x3), x0)$	15,949	63,970	4430	

In Tables II and III we compare the performance of Vampire with that of Otter (Otter uses discrimination trees) on forward subsumption. The Table is for the condensed detachment problem (unit clauses only).

Table III demonstrates a series of experiments with multi-literal clauses. It shows that the difference in perfomance between Vampire and Otter is growing when the clauses have more literals.

The minimal ratio in this Table III is based on the number of subsumed clauses per second, the maximal ratio takes into account the size of the database (i.e., the number of kept clauses).

In the case of unit clauses, the difference in performance is not high, since Otter uses a version of discrimination trees that encodes essentially the same information as code trees. The two provers use similar ways of clause generation, so the number of kept/subsumed clauses is essentially the same. In the case of the Steamroller problem, the binary resolution algorithms used by the two provers were very different, which resulted in very different numbers of subsumed clauses. The difference is mostly due to the treatment of factoring in Vampire and Otter. However, the average size of generated clauses was about the same, so the comparison of performance is fair.

Acknowledgements

I am grateful to all people who made a direct or indirect influence on this paper. These are Anatoli Degtyarev, Ewing Lusk, Bill McCune, Micha Meier, Ross Overbeek, Mark Wallace, David S. Warren, and Larry Wos. The motivation to efficiently implement some basic procedures is mostly due to the study of the excellent performance of Otter implemented by Bill McCune.

References

Astrakhan, O. L. and Stickel, M. E.: Caching and lemmaizing in model elimination theorem prover in D. Kapur (ed.), *11th Int. Conf. on Automated Deduction*, Vol. 607 of Lecture Notes in Artificial Intelligence, Saratoga Springs, NY, June 1992. Springer Verlag, pp. 224–239.

Bancilhon, F., Maier, D., Sagiv, Y., and Ullman, J. D.: Magic sets and other strange ways to implement logic programs, in *Proc. of the 5th ACM SIGMOD-SIGACT Symposium on Principles of Database Systems*, Cambridge, MA, March 1986, pp. 1–15.

Beeri, C., Naqvi, Sh., Ramakrishnan, R., Shmueli, O., and Tsur, Sh.: Sets and negation in a logic database language (LDL1), in *Proc. 6th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, ACM Press, 1987, 21–36.

Beeri, C. and Ramakrishnan, R.: On the power of Magic, in *Proc. 6th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, ACM Press, 1987, pp. 269–283.

Codish, M. and Demoen, B.: Analyzing logic programs using "Prop"-ositional logic programs and a magic wand, in Dale Miller (ed.), *Logic Programming – Proc. of the 1993 Int. Symposium*, The MIT Press, 1993, pp. 114–129.

Das, S. K.: *Deductive Databases and Logic Programming*, Addison-Wesley, 1992.

Decker, H.: Integrity enforcements on deductive databases, in *Proc. of the 1st Int. Conf. on Expert Database Systems*, Charleston, South Carolina, April 1986, pp. 271–285.

Garey, M. R. and Johnson, D. S.: *Computers and Intractability*, Freeman, San Francisco, 1979.

Goller, C., Letz, R., Mayr, K., and Schumann, J.: SETHEO V3.2: Recent developments, in A. Bundy (ed.), *Automated Deduction – CADE-12. 12th Int. Conf. on Automated Deduction*, Vol. 814 of Lecture Notes in Artificial Intelligence, Nancy, France, June/July 1994, pp. 778–782.

Gottlob, G. and Leitsch, A.: On the efficiency of subsumption algorithms, *J. Association for Computing Machinery* **32**(2) (1987), 280–295.

Graf, P.: Extended path-indexing, in A. Bundy (ed.), *Automated Deduction – CADE-12. 12th Int. Conf. on Automated Deduction*, Vol. 814 of Lecture Notes in Artificial Intelligence, Nancy, France, June/July 1994, pp. 514–528.

Lusk, E. L.: Controlling redundancy in large search spaces: Argonne-style theorem proving through the years, in A. Voronkov (ed.), *Logic Programming and Automated Reasoning. International Conference LPAR'92*, Vol. 624 of Lecture Notes in Artificial Intelligence, St. Petersburg, Russia, July 1992, pp. 96–106.

Lloyd, J. W., Sonenberg, E. A., and Topor, R. W.: Integrity constraint checking in stratified databases. Technical Report 86/5, Department of Computer Science, University of Melbourne, 1986.

Lloyd, J. W. and Topor, R. W.: A basis for deductive database systems, *J. Logic Programming* **2**(2) (1985), 93–109.

Maeda, A. M., Aoe, J.-I., and Tomabechi, H.: Signature-check based unification filter, *Software – Practice and Experience* **24**(7) (1994), 603–622.

Manthey, R. and Bry, F.: SATCHMO: A theorem prover implemented in Prolog, in *CADE'88 (9th Int. Conf. on Automated Deduction)*, Lecture Notes in Computer Science, Argonne, IL, May 1988, pp. 179–216.

Maslov, S. Yu.: An inverse method for establishing deducibility of nonprenex formulas of the predicate calculus, in J. Siekmann and G. Wrightson (eds), *Automation of Reasoning (Classical Papers on Computational Logic)*, Vol. 2, Springer Verlag, 1983, pp. 48–54.

McCharen, J., Overbeek, R., and Wos, L.: Complexity and related enhancements for automated theorem-proving programs, *Computers and Mathematics with Applications* **2** (1976), 1–16.

McCharen, J., Overbeek, R., and Wos, L.: Problems and experiments for and with automated theorem-proving programs, *IEEE Transactions on Computers* **C-25**(8) (1976), 773–782.

McCune, William W.: An indexing method for finding more general formulas, *Association for Automated Reasoning Newsletter* **1**(9) (1988), 7–8.

McCune, William W.: OTTER 2.0 users guide. Technical report, Argonne National Laboratory, March 1990.

McCune, William W.: Experiments with discrimination-tree in indexing and path indexing for term retrieval, *J. Automated Reasoning* **9**(2) (1992), 147–167.

Neiman, V.: Refutation search for horn sets by a subgoal-extraction method, *J. Logic Programming* **9**(2) (1990), 267–284.

Overbeek, R.: An implementation of hyper-resolution, *Computers and Mathematics with Applications* **1** (1975), 201–214.

Pelletier, F. J.: Seventy-five problems for testing automatic theorem provers, *J. Automated Reasoning* **2**(2) (1986), 191–216.

Robinson, J. A.: Automatic deduction with hyper-resolution, *Int. J. Computer Mathematics* **1** (1965), 227–234.

Robinson, J. A.: A machine-oriented logic based on the resolution principle, *J. the Association for Computing Machinery* **12**(1) (1965), 23–41.

Robinson, G. and Wos, L. T.: Paramodulation and theorem-proving in first order theories with equality, in *Machine Intelligence*, Vol. 4. Edinburgh University Press, Edinburgh, 1969.

Stickel, M.: A PROLOG technology theorem prover: Implementation by an extended Prolog compiler, *J. Automated Reasoning* (4) (1988), 353–380.

Stickel, M.: The path indexing method for indexing terms. Technical Report 473, Artificial Intelligence Center, SRI International, Menlo Park, Calif., October 1989.

Sudarshan, S. and Ramakrishnan, R.: Optimizations of bottom-up evaluation with non-ground terms (extended abstract), in Dale Miller (ed.), *Logic Programming. Proc. of the 1993 Int. Symp.*, The MIT Press, 1993, pp. 557–574.

Tambe, M. and Rosenbloom, P. S.: Investigating production system representations for non-combinatorial match, *Artificial Intelligence* **68** (1994), 155–190.

Tamaki, H. and Sato, T.: OLDT resolution with tabulation, in *Int. Conf. on Logic Programming*, 1986, pp. 84–98.

Vieille, Laurent: Recursive query processing: The power of logic, *Theoretical Computer Science* **69** (1989), 1–53.

Voronkov, A.: LISS – The Logic Inference Search System, in Mark Stickel (ed.), *Proc. Int. Conf. on Automated Deduction*, Vol. 449 of Lecture Notes in Computer Science, Kaiserslautern, Germany, 1990. Springer Verlag, pp. 677–678.

Voronkov, A.: Theorem proving in non-standard logics based on the inverse method, in D. Kapur (ed.), *11th Int. Conf. on Automated Deduction*, Vol. 607 of Lecture Notes in Artificial Intelligence, Saratoga Springs, NY, USA, June 1992. Springer Verlag, pp. 648–662.

Warren, David H. D.: An abstract Prolog instruction set. SRI Tech. Note 309, SRI Intl., Menlo Park, Calif., 1983.

Warren, D. S.: Memoing for logic programs, *Communications of the ACM (invited paper)* **35**(3) (1992), 93–111.

Wos, Larry, Overbeek, Ross, and Lusk, Ewing: Subsumption, a sometimes undervalued procedure, in Jean-Louis Lassez and Gordon Plotkin (eds), *Computational Logic. Essays in Honor of Alan Robinson,* The MIT Press, Cambridge, MA, 1991, pp. 3–40.

Wos, Larry: Note on McCune's article on discrimination trees, *J. Automated Reasoning* **9**(2) (1992), 145–146.