

A DEFINITION-DRIVEN THEOREM PROVER

George W. Ernst
Department of Computing and Information Sciences
Case Western Reserve University
Cleveland, Ohio 44106

Abstract

This paper describes a theorem prover, running on a PDP-10-TENEX system, that can prove some theorems whose statements involve a relatively large number of definitions. Such theorems require special methods because (1) their statements contain a large number of clauses and (2) their proofs are quite long although straight-forward.

A theorem is proven by first subdividing it into "simple" subgoals and then using a standard resolution theorem prover to prove the subgoals. The first part of this process involves the substitution of definitions for defined quantities and the use of logical simplification. This process which is more similar to a natural deduction system than a resolution system, is shown to be complete when restricted to first-order logic. However, the theorem prover can deal with some interesting higher-order theorems as is shown by an example.

1. Introduction

The motivation for this research stems from our view of how theorem provers will be used in solving real-life problems. We believe that the first few applications of theorem provers will be man-machine oriented because, at the current state of the art, theorem provers can only solve relatively easy problems and it will probably be quite a while before they can solve truly difficult problems on their own.

It also seems to us that theorem provers will be faced with a large number of problem-oriented predicates due to diversity of real world data. Some of these predicates will be sufficiently basic that they will of necessity be built into the theorem prover. However, most of them, due to their abundance, will have to be defined in terms of the basic predicates and logical connectives.

On the basis of these assumptions, we feel that the role of the machine will be primarily a big data filter, doing the relatively straight-forward part of the theorem proving which might be quite lengthy. For the most part, this amounts to "wading" through definitions and previously proven theorems (and perhaps meta-theorems), many of which are quite long and tedious; proving the easy parts and isolating the non-trivial parts. The man then would select a non-trivial part to work on next and would give the machine "hints" on how to go about proving it. The hints will be of various forms, such as conjecturing induction hypothesis, lemmas, the value of important variables, etc.

This paper describes a theorem prover, TPI, that can prove some lengthy but straightforward theorems whose statement contains defined quantities. Although TPI is a totally mechanical, and not a man-machine theorem prover, the types of problems for which it is designed are similar to what we feel will be the kind of problems that a machine should do in a man-machine application of theorem proving. Since the approach of TPI is somewhat different than a standard resolution

theorem prover, in the next section we will give an example of how it proved a particular theorem. This will motivate Sec. 3 which describes TPI in fair detail.

2. An Example of TPI

In this section we give an example of how TPI proved the theorem that the direct product of 2 groups is a group. The statement of this theorem is*

$$\begin{aligned} & \forall ((G \text{ univ}) (+ (FN G G G)) (- (FN G G)) (O G) \\ & (H \text{ univ}) (* (FN H H H)) (I (FN H H)) (E H)) \quad (1) \\ & [Gr(G,+,0) \ \& \ Gr(H,*,I,E) \rightarrow Gr(GXH,\circ,J,(0,E))] \end{aligned}$$

The operator and inverse of the product group are $\circ$ and J , respectively, and are defined below. In the first line of (1) $(G \text{ univ})$ indicates that G is some element of the universe, and $(+(FN G G G))$ indicates that $+$ is a binary operator on G . The remainder of the first line indicates that $-:G \rightarrow G$ and O is some element of G . The second line of (1) states the same things about $H, *, I$ and E . Thus the first two lines universally quantify 8 variables subject to the restrictions that they must lie in certain domains. Gr is a defined predicate having 4 arguments, a group, its operator, its inverse and its identity, respectively. Thus, the last line of (1) states that, if G and H are groups then so is GXH (cartesian product of G and H). The operator $\circ$ of the latter is

$$\lambda((u \text{ GXH})(v \text{ GXH}))(1^{st}(u)+1^{st}(v), 2^{nd}(u)*2^{nd}(v)) \quad (2)$$

In this expression λ is the usual operator of abstraction. The meaning of (2) is that it is a two-place function. The string $(u \text{ GXH})$ indicates that the first argument which gets substituted for u must be in the set GXH . $1^{st}(u)$ is the first element of the ordered pair u while $2^{nd}(u)$ is the second element of u . The $,$ in (2) is the operator which takes two terms and makes an ordered pair out of them. An example of the use of (2) is $a \circ b$. This expression is defined if both a and b are in GXH and the term $a \circ b$ evaluates to the ordered pair $(1^{st}(a)+1^{st}(b), 2^{nd}(a)*2^{nd}(b))$. Note that $G, H, +, *$ are free variables in (2) which are quantified by the quantifiers in (1) because $\circ$ occurs in (1).

The inverse J of the product group is defined in a manner similar to (2) as

$$\lambda((u \text{ GXH}))(-1^{st}(u), I(2^{nd}(u))). \quad (3)$$

The identity of the product group is given in (1) as the ordered pair (O,E) .

To prove (1), TPI must know the definition of group, 1^{st} , etc. TPI has a repertoire of definitions and previously proven theorems which includes these definitions and other definitions, such as definition of homomorphism, that are not relevant to proving (1). We will not state the definitions because their content will become clear by their use.

First of all, TPI simplifies (1) which in this

*There are 3 differences between (1) and the formula given to TPI: An example of the first is that $Gr(G,+,0)$ is given to the machine as $(Gr G + - 0)$. The second is that infix notation is used in (1) and thirdly parentheses are omitted in (1) if their omission causes no ambiguity.

case involves putting Skolem functions in for the universally quantified variables in (1). For notational purposes we will use the same names as the variables and thus (1) becomes

$$\text{Gr}(G, +, -, 0) \ \& \ \text{Gr}(H, *, I, E) \rightarrow \text{Gr}(GXH, \circ, J, (0, E)). \quad (4)$$

In this formula $G, +$, etc. are constants and their domains are known because of the first 2 lines of (1), e.g., $+$ is a binary operator on G . Next, TPI sees from the definition of group that it is sufficient to prove that $\circ$ is associative; J is a right inverse; and $(0, E)$ is a right identity. The resulting formula is immediately simplified and stated as the following three subgoals (or lemmas) which TPI will work on separately:

$$\begin{aligned} &\text{Gr}(G, +, -, 0) \ \& \ \text{Gr}(H, *, I, E) \rightarrow \\ & (1^{\text{st}}(a) + 1^{\text{st}}(0, E), 2^{\text{nd}}(a) * 2^{\text{nd}}(0, E)) = a \end{aligned} \quad (5)$$

$$\begin{aligned} &\text{Gr}(G, +, -, 0) \ \& \ \text{Gr}(H, *, I, E) \rightarrow \\ & [1^{\text{st}}(b) + 1^{\text{st}}(-1^{\text{st}}(b)), \\ & I(2^{\text{nd}}(b)), 2^{\text{nd}}(b) * 2^{\text{nd}}(-1^{\text{st}}(b)), I(2^{\text{nd}}(b)))] \\ & = (0, E) \end{aligned} \quad (6)$$

$$\begin{aligned} &\text{Gr}(G, +, -, 0) \ \& \ \text{Gr}(H, *, I, E) \rightarrow \\ & [1^{\text{st}}(1^{\text{st}}(c) + 1^{\text{st}}(d), 2^{\text{nd}}(c) * 2^{\text{nd}}(d)) + 1^{\text{st}}(e), \\ & 2^{\text{nd}}(1^{\text{st}}(c) + 1^{\text{st}}(d), 2^{\text{nd}}(c) * 2^{\text{nd}}(d)) * 2^{\text{nd}}(e)] \\ & = [1^{\text{st}}(c) + 1^{\text{st}}(1^{\text{st}}(d) + 1^{\text{st}}(e), 2^{\text{nd}}(d) * 2^{\text{nd}}(e)), \\ & 2^{\text{nd}}(c) * 2^{\text{nd}}(1^{\text{st}}(d) + 1^{\text{st}}(e), 2^{\text{nd}}(d) * 2^{\text{nd}}(e))] \end{aligned} \quad (7)$$

In these formulae a, b, c, d, e are Skolem functions that are elements of $G \times H$. (5), (6), (7) state respectively that $(0, E)$ is a right identity, J is a right inverse and $\circ$ is associative. The reader will note that (5), (6) and (7) are more complicated than textbook formulae because they contain terms that can obviously be simplified such as $1^{\text{st}}(0, E)$ which is just 0. However, these formulae were produced directly from the definition of a group and such simplification will occur later in the proof.

TPI next processes (5) in much the same way that it processed (1). It sees from the definition of an ordered pair that two ordered pairs are equal if their first components are equal and if their second components are equal. Using this definition on (5), the following two new subgoals are produced:

$$\begin{aligned} &\text{Gr}(G, +, -, 0) \ \& \ \text{Gr}(H, *, I, E) \rightarrow \\ & 1^{\text{st}}(a) + 1^{\text{st}}(0, E) = 1^{\text{st}}(a) \end{aligned} \quad (8)$$

$$\begin{aligned} &\text{Gr}(G, +, -, 0) \ \& \ \text{Gr}(H, *, I, E) \rightarrow \\ & 2^{\text{nd}}(a) + 2^{\text{nd}}(0, E) = 2^{\text{nd}}(a) \end{aligned} \quad (9)$$

When TPI processes (8) it notes that (8) cannot be further divided into subgoals, and thus calls on a resolution-based theorem prover called the subgoal solver (SGS), to solve it. But first it replaces defined quantities by their definitions and then reduces the formula to normal form. Fig. 1 gives the result of processing (8) in this way. In Fig. 1 the variable x is restricted to be in G . This information is checked implicitly in TPI whereas a standard resolution theorem prover would append an $x \notin G$ to each clause containing an x , e.g., the first clause in Fig. 1 would be $x \notin G \vee x + 0 = x$. Similarly y is restricted to be in H while u and v are not restricted.

The definition of a group indicates that $+$ and $*$ are associative. Instead of representing this information as clauses, the operators are "flagged" and the SGS processes them as n -ary operators.

As the reader can easily see, Fig. 1 gives a

contradiction in just two steps. From clauses 9 and 12 in Fig. 1 we get $1^{\text{st}}(a) + 0 \neq 1^{\text{st}}(a)$. From this and clause 1 in Fig. 1 we get $1^{\text{st}}(a) \neq 1^{\text{st}}(a)$ which is contradiction.

TPI processes (9) almost identically to the way in which it processed (8). The normal form of (9) is the first 11 clauses in Fig. 1 and $2^{\text{nd}}(a) + 2^{\text{nd}}(0, E) \neq 2^{\text{nd}}(a)$ which also has a two-step contradiction.

TPI then proceeds to (6) and uses the definition of ordered pair to divide (6) and into the following two subgoals:

$$\begin{aligned} &\text{Gr}(G, +, -, 0) \ \& \ \text{Gr}(H, *, I, E) \rightarrow \\ & 1^{\text{st}}(b) + 1^{\text{st}}(-1^{\text{st}}(b)), I(2^{\text{nd}}(b)) = 1^{\text{st}}(0, E) \end{aligned} \quad (10)$$

$$\begin{aligned} &\text{Gr}(G, +, -, 0) \ \& \ \text{Gr}(H, *, I, E) \rightarrow \\ & 2^{\text{nd}}(b) * 2^{\text{nd}}(-1^{\text{st}}(b)), I(2^{\text{nd}}(b)) = 2^{\text{nd}}(0, E) \end{aligned} \quad (11)$$

(10) equates the first components of the left and right sides of the $=$ in (6) while (11) equates their 2nd components.

Since (10) cannot be divided into subgoals, it is reduced to normal form, which is just the first 11 clauses in Fig. 1 and

$$1^{\text{st}}(b) + 1^{\text{st}}(-1^{\text{st}}(b)), I(2^{\text{nd}}(b)) \neq 1^{\text{st}}(0, E) \quad (12)$$

This is passed to the SGS which easily solves it in 3 steps. From (12) and clause 9 in Fig. 1, $1^{\text{st}}(b) + (-1^{\text{st}}(b)) \neq 1^{\text{st}}(0, E)$ is deduced. This and clause 3 in Fig. 1 gives $0 \neq 1^{\text{st}}(0, E)$ which by clause 9 in Fig. 1 implies $0 \neq 0$ which is a contradiction.

TPI processes (11) in much the same way as (10) proving it also in three steps. Proceeding onto (7), TPI breaks it into 2 subgoals by separately equating the two components of the left and right of the $=$ in (7). Since the generation of these subgoals is so similar to the generation of the above subgoals we will not state the subgoals in this paper. Also their proofs will be omitted because they are just 3 steps apiece.

This concludes the proof of (1). In summary then, TPI divides (1) into 6 terminal subgoals or lemmas and proves these 6 subgoals independently which constitutes a proof of (1). The purpose (5), (6) and (7), which are intermediate subgoals, is to facilitate the generation of the 6 terminal subgoals; i.e., (5), (6) and (7) are not directly proven by the SGS.

3. Description of TPI

The purpose of this section is to describe the internal mechanisms of TPI which is a LISP program currently running on a PDP-10 under TENEX. Due to space limitations, the description will be very brief. For this reason the interested reader is referred to ³ which is a fairly detailed description of TPI.

TPI operates from a repertoire of definitions not all of which are needed in any particular theorem. Currently, there are over a dozen defined quantities in TPI's repertoire. Incorporated into the definitions are also some previously proven theorems. All of this information is contained on two lists BD (backward definitions) and FD (forward definitions).

All statements on BD are of the form $Q(A \rightarrow B)$ while FD statements are of the form $Q(B \Leftarrow C)$. B is the defined quantity and Q is a string of restricted quantifiers. A together with C comprises the definition of B . In addition, certain previously proven theorems about B may also be incorporated into these statements as described in Sec. 4. As an example, consider the definition of a group on BD. In this case B is $\text{Gr}(G, +, -, 0)$ and A is

$$\forall(x G)(y G) (z G)[x+0=x \ \& \ x+(-x)=0 \ \& \ (x+y)+z=x+(y+z)] \quad (13)$$

and Q is $\forall((G \text{ univ})+(FN \ G \ G \ G))(- (FN \ G \ G))(O \ G))$. The intent of these formulae should be clear from the discussion in the previous section. FN and univ are primitives whose semantics are built into TPI. Although the relationship between A and B in this example is 'if and only if', in general, A will only imply B for the statements on BD.

Since the relationship is 'if and only if' we could use the same statement on FD, i.e., we could put $Q(B \equiv A)$ on FD. Note that our use of restricted quantification is somewhat non-standard. The restrictions are processed as if they are conjuncts of the second argument of $\equiv$. For example, $\forall((x \ a)(y \ b))E \equiv F$ is taken to mean $\forall xy(E \equiv xea \ \& \ yeb \ \& \ F)$.

TPI attempts to prove theorems by generating an AND/OR tree of subgoals in much the same way as described in ². The root of the tree is the theorem to be proven. TPI processes a node in the tree by first applying some logical simplifications such as "λ-conversion," transforming a theorem of the form $D \rightarrow (E \rightarrow F)$ into $D \& E \rightarrow F$, and removing quantifiers at the outermost level by putting in Skolem functions. (Note that, since the theorem has not been negated, the Skolem functions are put in for universally quantified variables instead of existentially quantified variables.) Then, if the simplified theorem is of the form $E \rightarrow (F_1 \& F_2 \& \dots)$, the subgoals $E \rightarrow F_1$, $E \rightarrow F_2 \dots$ are generated provided that the F_i have no variables in common.

Next for each subgoal of the form $G \rightarrow F_i$, TPI generates the subgoals $G \rightarrow \sigma_{ij} A_{ij}$ where $Q_{ij}(A_{ij} \rightarrow B_{ij})$ are statements on BD and σ_{ij} is the most general unifier of B_{ij} and F_i . These are subnodes of an OR node while the above subgoals come from an AND node. Fig. 2 shows the subgoal tree generated so far. Now the whole process is repeated for the subgoals: logical simplification, generation of AND subnodes, generation of OR subnodes via BD.

Actually this picture is highly simplified. The most notable omission is that TPI checks that the restricted quantification of Q_{ij} is satisfied. Experience has shown that most of the variables are known to be in the right domain (this is always the case in the example in Sec. 2) but for the others separate subgoals must be generated to prove that they are in the correct domain. Examples of such subgoals are given in Sec. 5.

TPI attempts to prove the terminal nodes in the subgoal tree as follows: Each occurrence of a B in the subgoal is replaced by σC where $Q(B' \equiv C)$ is an element of the FD list and σ is the most general unifier of B and B' . The restrictions on the variables specified by Q are checked and appropriate literals are added to the subgoals when the restrictions are not already known to be satisfied. The replacement of defined quantities, as specified by FD, is a recursive process since the definition of a quantity may contain other defined quantities. In addition, some defined quantities cannot be removed and the definitions of such quantities are added to the subgoal. For example, the definition of ordered pair was added to the subgoals of the example of Sec. 2 because the $\circ$ operator could not be removed.

The above process is continued until all pertinent definitions have been incorporated into the subgoal. Then it is converted into clause form which is passed to a resolution-based theorem prover, called the subgoal solver (SGS), for proof. Note that the SGS never refers to FD and BD. Instead TPI, selects, via generating the tree of subgoals, the relevant elements of FD and BD to be incorporated into the terminal subgoals before they are passed to the SGS. This is important because most of FD and BD are not used in proving any one theorem.

It is quite easy to see the conditions under which the TPI is complete. Suppose that for each defined quantity B, there is an entry on both BD and FD and one of the entries $Q_i(A_i \rightarrow B)$ on BD has the property that $Q_i(B \rightarrow A_i)$ is also true. Then some subtree of the AND/OR subgoal tree will be equivalent to the original theorem, in the sense that it is valid if and only if all of the terminal nodes in the subtree are valid. The reason is that each OR node has a subnode which is equivalent to it by the above assumption. Since the elements on FD all have $\equiv$ as a main connective, the clause form of a terminal node which is passed to the SGS, is equivalent to the node.

The above is a brief argument showing that TPI is complete if the SGS is complete and the unification algorithm is complete. Both of these conditions are violated by the current version of TPI. The SGS does not process equality in a complete way -- a difficulty which can easily be removed. The difficulty with the unification algorithm can only be removed by removing all higher-order variables because there is no unification algorithm for higher-order logic.⁴ Although TPI is not complete, its underlying approach (i.e., an AND/OR tree of subgoals) is complete. We feel it is important that we know in what ways TPI is incomplete and what are the difficulties in the way of making TPI complete.

4. Formulation of BD and FD

This section discusses certain pragmatic considerations in the formulating of FD and BD.

Consider modifying the definition of group in Sec. 3 by "anding" right cancellation to (13). Would this make TPI's job easier or harder? If the definition on BD were modified this way TPI would always have to do more work. The reason is that BD is only used when trying to prove that something is a group, i.e., when $Gr(\dots)$ is the consequent of some subgoal. It would be ridiculous to try to prove right cancellation in addition to the 3 conjuncts of (13) since the latter are sufficient. Thus, a statement on BD should be "minimal" in the sense that removing any conjunct from its antecedent causes the statement to become invalid.

Adding right cancellation to the definition of group on FD may be quite useful because this statement is used when TPI is "given" that something is a group. Thus, this modification will allow the SGS of TPI to use right cancellation of the given group in proving a subgoal. In general, for any definition $Q(B \equiv C)$ on FD, a previously proven theorem of the form $Q(B \rightarrow D)$ can be incorporated into FD by using the new definition $Q(B \equiv C \& D)$ on FD. Note that the logical connective of this new definition is still $\equiv$. Of course, if too many previously proven theorems are incorporated into a definition on BD, TPI may become bogged down. But adding a few carefully selected previously proven theorems can be very beneficial.

Another beneficial way to express things is to remove certain existential quantifiers from definitions. Although quantifiers cannot be "moved across" an $\equiv$ (which statements on FD contain) extra parameters can be added to definitions which allow the existential quantifiers to be removed. For example, suppose we defined group as a predicate with two arguments: set and operator. Then the definiens would contain "there exists an identity and an inverse such that..." Making the identity and the inverse arguments to the predicate, as in Sec. 3, removes these two existential quantifiers. This formulation allows one to give TPI identities and inverses in the statement of theorems as in Sec. 2. But more importantly, it allows other definitions to refer to these quantities. For example, the predicate subgroup can be defined by

$$\begin{aligned} \text{Sg}(S, G, +, -, 0) \equiv \text{Gr}(G, +, -, 0) \ \& \ \forall(x \ S)(y \ S) \\ (-x \in S \ \& \ x + y \in S) \end{aligned} \quad (14)$$

If group were a 2-place predicate this definition would not be possible.

5. An Example with More Complex Higher-Order Terms

Although TPI processes higher-order variables incompletely, it can still handle the example, below, which is to prove that the set of homomorphisms from one abelian group to another is a group. For readability quantifiers and domains of quantified variables will be omitted from the formulae below. Rather we will note in the text the domain of variables and constants whenever this information is important.

The theorem is

$$\begin{aligned} \text{Ab}(G, +, -, 0) \ \& \ \text{Ab}(H, *, I, E) \rightarrow \text{Gr}(\text{homo}(G, H, +, *), \\ \alpha, \beta, \gamma) \end{aligned} \quad (15)$$

where

$$\alpha \text{ is } \lambda(uv)(\lambda x(u(x)*v(x))) \quad (16)$$

$$\beta \text{ is } \lambda u(\lambda x(I(u(x)))) \quad (17)$$

$$\gamma \text{ is } \lambda uE \quad (18)$$

In the above, u and v are homomorphisms, $x \in G$, and Ab is the predicate stating its first argument is an abelian group with respect to its other 3 arguments. Similarly Gr is the group predicate (see Section 2).

The definition of homomorphism (on both FD and BD) states

$$f \in \text{homo}(G, H, +, *) \equiv \forall(x \ G)(y \ G)f(x+y)=f(x)*f(y) \quad (19)$$

where $f: G \rightarrow H$.

In (19) we have used for arguments of homo , the same names as used in (15) for readability, but in BD and FD the arguments of homo are universally quantified.

The theorem (15) is eventually broken into the 6 terminal subgoals shown in Fig. 3. Space only permits us to describe the highlights of the generation of these subgoals.

The generation of 2 in Fig. 3 is typical of the generation of the first 3 subgoals. From the definition of group we get the AND subnode whose consequent is

$$\alpha(f, \beta(f)) = \gamma \quad (20)$$

where u is any homomorphism. This just states that β is the inverse, and by " λ -conversion" immediately simplifies to

$$\lambda x(f(x)*I(f(x))) = \gamma \quad (21)$$

Both the left and the right of (21) are functions from G to H ; i.e., their domain are $(\text{FN } G \ H)$. As mentioned in Sec. 3, FN is a primitive built into TPI. To prove that two functions are equal TPI invokes functional extensionality,

$$(\forall x(f(x)=g(x))) \rightarrow f=g, \quad (22)$$

an axiom built into TPI. This produces the subgoal

$$(\lambda x(f(x)*I(f(x))))(a) = \gamma(a) \quad (23)$$

where a is a constant in G . (23) immediately simplifies to $f(a)*I(f(a))=E$.

Subgoals 4, 5, and 6 in Fig. 3 are generated because the terms substituted for the variables in the definition of a group are not known to be in the required domain. We will describe the generation of subgoal 5, which is typical of the other two. The domain

of β is required to be a function from homomorphisms to homomorphisms according to the definition of a group on BD. However, according to (17), TPI only knows that the domain of β is $(\text{FN}(\text{FN } G \ H)(\text{FN } G \ H))$. Thus the subgoal $\beta \in (\text{FN } \text{homo}(G, H, +, *) \text{homo}(G, H, +, *))$ is generated. TPI uses the built-in axiom, $\forall x(x \in A \rightarrow f(x) \in B) \rightarrow f \in (\text{FN } A \ B)$, to generate the subgoal,

$$\beta(f) \in \text{homo}(G, H, +, *), \quad (24)$$

where f is a constant homomorphism. Simplifying (24) with " λ -conversion" we get

$$\lambda x(I(f(x))) \in \text{homo}(G, H, +, *) \quad (25)$$

Replacing homo by its definition and simplifying we get subgoal 5 in Fig. 3.

Currently, TPI cannot prove (15) because the commutativity axioms give rise to too many resolvents at each step. Although this is no problem for the first 4 subgoals in Fig. 3, it prevents the SGS from proving 5 and 6. We have included this example because the emphasis in this work has been on the generation of simple subgoals which TPI does in this example. We believe that existing resolution theorem provers can prove 5 and 6 in Fig. 3 and thus if TPI had a more sophisticated SGS, it could prove (15).

6. Conclusion

Most contemporary theorem provers have difficulty proving a theorem whose proof is relatively long or whose statement is relatively long (i.e., large number of clauses). The goal of our research is to develop a theorem prover which is somewhat insensitive to the length of the proof and the length of the statement of the theorem.

The three main methods in TPI for achieving this goal are: (1) dividing the theorem into independent subgoals, (2) the special treatment of definitions, and (3) implicit use of restricted quantification. The purpose of (1) is to make TPI sensitive to the length of the proof of the "longest" subgoal as opposed to the length of the total theorem. Empirical results indicate that (1) achieves this purpose quite well.

The main purpose of (2) is to disregard those definitions that are not necessary for the proving theorem. Currently, TPI is given 14 definitions and no theorem has used more than 4. Thus, the extraneous definitions are considerably longer than the theorem and the pertinent definitions. The processing of definitions also interfaces very nicely with (1), allowing TPI to subdivide subgoals into independent subgoals.

The purpose of (3) is to give high priority to inferences which prove that variables are in the correct domain. This heuristic allows TPI to discover longer proofs because such inferences are done implicitly. Of course, the price paid for this ability is an increase in the amount of time for each explicit inference.

The example in Sec. 2 shows that TPI has achieved, to some extent, its goal. The statement of the theorem given to TPI (including extraneous definitions) is 48 clauses, and the shortest proof that we know of is 27 binary resolutions.

The methods of TPI form a complete system of first-order logic. Although TPI cannot process higher-order terms in a complete way, its methods are sufficient for some interesting theorems containing higher-terms. Initial results indicate that the use of lambda notation and λ -normal-form are very powerful theorem-proving techniques.

Other theorem provers have used methods which are quite similar to those of TPI. Subdividing a theorem into parts has been used by Bledsoe,¹ Nivens,⁵ Norton,⁶ and Slagle and Koniver,⁷ and probably others. Norton⁶

uses definitions in a way which is very similar to TPI and Bledsoe's¹ use of definitions is somewhat similar. Although the overall characteristics of all these theorem provers (including TPI) are quite different, they all proved some relatively impressive theorems. In particular, each can prove theorems that the others cannot prove. However, this author believes that in each case much of the power stems from the same few basic principles even though the methods of these theorem provers were independently developed along quite different lines.

Acknowledgments

This research was supported by the National Science Foundation under grant GJ-1135, and the Air Force Office of Scientific Research under grant AFOSR-71-2110C.

References

1. Bledsoe, W. W., "Splitting and Reduction Heuristics in Automatic Theorem Proving," Artificial Intelligence 2, 1971, pp. 55-77.
2. Ernst, G. W., "The Utility of Independent Subgoals in Theorem Proving," Information and Control, April, 1971.
3. Ernst, G. W., "A Definition-Driven Theorem Prover," Report No. 1124, Computing and Information Sciences, Case Western Reserve University, 1973.
4. Huet, G. P., "The Undecidability of the Existence of a Unifying Substitution in the Simple Theory of Types," Report 1120, Jennings Computing Center, Case Western Reserve University, 1972.
5. Nevins, A. J., "A Human Oriented Logic for Automatic Theorem Proving," AI Memo No. 268, Mass. Inst. of Tech., Cambridge, Mass., 1972.
6. Norton, L. M., "Adept: A Heuristic Program for Proving Theorems of Group Theory," Report MAC-TR-33, Mass. Inst. of Tech., Cambridge, Mass., 1966.
7. Slagle, J. R., and Koniver, D., "Finding Resolution Graphs and Using Duplicate Goals in AND/OR Trees," Information Sciences, Vol. 3, No. 4, Oct., 197., pp. 315-342.




Figure 2 The General Form of the AND/OR Tree Generated by TPI

1.	x=0=x	0 is identity
2.	0+x=x	
3.	x+(-x)=0	- is inverse
4.	-x+x=0	
5.	y*E=y	E is identity
6.	E*y=y	
7.	y*I(y)=E	I is inverse
8.	I(y)*y=E	
9.	1st(u,v)=u	definition of ordered pair
10.	2nd(u,v)=v	
11.	1st(w)=u & 2nd(w)=v ⇒ v=1st(u,v)=w	
12.	1st(a)+1st(0,E) ≠ 1st(a)	negation of consequent of (8)

Figure 1 Subgoal (8) after negation and reduction to normal form. The domain of x is G and the domain of y is H whereas u, v and w are not restricted.

1.	f(a)*E=f(a)	γ is identity
2.	f(a)*I(f(a))=E	β is inverse
3.	f(a)*(g(a)*h(a))=(f(a)*g(a))*h(a)	α is associative
4.	E*E=E	γ is a homomorphism
5.	I(f(a+b))=I(f(a)*I(f(b)))	closed under β
6.	f(a+b)*g(a+b)=(f(a)*g(a))*(f(b)*g(b))	closed under α

Figure 3 The consequents of 6 terminal suggoals whose conjunction is equivalent to (15). All of the symbols in this figure are constants. f, g and h are homomorphisms while a and b are elements of G.