TLA$^+$ Version 2

A Preliminary Guide

Leslie Lamport

11 December 2018

Contents

1 Introduction

The original version of the TLA$^+$ language was released at the beginning of the millennium and described in the book [Specifying Systems](#), published in 2002. The current version, released around 2006, is Version 2. It is the version supported by the latest versions of the TLA$^+$ tools and described in documentation written since then, including [the video course](#) and [the hyperbook](#). TLA$^+$ now means TLA$^+$ Version 2. In this document, it is called TLA^{+2} for short. The original version of TLA$^+$ is here called TLA^{+1}. This document explains the differences between TLA^{+2} and TLA^{+1}.

Most of the additions to the language in TLA^{+2} are for writing proofs that can be checked with [TLAPS](#), the TLA$^+$ proof system. The major change that affects specifications is that you can now write recursive operator definitions, as described in Section 2. Another change is the introduction of *lambda* expressions, explained in Section 3.

Almost all legal TLA^{+1} specifications are legal TLA^{+2} specifications. Two rather arcane changes have been made to instantiation; they are explained in Section 5. The only other change that affects TLA^{+1} specifications is that the following new keywords have been added in TLA^{+2}, and thus cannot be used as identifiers.

ACTION	HAVE	PICK	SUFFICES
ASSUMPTION	HIDE	PROOF	TAKE
AXIOM	LAMBDA	PROPOSITION	TEMPORAL
BY	LEMMA	PROVE	USE
COROLLARY	NEW	QED	WITNESS
DEF	OBVIOUS	RECURSIVE	
DEFINE	OMITTED	STATE	
DEFS	ONLY		

2 Recursive Operator Definitions

The only recursive definitions allowed in TLA^{+1} were recursive function definitions. This restriction was inconvenient for the following reasons: (i) specifying the function's domain was sometimes difficult, (ii) checking that the function was applied to an element in the domain could significantly slow down TLC, and (iii) there was no provision for mutual recursion. I did not allow recursive operator definitions in TLA^{+1} because I didn't know how to assign a sensible meaning to them—for example, what should be the

meaning of this silly definition?

$$F \;\triangleq\; \text{CHOOSE } v \,:\, v \neq F$$

Georges Gonthier and I have figured out how to define recursive operator definitions so they have the expected meaning when you expect them to be meaningful—namely, when the value can be computed by expanding the definition a finite number of times. The precise definition is complicated and I hope will eventually appear elsewhere.

In TLA^{+2}, the use of a defined operator must come after either its definition or its declaration by a RECURSIVE statement. For example,

RECURSIVE $fact(\_)$
$fact(n) \;\triangleq\;$ IF $n = 0$ THEN 1 ELSE $\; n * fact(n-1)$

defines $fact(n)$ to equal $n!$ if n is a natural number. I have no idea what it defines $fact(-2)$ or $fact(\text{"abc"})$ to equal. (Without the RECURSIVE declaration, $fact$ could be used only after its definition, so its use in the right-hand side of the definition would be illegal.)

The syntax of the RECURSIVE statement is the same as that of the CONSTANT statement, allowing multiple declarations separated by commas. The RECURSIVE statement can come anywhere before the first use of the operators it declares, so it's easy to write mutually recursive definitions. However, you should put a RECURSIVE statement as close as possible to the definitions of the operators it declares. A tool might treat as recursive any definitions that come between an operator's RECURSIVE declaration and its definition.

A RECURSIVE statement can be used in a LET expression to permit recursive definitions local to the LET. A symbol declared in a RECURSIVE statement must later be defined to be an operator taking the correct number of arguments. Thus, recursive instantiations are *not* allowed; you cannot write

RECURSIVE $Ins(\_)$
$Ins(n) \;\triangleq\;$ INSTANCE M WITH $\ldots$

TLA^{+1} has the nice property that operator definitions are like macros. If F is defined by

$$F(x) \;\triangleq\; \ldots$$

then $F(exp)$ is simply the expression obtained from the right-hand side of the definition by replacing every instance of x with exp. In TLA^{+2}, this

2

is not true for recursively-defined operators. We do not know if $fact(-2)$ equals

$$\text{IF } n = 0 \text{ THEN } 1 \text{ ELSE } n * fact(-3)$$

It can be proved that

$$fact(42) \triangleq \text{ IF } n = 0 \text{ THEN } 1 \text{ ELSE } n * fact(41)$$

However, because *fact* is defined recursively, this must be proved. The method of proving it is fairly standard; I won't discuss it here.

3   Lambda Expressions

TLA$^+$ allows you to define higher-order operators—that is, ones that take operators as arguments, such as

$$F(Op(\_, \_)) \triangleq Op(1, 2)$$

The argument of F is an operator that takes two arguments. In TLA^{+1}, such an argument had to be the name of an operator. For example, we might define

$$Id(a, b) \triangleq a + 2 * b$$

and write $F(Id)$. TLA^{+2} allows you to use F without having to define an operator to use as its argument. Instead of defining Id in this way and writing $F(Id)$, you can write

$$F(\text{LAMBDA } a, b : a + 2 * b)$$

The LAMBDA expression is the operator that Id is defined to equal.

A LAMBDA expression can also be used in an *instance* statement to instantiate an operator parameter. For example, with the definition of Id given above, the following two statements are equivalent.

$$\text{INSTANCE } M \text{ WITH } Op \leftarrow Id$$
$$\text{INSTANCE } M \text{ WITH } Op \leftarrow \text{LAMBDA } a, b : a + 2 * b$$

Syntactically, a LAMBDA expression consists of the keyword LAMBDA followed by a comma-separated list of identifiers, followed by ":", followed by an expression. A LAMBDA expression can be used only as the argument of a higher-order operator or to the right of a "$\leftarrow$" in an *instance* statement.

4 Theorems and Assumptions

4.1 Naming

There is no need for theorem or assumption names in a specification, since the name would be equivalent to TRUE. However, theorem and assumption names are used in writing proofs. In TLA^{+2}, you can name a theorem or assumption by inserting an optional "*identifier* $\triangleq$ " right after THEOREM or ASSUME, as in

$$\text{THEOREM } Fermat \;\triangleq\; \neg\exists\, n \in Nat \setminus (0 \,..\, 2) : \ldots$$

This is equivalent to

$$Fermat \;\triangleq\; \neg\exists\, n \in Nat \setminus (0 \,..\, 2) : \ldots$$
$$\text{THEOREM } Fermat$$

A theorem cannot have parameters.

TLA^{+2} allows LEMMA and PROPOSITION as synonyms for THEOREM and ASSUMPTION as a synonym for ASSUME. TLA^{+2} also allows AXIOM as almost a synonym for ASSUME and ASSUMPTION; it differs only in that (in Toolbox releases later than Version 1.1.2) TLC does not check assumptions labeled AXIOM. This is useful when writing assumptions that TLC can't check, for use in a proof.

4.2 ASSUME/PROVE

In TLA^{+2}, a theorem can assert either a formula or an ASSUME/PROVE. A formula is a Boolean-valued expression. However, since TLA^{+} is untyped, the silly statement "THEOREM 42" is a legal (but unprovable) theorem. (See Section 16.1.3 of *Specifying Systems*.)

An ASSUME/PROVE asserts a proof rule. Here is how it is used to assert a well-known rule of elementary logic that we can prove $P \Rightarrow Q$ by assuming P and proving Q.

$$
\begin{array}{ll}
\text{THEOREM } DeductionRule \;\triangleq\; & \text{ASSUME} \quad \text{NEW } P, \; \text{NEW } Q, \\
& \qquad \text{ASSUME} \quad P \\
& \qquad \text{PROVE} \quad\;\, Q \\
& \text{PROVE} \quad\; P \Rightarrow Q
\end{array}
$$

Logicians often use "$\vdash$" to express such a rule, writing this as $(P \vdash Q) \vdash (P \Rightarrow Q)$. In TLA^{+}, we need to declare identifiers like P and Q before they can be used. Here is a standard proof rule of predicate logic; it asserts that

we can prove $\forall\, x \in S \,:\, P(x)$ by choosing a brand-new identifier x, assuming $x \in S$, and proving $P(x)$.

> THEOREM ASSUME   NEW $P(\_)$,   NEW S,
>   ASSUME   NEW $x \in S$
>   PROVE $P(x)$
>   PROVE $\forall x \in S \,:\, P(x)$

The third assumption of this rule,

> ASSUME NEW $x \in S$ PROVE $P(x)$

is an abbreviation for

> ASSUME   NEW x, $x \in S$
> PROVE $P(x)$

Here are a couple of proof rules of TLA. The first asserts that a primed constant equals itself.

> THEOREM $Constancy$ $\;\stackrel{\Delta}{=}\;$ ASSUME   CONSTANT C
>   PROVE $C' = C$

Here is a standard temporal-logic rule:

> THEOREM ASSUME   TEMPORAL F,   TEMPORAL G
>   PROVE $\Box(F \wedge G) \;\equiv\; \Box F \wedge \Box G$

These theorems assert a rule that is valid whenever expressions or operators of the specified (or lower) level are substituted for the declared identifiers. For example, Theorem $Constancy$ implies $(2+N)' = (2+N)$ if N is declared to be a constant parameter of the module. See Section 17.2 of *Specifying Systems* for an explanation of levels. (The *action* level is called transition-level there.)

The declaration NEW is equivalent to CONSTANT. If all the expressions and identifiers that appear in a theorem have constant level, then the theorem is valid when expressions of any level are substituted for the declared identifiers.

You can also use a VARIABLE declaration in an ASSUME to state that some identifier is a TLA$^+$ variable. To illustrate the difference between a *variable* and a *state* declaration, consider this valid TLA$^+$ rule.

> THEOREM ASSUME   VARIABLE x,   VARIABLE y
>   PROVE   ENABLED $x' \neq y'$

The theorem would not be valid if "VARIABLE" were replaced by "STATE" because the resulting theorem would allow any state-level expressions to be substituted for x and y. Substituting the variable z for both x and y would then yield the conclusion ENABLED $z' \neq z'$, which is false.

You have probably inferred most of the grammar of ASSUME/PROVE assertions:

- An ASSUME/PROVE consists of the token ASSUME, followed by a comma-separated list of *assumptions*, followed by the token PROVE, followed by an expression.

- An *assumption* is an expression, a *declaration*, or an ASSUME/PROVE.

- A *declaration* may be:

    - The same as a CONSTANT statement in the body of the module that declares a single constant parameter, except that the keyword CONSTANT may optionally be replaced by NEW, STATE, ACTION, or TEMPORAL.
    - The token NEW or CONSTANT, followed by an identifier, the token $\in$, and an expression.
    - The token VARIABLE followed by an identifier.

    An optional NEW token may precede any of these declarations except for one beginning with a NEW token. (The unnecessary "NEW" may help some people understand the meaning of the declaration.)

Indentation is not significant. (In TLA^{+2} as in TLA^{+1}, indentation matters only in bulleted lists of conjuncts and disjuncts.)

The meaning of ASSUME/PROVE assertions is subtle when they contain temporal or action-level formulas. See Section 8.3 on page 35 for an explanation.

5 Instantiation

Two minor changes to instantiation have been made in TLA^{+2}: (i) there is a different syntax for instantiated in-, pre-, and postfix operators, and (ii) operator instantiation has been restricted to allow instantiation only with "Leibniz" operators, which are defined below.

5.1 Instantiating *fix Operators

If module M defines an infix operator such as &&, then in TLA^{+1} the statement

$$Foo \;\triangleq\; \text{INSTANCE } M \text{ WITH } \ldots$$

defines an infix operator Foo!&& that would be used in such strange expressions as

$$1\ Foo!\&\&\ 2$$

TLA^{+2} eliminates this awkward syntax. Instead, the operator Foo!&& is a "normal" *nonfix* operator and not an infix one, so you write this expression as Foo!&&$(1, 2)$. If this were a parameterized instantiation, so Foo took an argument, then you would write something like $Foo(42)$!&&$(1, 2)$.

The analogous change has been made to postfix operators and the prefix operator unary "$-$", which must be written as "`-.`" after a "`!`".

For the sake of uniformity, TLA^{+2} permits any infix or postfix operator to be used as a nonfix operator. For example, `+(1,2)` is another way of writing `1+2`. (Prefix operators could always be written this way.) This alternate syntax does not apply to the left-hand side of a definition. For example, the only way to define the infix operator && is to write something like

$$a\ \&\&\ b\;\triangleq\;\ldots$$

Because of a bug that is unlikely to be fixed, the current SANY parser does not accept this alternate syntax for the infix operator "$-$"; it accepts only `2-1` and not `-(2,1)`.

5.2 Leibniz Operators and Instantiation

Consider the following module.

```
┌──────────────── MODULE M ────────────────┐
│ CONSTANTS C, D, F(_)                      │
│ THEOREM (C = D)  ⇒  (F(C) = F(D))         │
│                                           │
└───────────────────────────────────────────┘
```

The perfectly reasonable theorem in this module is not valid in TLA^{+1} for the following reason. The semantics of TLA^{+} requires that any instantiation of a valid theorem be valid. Now consider

7

$$\text{VARIABLES } x,\ y$$
$$Prime(p) \;\triangleq\; p'$$
$$\text{INSTANCE } M \text{ WITH } C \leftarrow x,\ D \leftarrow y,\ F \leftarrow Prime$$

This imports the theorem from module M as

$$\text{THEOREM } (x = y) \Rightarrow (x' = y')$$

which is not valid. (Equality of the values of x and y in the current state doesn't imply that they are equal in the next state.)

In TLA^{+2}, the theorem of module M is valid, which means that this INSTANCE M statement is illegal. It is illegal because TLA^{+2} allows instantiation of an operator parameter only by a Leibniz operator, and F is non-Leibniz. An operator F of a single argument is defined to be *Leibniz* iff $e = f$ implies $F(e) = F(f)$, for any expressions e and f. (Logicians generally use the term *substitutive* rather than *Leibniz*.) For an operator F that takes k arguments, F is *Leibniz* iff the value of $F(e_1, \ldots, e_k)$ remains unchanged if any of the expressions e_i is replaced by an equal expression. Constant parameters are assumed to be Leibniz, so one constant parameter can be instantiated by another.

In TLA$^+$, all built-in and definable constant operators are Leibniz. The only built-in TLA$^+$ operators that are not Leibniz are the action operators and the temporal operators, listed in Tables 3 and 4 of *Specifying Systems*. In a non-constant module, a constant parameter can be instantiated only by a constant operator. Thus, the restriction added in TLA^{+2} is automatically satisfied except when substituting non-constant operators in a constant module. However, a non-constant operator can be Leibniz—for example, the Leibniz operator G defined by

$$G(a) \;\triangleq\; x' = [x \text{ EXCEPT } ![a] = y']$$

For a defined operator to be non-Leibniz, one of its parameters must appear in the definition within an argument of a non-Leibniz operator like $'$ (prime).

6 Naming Subexpressions

When writing proofs, it is often necessary to refer to subexpressions of a formula. In theory, one could use definitions to name all these subexpressions. For example, if

$$Foo(y) \;\triangleq\; (x + y) + z$$

and we need to mention the subexpression $(x + 13)$ of $Foo(13)$, we could write

$$Newname(y) \triangleq (x + y)$$
$$Foo(y) \triangleq NewName(y) + z$$

This doesn't work in practice because it results in a mass of non-locally defined names, and because we may not know which subformulas need to be mentioned when we define the formula.

TLA^{+2} provides a method of naming subexpressions of a definition. If F is defined by $F(a, b) \triangleq \ldots$, then any subexpression of the formula obtained by substituting expressions A for a and B for b in the right-hand side of this definition has a name beginning "$F(A, B)$!". (Although this is a new use of the symbol "!", it is a natural extension of its use with module instantiation.)

You can use subexpression names in any expression. When writing a specification, you can define operators in terms of subexpressions of the definitions of other operators. Don't! Subexpression names should be used only in proofs. In a specification, you should use definitions to give names to the subexpressions that you want to re-use in this way.

6.1   Labels and Labeled Subexpression Names

Any subexpression of a definition can be labeled. The syntax of a labeled expression is

$label$ :: $expression$

(The symbol "::" is typed "$::$".) The label applies to the largest possible expression that follows it. In other words, the end of the labeled expression is the same as the end of the expression that you would get by replacing the "$label$ ::" with "$\forall x$:". However, the expression is illegal if removing the label would change the way the expression is parsed. For example,

$a + lab :: b * c$

is legal because it is parsed as $a + (lab :: (b * c))$, which is how it would be parsed if the label lab were not there. However,

$a * lab :: b + c$

is illegal because it would be parsed as $a * (lab :: (b + c))$ and removing the label causes the expression to be parsed as $(a * b) + c$.

Label parameters are required if labels occur within the scope of bound identifiers. Here is an example.

$$F(a) \;\triangleq\; \forall\, b\,:\, l1(b) :: (a > 0) \Rightarrow$$
$$\wedge \ldots$$
$$\wedge\, l2 :: \exists\, c\,:\, \wedge \ldots$$
$$\wedge\, \exists\, d\,:\, l3(c,d) :: a - b > c - d$$

For this example, $F(A)!l1(B)!l2!l3(C,D)$ names the expression $A - B > C - D$. Note how the parameters of each label are the bound identifiers introduced between it and the next outer-most label. Those identifiers can appear in any order. For example, if the label $l3(c,d)$ were replaced by $l3(d,c)$, then $F(A)!l1(B)!l2!l3(C,D)$ would name the expression $A - B > D - C$.

In this example, a reference to the subexpression labeled by $l3(c,d)$ from outside the definition of F, must specify the values of all the bound identifiers a, b, c, and d. That's why labels must include the bound identifiers as parameters. Also observe that to name a labeled subexpression, we have to name all the labeled subexpressions within which it lies. We're not even allowed to eliminate the label $l2$, even though it is superfluous in this example.

Label names do not conflict with operator names. In this example, any one of the label names $l1$, $l2$, or $l3$ could be replaced by F. The rule for name conflict is the obvious one needed to guarantee that there's no ambiguity in a subexpression name (where we are not allowed to use the number of parameters to disambiguate). Thus, we cannot label the first conjunct of the $\exists\, c$ expression with $l3(c)$, but we could label it with $l1(c)$ or $l2(c)$.

For subexpressions of the definition of an infix, postfix, or prefix operator, we use the "nonfix" form. For example, a subexpression of the definition of $\&\&$ would have the form $\&\&(A,B)\,!\ldots$.

We can also name subexpressions of definitions in instantiated modules. For example, if we have

$$Ins(x) \;\triangleq\; \text{INSTANCE } M \text{ WITH } \ldots$$

and ν is the name of any subexpression of a definition in module M, then $Ins(exp)!\nu$ is the name of the subexpression of the instantiated definition obtained when exp is substituted for x.

We call a subexpression name having one of the forms described here a *labeled subexpression name*. We include in this category the trivial case in which there is no label name, only the name of a defined operator—possibly

in an instantiated module. The precise definition is contained in the "fine print" below. You probably don't want to read it.

The Fine Print

Here is the general definition explained above with examples. We say that label $lab1$ is the *containing label* of $lab2$ iff (i) $lab2$ lies within the expression labeled by $lab1$ and (ii) if $lab2$ lies within the expression labeled by any other label, then $lab1$ also lies within that expression.

We use the notation that $f(e_1, \ldots, e_k)$ denotes f when $k = 0$. A label lab has the form $id(p_1, \ldots, p_k)$ where id and the p_i are identifiers, the p_i are all distinct, and $\{p_1, \ldots, p_k\}$ is the set of all bound identifiers p_i such that:

- Label lab lies within the scope of p_i.

- If lab has a containing label $labc$, then the expression that introduces p_i lies within the expression labeled by $labc$.

We call id the *name* of the label. Two labels that either have no containing label or have the same containing label must have different names.

A *simple labeled subexpression name* of a module M has the form $prefix \, ! \, labexp_1 \, ! \ldots ! \, labexp_n$, where $prefix$ has the form $Op(e_1, \ldots, e_{k[0]})$, each $labexp_i$ has the form $id_i(e_1, \ldots, e_{k[i]})$, Op and the id_i are identifiers, and the e_j are expressions. It must satisfy:

- The definition

$$Op(p_1, \ldots, p_{k[0]}) \;\triangleq\; \ldots$$

   occurs at the top level (not inside a LET or inner module) of M.

- id_1 must be the name of a label lab_1 in the definition of Op that has no containing label.

- If $i > 1$, then id_i must be the identifier of a label lab_i whose containing label is lab_{i-1}.

- $k[i]$ must equal the number of parameters in lab_i, for each $i > 0$.

This labeled subexpression name denotes the expression obtained from the expression labeled with lab_n by substituting for each parameter of Op and of each lab_i the corresponding argument of $prefix$ and $labexp_i$, respectively.

A *labeled subexpression name* of a module M is either a simple labeled subexpression name of M or else has the form $Id(e_1, \ldots, e_k) \, ! \, \lambda$ where there is a statement

$$Id(e_1, \ldots, e_k) \;\triangleq\; \text{INSTANCE } N \ldots$$

at the outermost level of M and λ is a labeled subexpression name of module N.

6.2 Positional Subexpression Names

Instead of using labels, we can name subexpressions of a definition by a sequence of *positional selectors* that indicate the position of the subexpression in the parse tree. Consider this example

$$F(a) \;\triangleq\; \wedge \ldots$$
$$\wedge \ldots$$
$$\wedge \; Len(x[a]) > 0$$
$$\wedge \ldots$$

Here are how some of the subexpressions of this definition are named, where A is an arbitrary expression:

- $F(A)!3$ names $Len(x[A]) > 0$, the third conjunct of $F(A)$—that is, of the right-hand side of the definition with A substituted for a. We think of this conjunct list as the application of a conjunction operator that takes four arguments, the third being $Len(x[A]) > 0$.

- $F(A)!3!1$ names $Len(x[A])$, the first argument of $>$, the top-level operator of the expression $F(A)!3$

- $F(A)!3!1!1$ names $x[A]$, the first (and only) argument of the top-level operator of the expression $F(A)!3!1$.

- The naming of subexpressions of $x[A]$ is based on the realization that this expression represents the application of a function-application operator to the two arguments x and A. Thus, $F(A)!3!1!1!1$ names x and $F(A)!3!1!1!2$ names A

The positional selector "!$\langle$" is always synonymous with !1, and "!$\rangle$" is synonymous with !2 when selecting the second argument of an operator that takes two arguments. Thus, instead of $F(A)!3!1!1!2$, we could write $F(A)!3!\langle!\langle!\rangle$ or $F(A)!3!\langle!1!\rangle$ or $F(A)!3!1!\langle!2$ or .... As usual, "$\langle$" is typed "<<" and "$\rangle$" is typed ">>".

The use of positional selectors to pick an argument of an operator is self-evident for most operators that do not introduce bound identifiers. Here are the cases that are not obvious.

- In $\;\;[f \text{ EXCEPT } ![a] = g, ![b].c = h]\;\;$ we select f with !1, g with !2, and h with !3. No other subexpressions of the EXCEPT construct can be named.

- $r.fld$ is an application of a record-field selector operator to the two arguments r and "fld", so !1 selects r. (You can also use !2 to select "fld", but there's no reason to name a simple string constant with a subexpression name.)

- In $[fld_1 \mapsto val_1, \ldots, fld_n \mapsto val_n]$ and $[fld_1 : val_1, \ldots, fld_n : val_n]$ the selector $!\,i$ names the subexpression val_i for $i \in 1 \,.. \, n$. The field names fld_i cannot be selected. (There is no point naming fld_i, since it's just a string constant.)

- In IF p THEN e ELSE f the selector !1 names p, the selector !2 names e, and the selector !3 names f.

- In CASE $p_1 \to e_1 \,\square\, \ldots \,\square\, p_n \to e_n$ the selector $!i!1$ names p_i and $!i!2$ names e_i. If p_n is the token OTHER, then it cannot be named.

- In $\mathrm{WF}_e(A)$ and $\mathrm{SF}_e(A)$ the selector !1 names e and !2 names A.

- In $[A]_e$ and $\langle A \rangle_e$ the selector !1 names A and !2 names e.

- In LET ... IN e the selector !1 names e. This is rather subtle because we are naming an expression that contains operators defined in the LET clause that are not defined in the context in which the subexpression name appears. Consider this example

$$
\begin{aligned}
F &\triangleq \text{ LET } G \triangleq 1 \text{ IN } G + 1 \\
G &\triangleq 22 \\
H &\triangleq F!1
\end{aligned}
$$

The $F!1$ in the definition of H names the expression $G + 1$ in which G has the meaning it acquires in the LET definition. Thus, H is equal to 2, not to 23.

We will see below how to name subexpressions of LET definitions, such as the first (local) definition of G above.

I now describe selectors for subexpressions of constructs that introduce bound identifiers. Consider this example:

$$
R \triangleq \exists x \in S, \, y \in T : x + y > 2
$$

- $R!(X, Y)$ names $X + Y > 2$, for any expressions X and Y.

- $R!1$ names S.

- $R!2$ names T.

In general, for any construct that introduces bound identifiers:

- $!(e_1, \ldots, e_n)$ selects the body (the expression in which the bound identifiers may appear) with each expression e_i substituted for the i^{th} bound identifier.

- If the bound identifiers are given a range by an expression of the form " $\in S$", then $!i$ selects the i^{th} such range S.

For example, in the expression

$$[x, y \in S,\ z \in T \mapsto x + y + z]$$

the selector $!1$ names S, the selector $!2$ names T, and the selector $!(X, Y, Z)$ names $X + Y + Z$.

Parentheses are "invisible" with respect to naming. For example, it doesn't matter if ν names the subexpression $a + b$ or the subexpression $((a + b))$; in either case, $\nu!\langle$ names a.

We usually don't need to name the entire expression to the right of a "$\triangleq$" because the operator being defined names it. However, as observed in Section 2, this is not true for recursively defined operators. If Op is recursively defined by

$$Op(p_1, \ldots, p_k) \;\;\triangleq\;\; exp$$

then "$Op(P_1, \ldots, P_k)!:$" names exp with P_i substituted for p_i, for each i in $1 \ldots k$.

A *positional subexpression name* consists of a labeled subexpression name (defined in Section 6.1 above) followed by a sequence of positional selectors. For example, in

$$F(c) \;\;\triangleq\;\; a \;*\; lab :: (b + c * d)$$

$F(7)!lab!\rangle$ names $7 * d$. Remember that a labeled subexpression need not contain labels—for example, $F(7)$ is a labeled subexpression name.

6.3 Subexpressions of LET Definitions

If a positional subexpression name ν names a LET/IN expression and Op is an operator defined in the LET clause, then $\nu! \, Op(e_1, \ldots, e_n)$ is the name of the expression $Op(e_1, \ldots, e_n)$ interpreted in the context determined by ν. For example, in

$$F(a) \;\triangleq\; \wedge \ldots$$
$$\wedge \text{ LET } G(b) \;\triangleq\; a + b$$
$$\text{IN} \;\; \ldots$$

$F(A)!2!\,G(B)$ names the expression $G(B)$, where the definition of G is interpreted in a context in which A is substituted for a. This expression of course equals $A + B$. (However, if G were recursively defined, $F(A)!2!\,G(B)$ might not be so simply related to the expression to the right of the "$\triangleq$" in G's definition.) We can also name subexpressions of the definition of G. For example, $F(A)!2!\,G(B)!\rangle$ names B. The naming process can be continued all the way down, naming subexpressions of LET definitions contained within LET definitions contained within $\ldots$.

If the LET/IN expression is labeled, then it can be named by a labeled subexpression name λ. In that case, $\lambda! \, Op(e_1, \ldots, e_n)$ is a labeled subexpression name that names a subexpression of the *in* clause with label $Op(p_1, \ldots, p_n)$. To refer to the operator Op defined in the LET clause, just add a "!:" to the end of λ, writing $\lambda!:!\,Op(e_1, \ldots, e_n)$. In particular, if H is defined to equal the LET/IN expression, then we write $H!:!\,Op(e_1, \ldots, e_n)$, even if H is not recursively defined.

6.4 Subexpressions of an ASSUME/PROVE

If we have

$$\text{THEOREM } Id \;\triangleq\; \text{ASSUME } A_1, \ldots, A_n \text{ PROVE } G$$

then Id is not an expression and cannot be used as one. Subexpressions of an ASSUME/PROVE can be named with labels or positionally, where $Id!i$ names A_i if $1 \leq i \leq n$, and $Id!n{+}1$ names G. However, the assumptions can contain declarations like NEW C, so it is possible to name a subexpression of an ASSUME/PROVE that contains identifiers declared within the ASSUME/PROVE. Such a name can be used only within the scope of those

declarations. For example, consider

$$\text{THEOREM } T \ \triangleq \ \text{ASSUME} \quad x > 0, \text{ NEW } C \in Nat, \ y > C$$
$$\phantom{\text{THEOREM } T \ \triangleq \ } \text{PROVE} \quad x + y > C$$

$$\vdots$$
$$Foo \ \triangleq \ \dots$$

Then $T!1$ names the expression $x > 0$, which can be used in the definition of *Foo*. However, $T!3$ names the expression $y > C$ that contains the constant C, and the definition *Foo* is not within the scope of the declaration of C, so $T!3$ cannot be used within the definition of *Foo*. In fact, $T!3$ can be used only within the proof of T. (Proofs are discussed in Section 7.)

6.5 Using Subexpression Names as Operators

Subexpression names can be used as operator names by replacing every part of the form $!id(e_1, \dots, e_n)$ by $!id$, and every selector $!(e_1, \dots, e_n)$ by $!@$. For example, consider:

$$F(Op(\_,\_,\_)) \ \triangleq \ Op(1, 2, 3)$$
$$G \ \triangleq \ \forall x : P \subseteq \{\langle x, \ y{+}z \rangle : y \in S, \ z \in T\}$$

Then $G!(X)!\rangle!(Y, Z)$ is the expression $\langle X, \ Y{+}Z \rangle$, so $G!@!\rangle!@$ is the operator

$$\text{LAMBDA } x, \ y, \ z : \langle x, \ y{+}z \rangle$$

and $F(G!@!\rangle!@)$ equals $\langle 1, \ 2{+}3 \rangle$.

7   The Proof Syntax

This section describes the syntax of proofs and how proofs are checked by TLAPS, the TLA$^+$ proof checker.

7.1 The structure of a proof

A theorem is optionally followed by a proof. A proof is either a terminal proof or a sequence of steps, some of which have proofs. Figure 1 shows a possible proof structure, where the actual assertions made by the steps or by the terminal proofs are elided. This example is a proof having level number 1 and consisting of three steps named $\langle 1 \rangle 1$, $\langle 1 \rangle 2$, and $\langle 1 \rangle 3$. Step $\langle 1 \rangle 1$ has a level-2 proof that consists of three steps, one named $\langle 2 \rangle 4a$, an unnamed step

```
⟨1⟩1.  . . .
        PROOF
            ⟨2⟩4a. . . .
                        OBVIOUS
            ⟨2⟩      . . .
                        ⟨17⟩    . . .
                                    PROOF OMITTED
                        ⟨17⟩1.  . . .
                        ⟨17⟩    . . .
                        ⟨17⟩ab  QED
            ⟨2⟩11   QED
                        BY . . .
⟨1⟩2.  . . .
        BY . . .
⟨1⟩3.   QED
```

Figure 1: The structure of a simple proof.

(marked by the token "⟨2⟩"), and a QED step named ⟨2⟩11. Step ⟨2⟩4a has a terminal proof. The unnamed level-2 proof step has a four-step proof with level number 17. Only its first step has a proof—a terminal proof asserting that the actual proof is omitted.

A proof may optionally begin with the token PROOF. Thus, the PROOF token that begins the proof of step ⟨1⟩1 and that precedes the token OMITTED could be removed, and a PROOF token could be added before step ⟨1⟩1, before the "OBVIOUS" terminal proof, before the first level ⟨17⟩ step, and before either of the BY proofs. The formatting is for readability only; indentation has no significance.

In general, a proof consists of the optional keyword PROOF followed by either a terminal proof or else by a sequence of steps followed by a QED step. A step or a QED step may have a proof, which is called a *subproof* of the proof containing the step. A terminal proof consists of the keyword OBVIOUS or OMITTED or else begins with the keyword BY.

Each step begins with a *step-starting token* that consists of a *step name* followed by an optional sequence of periods. A step name consists of

- < (printed as "⟨")

- a number called the step's *level number* or a + or ∗ character. (The meaning of + and ∗ is explained below.)

17

- **>** (printed as "⟩")

- an optional string of letters and/or digits. If this string is present, then the step is said to be *named* and its *step name* consists of the entire token up to and including this string.

Since a step-starting token is a single token, it may not contain spaces. (Note that a step-starting token is the one place in which "⟨" and "⟩" are typed "<" and ">" rather than "<<" and ">>".) All the steps of a proof have the same level number, which is less than that of any of its subproofs. A step with a greater level number than the preceding step begins the proof of that preceding step, whether or not it is preceded by a PROOF token.

Named steps are referred to by their step names. The scope of a level k step name (the part of a proof within which it can be used) consists of the step's proof (if it has one), all the level-k steps in the same proof that follow it and in those steps' proofs. A step name cannot be used within its scope to label another step. However, the same step name can be used in different subproofs of a proof. For example, step names $\langle 2 \rangle 4a$ and $\langle 17 \rangle 1$ could be used in a proof of step $\langle 1 \rangle 3$.

The level number of a step may be written implicitly with a "$*$" or a "$+$". To explain the meaning of such a level number, let us define the *current level* at a proof step to equal -1 for the first step of the entire proof, and otherwise to equal the level of the latest preceding step that is neither a QED step nor followed by a QED step of the same level. In the example above, the current level at step $\langle 1 \rangle 1$ is -1, the current level at step $\langle 2 \rangle 4a$ is 1, and the current level at step $\langle 2 \rangle 11$ is 2. Let L be the current level at a step whose step-starting token begins with "$\langle * \rangle$" or "$\langle + \rangle$". Then

- a "$+$" is equivalent to the number $L + 1$, and

- a "$*$" is equivalent to the number $L + 1$ if it immediately follows a PROOF token or is at the beginning of the entire proof; otherwise it is equivalent to the number L.

In the above example, $\langle 1 \rangle 1$ can be replaced by either $\langle + \rangle 1$ or $\langle * \rangle 1$; $\langle 2 \rangle 4a$ can be replaced by $\langle + \rangle 4a$ or $\langle * \rangle 4a$; and either of the other two "$\langle 2 \rangle \ldots$" tokens could be replaced by "$\langle * \rangle \ldots$". If the PROOF token before it were missing, then $\langle 2 \rangle 4a$ could be replaced only by $\langle + \rangle 4a$ and not by $\langle * \rangle 4a$. In all cases, it makes no difference if we use the "$*$" or "$+$" or the equivalent explicit level number.

A "$*$" can also be used instead of a level number in a reference to a proof step, in which case it stands for the current level. For example, you

can write $\langle * \rangle 4a$ instead of $\langle 2 \rangle 4a$ in the *by* statement that is the proof of step $\langle 2 \rangle 11$. Again, it makes no difference if you write "$*$" or the equivalent explicit level number.

> TLAPS (the TLA$^+$ prover) does not yet support references to proof steps that use "$*$" as level number.

7.2   USE, HIDE, and BY

7.2.1   USE and HIDE

At any point in a module, there is a set of *current declarations*, a set of *current definitions*, and a set of *known facts*. Outside a proof, the current declarations come from CONSTANT or VARIABLE declarations within the module and within modules it extends; the current definitions come from definitions within the module and within extended or instantiated modules; and the facts come from assumptions and theorems asserted thus far in the module and in extended modules, and from assertions imported thus far by instantiation. Each theorem in an instantiated module yields the assertion that the instantiated theorem follows from the instantiation of the module's assumptions. For example, if module M contains the single assumption

> ASSUME A

and the theorem

> THEOREM $Thm \triangleq T$

then the statement

> $Mod \triangleq$ INSTANCE M WITH ...

imports a theorem named $Mod!Thm$ that asserts

> ASSUME $\overline{A}$
> PROVE $\overline{T}$

where $\overline{A}$ and $\overline{T}$ are the formulas obtained from A and T by performing the substitutions specified by the INSTANCE statement's WITH clause.

There are also subsets of the sets of current definitions and known facts called the *usable definitions* and the *usable facts*. These are the definitions that TLAPS expands and the facts that it tries to apply when trying to prove something. (The definitions referred to here are "outer-level" definitions and not LET definitions, which are always expandable.) Here are the default values of these subsets at points in a module outside a proof.

19

- Only the definitions of theorem names are usable. (Section 4.1 explains how theorems are named.)

- No theorems or assumptions are usable.

The defaults can be overridden by *use* and *hide* statements. Such statements can appear anywhere in the body of the module—that is, at the "top level", not inside any other statements. A *use* or *hide* statement consists of the keyword USE or HIDE followed by an optional list of facts, optionally followed by the keyword DEF or DEFS and a list of definition specifiers. (It must include at least one fact or definition specifier.)

A fact is one of the following:

- The name of a theorem, assumption, or proof step.

- An arbitrary formula—but only in a USE statement, not a HIDE statement. The formula must be easily provable from the currently usable facts and the preceding facts in the USE statement. "Easily provable" means that a proof tool should be able to find the proof without any help from the user.

The parser also allows the following two kinds of "facts" in a USE or HIDE statement. However, they are not supported by TLAPS and are likely to be removed from the language.

- MODULE *Name*, indicating that all known facts obtained from the module *Name* are to be added or removed from the set of usable facts. The module name must appear in an EXTENDS or INSTANCE statement or else be the name of the current module.

- An identifier *Id* that appears in a statement of the form

  $$Id \;\triangleq\; \text{INSTANCE } M \;\ldots$$

  It adds or removes from the set of usable facts all facts imported from module M. The INSTANCE statement cannot have parameters—that is, it can't be of the form $Id(x) \;\triangleq\; \ldots$.

Theorems in certain special standard modules will direct TLAPS to use decision procedures or proof tactics. For example, there will be a theorem named *SimpleArithmetic* that causes TLAPS to apply a certain decision procedure for arithmetic when trying to prove something.

A definition specifier is the name of a defined operator—for example,

20

- F if the module contains the definition $F(x,\ y) \triangleq \dots$.

- $Ins!F$ if the current module contains $Ins(a) \triangleq$ INSTANCE $M \ \dots$ and F is defined in M.

The SANY parser also accepts the following two kinds of definition specifiers. However, they are not supported by TLAPS and will probably be eliminated from the language.

- MODULE *Name*, indicating that all definitions from the module *Name* are to be added or removed from the set of usable definitions. The module name must appear in an EXTENDS or INSTANCE statement or else be the name of the current module.

- An identifier *Id* that appears in a statement

$$Id(p_1, \dots, p_k) \quad \triangleq \quad \text{INSTANCE } M \ \dots$$

(possibly with $k = 0$). It indicates that all the definitions imported from the instantiation are to be added or removed.

7.2.2   BY

A terminal BY proof has the same syntax as a USE statement, except that it starts instead with the keyword BY. As explained below, at any point in a proof there will be sets of known and usable facts and of current and usable definitions. There will also be a current goal. A BY proof asserts that this goal follows easily from the set of usable facts together with the set of facts specified in the BY statement, using only those definitions contained in the set of usable definitions or specified by the statement. "Easily" means that a proof tool should be able to find the proof without any help from the user.

In addition to names of theorems, assumptions, and steps, a fact in a BY statement can be an arbitrary formula. Such a fact must follow easily from the set of usable facts together with the previous facts in the BY statement. For example, suppose the set of currently usable facts includes the fact $e \in S$. You might write

$\langle 3 \rangle 1. \ \forall x \in S : P(x)$
$\langle 3 \rangle 2. \ P(e+1)$
      BY $\langle 3 \rangle 1, \ P(e)$

The fact $P(e)$, which follows from $e \in S$ and $\langle 3 \rangle 1$, makes the proof easier to understand (and easier for a prover to check) by alerting the reader that

to prove $P(e+1)$ from usable facts and the fact $\langle 3 \rangle 1$, he (or it) should first note that $P(e)$ follows from these facts. Arbitrary expressions can also be used as facts in a USE, but not in a HIDE.

A BY ONLY proof begins with the keywords BY ONLY. Unlike in an ordinary BY proof, the current goal must follow easily from just the specified facts and the currently known domain formulas, without using any other usable facts. TLAPS uses all currently usable definitions plus the ones specified by the DEF clause.

TLAPS also allows the use in a BY or BY ONLY proof or in a USE statement of a fact that is trivially equivalent to a known (but not necessarily usable) fact. For example, if a module contains

$$\text{THEOREM } Elementary \;\triangleq\; 1+1=2$$

then the facts $Elementary$ and $1+1=2$ can be used interchangably anywhere within the scope of the definition of $Elementary$.

7.2.3 OBVIOUS **and** OMITTED

The terminal proof OBVIOUS asserts that the current goal follows easily from the set of known facts and the definitions contained in the set of usable definitions.

The terminal "proof" OMITTED means that the user is asserting the validity of the step without providing a proof. It asserts that the user has deliberately chosen not to provide a proof, and has not omitted it either accidentally or temporarily while writing other parts of the proof.

A proof is incomplete if it contains a statement with no proof. Incomplete proofs will be the norm while a user is developing the proof. TLAPS attempts to check a step only if the step has a proof other than the terminal "proof" OMITTED.

7.3 The Current State

At each point in a proof there is a current state that consists of:

- The set of current declarations.

- The set of current definitions and a subset consisting of the usable definitions.

- A set of currently known facts and a subset consisting of the usable facts.

- A *current* goal, which is a formula.

Recall that at the start of a theorem, there are sets of current declarations, current and usable definitions, and facts and usable facts described above. The state at the start of the theorem's proof is obtained by adding to these sets the following:

- If the theorem asserts a formula, then the formula becomes the current goal.

- If the theorem asserts an ASSUME/PROVE, then the declarations in the assumptions are added to the set of current declarations. The set of formulas and ASSUME/PROVEs asserted in the assumptions is added to the set of known facts; it is also added to the set of usable facts iff the theorem has no name. The PROVE formula becomes the current goal. (If an assumption is an ASSUME/PROVE, then the declarations of the inner ASSUME are not added to the set of current declarations.)

  Remember that the assumption NEW $C \in S$ is an abbreviation for the declaration NEW C and the assertion $C \in S$. An assertion of the form $C \in S$ obtained from a declaration is called a *domain formula*. Domain formulas are always added to the set of usable facts as well as to the set of known facts, even if the theorem is unnamed.

After the theorem's proof (if any), the current state reverts to the state right before the theorem, with the theorem added to the set of known facts iff it is named. An unnamed theorem can never be used in a proof.

To explain the meaning of a step, we describe the relation between the state of the proof at the (beginning of the) step and

- the state at the beginning of the statement's proof (if it has one), and

- the state immediately after the statement and its proof (if it has one).

7.4   Steps That Take Proofs

In the following descriptions, σ will be used to denote an arbitrary step-starting token.

7.4.1   Formulas and ASSUME/PROVE

A step that asserts a formula or an ASSUME/PROVE affects the state exactly the same way as a theorem. It makes either the formula or the PROVE assertion the current goal of the step's proof. The formulas and ASSUME/PROVE

assertions from the ASSUME clause are added to the set of usable facts iff the step is unnamed. However, domain formulas obtained from NEW clauses are always added to the set of usable facts.

After the step and its proof, the step's assertion is added to the set of usable facts iff the step is unnamed. (An unnamed step can never be referred to in a BY or USE, so the step's assertion must be put into the set of usable facts for it ever to be used.)

7.4.2   CASE

A CASE step consists of the step-starting token followed by the keyword CASE and a formula. The step "σ CASE F" is equivalent to

σ ASSUME F PROVE G

where G is the current goal. (Since G is already the current goal, this means that the current goal remains the same.)

7.4.3   @ Steps

A common method of proving an inequality is by proving a sequence of inequalities. For example, to prove $A \leq D$, we might prove $A \leq B \leq C \leq D$. Such a proof might appear inside a proof as follows (where the proofs of the individual steps are omitted).

    $\langle 2\rangle 3.$ $A \leq D$
        $\langle 3\rangle 1.$ $A \leq B$
        $\langle 3\rangle 2.$ $B \leq C$
        $\langle 3\rangle 3.$ $C \leq D$
        $\langle 3\rangle 4.$ QED
            BY $\langle 3\rangle 1, \langle 3\rangle 2, \langle 3\rangle 3$

It's a nuisance to have to write B and C twice if they're large formulas. TLA^{+2} provides the following abbreviated way of writing this proof.

    $\langle 2\rangle 3.$ $A \leq D$
        $\langle 3\rangle 1.$ $A \leq B$
        $\langle 3\rangle 2.$ @ $\leq C$
        $\langle 3\rangle 3.$ @ $\leq D$
        $\langle 3\rangle 4.$ QED
            BY $\langle 3\rangle 1, \langle 3\rangle 2, \langle 3\rangle 3$

This style of reasoning can be used with any transitive operator or combination of operators, such as

$$
\begin{array}{ccccccc}
A & = & B & = & C & = & D \\
A & \Rightarrow & B & \Rightarrow & C & \Rightarrow & D \\
A & \subseteq & B & \subseteq & C & \subseteq & D \\
A & \leq & B & < & C & \leq & D
\end{array}
$$

However, the token @ followed by an infix operator followed by an expression can be used in a step that follows any @ step or any formula step in which the formula's top-level operator is an infix operator. The "@" then refers to the right-hand side of the preceding step's formula. Although it's bad style and you shouldn't do it, you could write

$\langle 3 \rangle 4.\ A \leq B$
$\langle 3 \rangle 5.\ @ > C$

where the @ stands for B.

7.4.4 SUFFICES

The step σ SUFFICES A asserts that proving A proves the current goal, where A can be a formula or an ASSUME/PROVE. At the beginning of the step's proof, A is added to the set of known facts and to the set of usable facts. (The proof must prove the current goal.) After the step and its proof:

- If A is a formula, then it is made the current goal.

- If A is an ASSUME/PROVE, then:

  - The declarations in its assumptions are added to the set of current declarations and the domain formulas from those declarations are added to the sets of known and usable facts.

  - The assertions among its assumptions (the formulas and the ASSUME/PROVEs) are added to the set of known facts. They are added to the set of usable facts iff the step is not named.

  - The PROVE formula is made the current goal.

7.4.5 PICK

A PICK step has the same syntax as one that asserts a $\exists$ formula, except with the $\exists$ replaced by the token PICK. For example, the step

σ PICK $x \in S$, $y \in T$: $P(x, y)$

asserts that there exist values x in S and y in T satisfying $P(x, y)$, and then declares x and y to be equal to an arbitrary pair of such values. The state at the start of the step's proof is the same as for the formula obtained by replacing PICK by $\exists$. After the proof:

- CONSTANT declarations of the identifiers introduced by the step are added to the set of declarations and the domain formulas of those declarations are added to the sets of known and usable facts. (In this example, the domain formulas are $x \in S$ and $y \in T$.)

- The body of the PICK (in this example, the formula $P(x, y)$) is added to the set of known facts. It is added to the set of usable facts iff the step is not named.

A PICK step is effectively translated to two steps. For example, the step and its proof

σ PICK $x \in S$, $y \in T$: $P(x, y)$
        PROOF Π

are translated to

ρ $\exists\, x \in S$, $y \in T$: $P(x, y)$
        PROOF Π
σ SUFFICES ASSUME   NEW $x \in S$, NEW $y \in T$
                        $P(x, y)$
            PROVE G

and σ contains a proof asserting that it follows from ρ, where ρ is a new step name and G is the current goal. This translation is relevant to the meaning of the step name σ. (See Section 7.6.2 on page 31.)

7.4.6   QED

The state at the beginning of a QED step's proof is unchanged. After the step and its proof, the state is determined by the rule for the step whose proof the QED step ends.

7.5 Steps That Do Not Take Proofs

7.5.1 Definitions

In a *definition* step, the step-starting token is followed by the optional token DEFINE and a sequence of operator definitions, function definitions, and/or module definitions, where a module definition is something like

$$Ins(x) \;\triangleq\; \text{INSTANCE } M \text{ WITH } \ldots$$

It has the same effect on the state as the corresponding (top-level) statements. The definitions introduced by the step (which are the definitions of the imported and renamed operators for a module definition) are added to both the set of definitions and the set of usable definitions.

7.5.2 INSTANCE

An INSTANCE step consists of a step-starting token followed by an ordinary INSTANCE statement (one that begins with the keyword INSTANCE). It has the same effect on the state as the corresponding (top-level) statement.

7.5.3 USE and HIDE

A USE or HIDE step has the same syntax as the corresponding (top-level) statement, except preceded by the step-starting token. It affects the sets of usable facts and definitions the same way as the corresponding USE or HIDE statement. As explained in Section 7.6 below, a USE or HIDE step can name facts or definitions made in earlier steps.

There is also a USE ONLY step, in which the keyword USE is followed by the keyword ONLY. It sets the usable facts to be only known domain facts and facts specified by the step. It affects the set of usable definitions the same way as an ordinary USE step.

7.5.4 HAVE

A HAVE step consists of a step-starting token followed by HAVE and a formula. For the statement

$$\sigma \text{ HAVE } F$$

to be correct, the current goal must be syntactically of the form $H \Rightarrow G$ for some formulas H and G, and the formula $H \Rightarrow F$ must be an obvious

consequence of the known facts and usable definition. In that case, the step is equivalent to

> σ SUFFICES ASSUME F
> PROVE G

plus a BY ONLY proof that permits using only the fact ASSUME F PROVE B (a fact that must therefore be easily provable with no assumptions). Thus, this step means that we are going to prove the current goal by assuming F and proving G.

7.5.5 TAKE

A TAKE step consists of a step-starting token followed by TAKE followed by anything that could come between "$\forall$" and its matching ":"—for example

> σ TAKE $x, y \in S,\, z \in T$

This step is typically used when the current goal is

> $\forall\, x, y \in S,\, z \in T : G$

for some formula G. It means that we are going to prove this goal by declaring x, y, z to be constants, assuming $x \in S$, $y \in S$, and $z \in T$, and proving G. More precisely this TAKE statement is equivalent to

> σ SUFFICES ASSUME NEW CONSTANT $x \in S$,
> NEW CONSTANT $y \in S$,
> NEW CONSTANT $z \in T$
> PROVE G

followed by a proof that permits using only the domain formulas $x \in S$, $y \in S$, and $z \in T$.

In general, for the step σ TAKE τ to be correct, the current goal should be obviously equivalent to $\forall\, \tau : G$ for some formula G. (Again, the meaning of "obviously equivalent" is not specified.) In that case, G is made the current goal, constant declarations of the bound identifiers in τ are added to the current set of declarations, and any formulas of the form $id \in e$ in τ are added to the set of known facts and to the set of usable facts.

7.5.6   WITNESS

A WITNESS step consists of a step-starting token, followed by WITNESS, followed by a comma-separated list of expressions. A WITNESS step is used to prove an existentially quantified formula by specifying instantiations of its bound identifiers. There are two cases in which the statement σ WITNESS $e_1, \ldots, e_k$ is correct:

- The current goal is obviously equivalent to a formula $\exists\, id_1, \ldots, id_k :$ G. In this case, the WITNESS step is equivalent to

  > σ SUFFICES $\overline{G}$
  >   PROOF OBVIOUS

  where $\overline{G}$ is the formula obtained by substituting each e_j for id_j in G, for j in $1 \,.\,.\, k$.

- The current goal is obviously equivalent to a formula $\exists\, \iota_1 \in S_1, \ldots, \iota_k \in S_k : G$ where each ι_j is an identifier, there is some substitution of expressions for these identifiers that transforms each $\iota_j \in S_j$ to e_j, and each e_j is easily provable from the current set of usable facts. In this case, the formula obtained from G by the aforementioned substitution of expressions for the identifiers in the ι_j is made the current goal, and the domain formulas e_j are added to the set of known facts and to the set of usable facts. (Adding a fact that is easily provable to the set of usable facts might make additional facts easily provable from that set.) For example, if the current goal is

  > $\exists\, x, y \in S,\, z \in T : G(x, y, z)$

  then the step

  > $\langle 3 \rangle 4.$ WITNESS $expX \in S,\ expY \in S,\ expZ \in T$

is equivalent to

> $\langle 3 \rangle 4$. SUFFICES $G(expX,\ expY,\ expZ)$
> > $\langle 4 \rangle 1$. $expX \in S$
> > > PROOF OBVIOUS
> > $\langle 4 \rangle 2$. $expY \in S$
> > > PROOF OBVIOUS
> > $\langle 4 \rangle 3$. $expZ \in T$
> > > PROOF OBVIOUS
> > $\langle 4 \rangle 4$. QED
> > > BY ONLY $\langle 4 \rangle 1, \langle 4 \rangle 2, \langle 4 \rangle 3$

7.6 Referring to Steps and Their Parts

Within a proof, steps and their parts can be named in three contexts: as ordinary expressions, as facts in a BY, USE, or HIDE, and in the DEF clause of one of those statements. We now consider these three possibilities.

7.6.1 Naming Subexpressions

Formulas The name of a step that asserts a formula names that formula. For example, the step

> $\langle 2 \rangle 3$. $x + y = z$

defines $\langle 2 \rangle 3$ to equal $x + y = z$. The step name $\langle 2 \rangle 3$ can be used like any other defined symbol—for example:

> $\langle 2 \rangle 3 \wedge (z \in Nat) \ \Rightarrow\ (x + y - z = 0)$

We can also use labels and/or positional selectors to name subexpressions of $\langle 2 \rangle 3$ the same way we name subexpressions of other defined symbols—for example, $\langle 2 \rangle 3!\langle$ names the subexpression $x + y$. (See Section 6.)

ASSUME/PROVE **Steps** The parts of an ASSUME/PROVE step are named as explained in Section 6.4, where the step number names the ASSUME/PROVE. Thus, in

> $\langle 3 \rangle 4$. ASSUME $\quad P$, ASSUME $\quad Q$
> > PROVE $\quad R$
> PROVE $\quad S$

$\langle 3 \rangle 4!1$ names P, $\langle 3 \rangle 4!2!\langle$ names Q and $\langle 3 \rangle 4!3$ names S.

As explained in Section 6.4, subexpressions of an ASSUME/PROVE can be used only within the scope of any identifier that could appear in that subexpression (even if it that identifier doesn't actually appear in it).

CASE, HAVE, SUFFICES, and WITNESS Steps   Expressions within a CASE, HAVE, SUFFICES, or WITNESS step are named as if CASE, HAVE, SUFFICES, and WITNESS were prefix operators—CASE, HAVE, and SUFFICES taking a single argument and WITNESS taking an arbitrary number of arguments. Thus, in

$\langle 2 \rangle 3.$ CASE $x + y > 0$

$\langle 2 \rangle 4.$ WITNESS y, $x + 1$

$\langle 2 \rangle 3!1$ equals $x + y > 0$ and $\langle 2 \rangle 3!1!\langle$ equals $x + y$, while $\langle 2 \rangle 4!2$ equals $x + 1$. The "argument" of SUFFICES can be an ASSUME/PROVE, whose subexpressions are named as described above for an ASSUME/PROVE step.

PICK and TAKE   A subformula of a PICK step is named as if the PICK were replaced by $\forall$. For example, in

$\langle 3 \rangle 4.$ PICK $x \in S$, $y \in T : x + y > 0$

$\langle 3 \rangle 4!2$ names T and $\langle 3 \rangle 4!(e, f)$ names $e + f > 0$. The naming of a *take* step is similar, except that there is no "body" to name, only the sets that follow an " $\in$ ".

Note that the symbols introduced in a PICK step are not declared within the proof of that step, but they are declared after the proof. However, references to the body of the PICK are made the same way in both places.

7.6.2   Naming Facts

Syntactically, any expression can be used as a fact. (A proof tool might accept only a restricted set of expressions as facts.) Any named step that makes an assertion can also be used as a fact. The only kinds of steps that can *not* be used as facts are USE, HIDE, *definition*, INSTANCE, and QED.

The scope of a step name includes the proof of the step. Thus, it is legal to use a subexpression of a step named σ within that step's proof—for example, to name an assumption if σ is an ASSUME/PROVE step.

When used by itself (and not in the name of one of its subexpressions), a step name denotes the fact or facts that the step adds to the current set of known facts. We now explain exactly what that means.

We have described above how every step that makes an assertion is equivalent to one of the form A or SUFFICES A, where A is either a formula or an ASSUME/PROVE. If we consider a formula G to be equivalent to ASSUME TRUE PROVE G, then any step σ is equivalent to a step of the form $\sigma\ A$ or σ SUFFICES A for an ASSUME/PROVE A.

- For the step $\sigma\ A$, within the proof's step the step name σ denotes the set of assumptions of A; outside the proof it denotes A.

- For the step σ SUFFICES A, within the proof's step the step name σ denotes A; outside the proof it denotes the set of assumptions of A.

It is quite useful to have a step name σ refer to the known facts introduced into the current context by the step, since those facts are not automatically added to the set of usable facts. However, it has the unfortunate effect of making a proof look circular when σ is used as a fact within the proof of the step named σ. Readers and writers of TLA$^+$ proofs should quickly get used to this convention.

One may want to refer to a long formula G inside a step $\sigma\ G$. For example, we can assume $\neg G$ in a proof by contradiction of the step. However, by these rules, σ names TRUE within the proof of σ, so we cannot write $\neg G$ as $\neg\sigma$. We can write $\neg G$ as $\neg\sigma!$: instead.

7.6.3 Naming Definitions

Only the names of defined operators may appear in the DEF clause of a BY proof or a USE or HIDE step. These include the names of operators defined in LET clauses. The step name of a DEFINE step may also be used in a DEF clause. If the step defines more than one operator, then the step name applies to all of them—but not to any operators defined in LET clauses within those definitions.

Remember that an operator name does not contain any parameters or any parentheses. For example, the expression $Ins(42)!Foo(x, y)$ is an application of the operator named $Ins!Foo$ to the three arguments 42, x, and y. Section 6.5 explains how to name operators defined in a LET clause.

7.7 Referring to Instantiated Theorems

Suppose module M contains a theorem

THEOREM T

and another module *MI* imports *M* with the statement

$$I \triangleq \text{INSTANCE } M \text{ WITH } \ldots$$

As explained in Section 17.5.5 of *Specifying Systems*, this imports the theorem that we can write in TLA^{+2} as

ASSUME $\overline{A_1}, \ldots \overline{A_k}$
PROVE $\overline{T}$

where A_1, $\ldots$, A_k are the assumptions asserted by ASSUME statements in M, and $\overline{A_i}$ and $\overline{T}$ are the formulas obtained from A_i and T by performing the substitutions specified by the WITH clause.

If the theorem is named, as in

THEOREM *Thm* $\triangleq$ T

then *I!Thm* names the imported fact—the ASSUME / PROVE above. However, the rules of Section 6.4 above for naming parts of an ASSUME / PROVE do not apply to this fact. The name *I!Thm!*: refers to the formula $\overline{T}$. A formula $\overline{A_i}$ can be referred to in module *MI* only if it has been assigned a name in module M, in which case it is named *I!* ... as usual. Because TLC cannot do anything with an ASSUME / PROVE, it treats *I!Thm* (as well as *I!Thm!*:) as the name of $\overline{T}$.

A proof of module *MI* that uses the imported theorem *Thm* generally wants to use $\overline{T}$. To do that, it must prove all its hypotheses $\overline{A_i}$. Often, the formulas $\overline{A_i}$ are simple consequences of the assumptions of module *MI*. In that case, a proof can simply use *I!Thm* in a context in which those assumptions of *MI* are usable facts.

7.8   Temporal Proofs

Temporal-logic reasoning has not yet been completely implemented in TLAPS. Currently, TLAPS can perform only propositional temporal-logic reasoning, meaning it can't prove formulas that involve quantification over temporal formulas. We expect that a complete implementation will require no changes to anything described in this document. However, since temporal logic is different from ordinary logic, the reasoning involved in temporal proofs is somewhat different from what mathematicians are used to. This is explained in Section 8.3 on page 35.

8 The Semantics of Proofs

8.1 The Meaning of Boolean Operators

As discussed in Section 16.1.3 of *Specifying Systems*, there are various ways to define the meanings of the Boolean operators on non-Boolean arguments. For example, we know that $x \wedge y$ equals $y \wedge x$ for any Booleans x and y. However, is $5 \wedge 7$ equal to $7 \wedge 5$? TLAPS uses what is called in *Specifying Systems* the *liberal* interpretation. In this interpretation, $(5 \wedge 7) = (7 \wedge 5)$ is true, and TLAPS will prove it.

The precise interpretation of the Boolean operators is in terms of an operator *ToBoolean* such that *ToBoolean*(x) is some Boolean that equals x if x is a Boolean. More precisely, *ToBoolean* is assumed to satisfy

$$\wedge \ \forall x \ : \ \mathit{ToBoolean}(x) \in \text{BOOLEAN}$$
$$\wedge \ \forall x \in \text{BOOLEAN} \ : \ \mathit{ToBoolean}(x) = x$$

For example, we can define conjunction $\wedge$ by

$$x \wedge y \ \triangleq \ \mathit{ToBoolean}(x) \underline{\wedge} \mathit{ToBoolean}(y)$$

where $\underline{\wedge}$ is ordinary conjunction on Booleans. The operator *ToBoolean* exists only in the semantics and is not defined at the TLA$^+$ level. However, it follows from the assumptions about *ToBoolean* that it can be defined in TLA$^+$ by

$$\mathit{ToBoolean}(x) \ \triangleq \ (x \equiv \text{TRUE})$$

8.2 The Meaning of ASSUME/PROVE

An ASSUME/PROVE essentially asserts that its assumptions imply its PROVE formula. For example,

$$\text{ASSUME} \quad A_1, \ \ \text{NEW } x \in S, \ \ A_2$$
$$\text{PROVE} \quad B$$

asserts the formula

$$\forall x \ : \ A_1 \wedge (x \in S) \wedge A_2 \ \Rightarrow \ B$$

except that x cannot be a free identifier in A_1 in the ASSUME/PROVE.

If an assumption declares an operator, as in NEW $op(\_,\_)$, then translating the ASSUME/PROVE to a formula would require quantification over the operator op, which can't be done in first-order logic. However, an AS-SUME/PROVE (even with assumptions consisting of an ASSUME/PROVE) can

be translated to a formula in which quantification over operators occurs only in outermost $\forall$ quantifiers. Since an ASSUME/PROVE can appear only as the statement of a theorem or a proof step (which is a theorem asserted in a certain context), TLA$^+$ remains essentially in the realm of first-order (temporal) logic. It raises none of the issues associated with full second-order logic.

8.3 Temporal Proofs

The meaning of temporal TLA$^+$ formulas is explained in Chapter 8 of *Specifying Systems*. A temporal formula is a predicate on behaviors, which are sequences of states. For a temporal formula F, the statement THEOREM F asserts that F is true on all behaviors. As explained in Section 8.3 of that book, TLA$^+$ satisfies the following law:

> **Necessitation Rule** For any formula F: if THEOREM F is true, then THEOREM $\Box F$ is true.

(The book called it the *Generalization Rule*, but it's called the *Necessitation Rule* by logicians.) The Necessitation Rule is very different from:

> THEOREM ASSUME F
> PROVE $\Box F$

This statement is equivalent to

> THEOREM $F \Rightarrow \Box F$

which is not true for all formulas F.

The Necessitation Rule is important because many temporal proof rules have assumptions of the form $\Box F$. For example one commonly used rule is:

> THEOREM $BoxImplies$ $\triangleq$
> ASSUME TEMPORAL F, TEMPORAL G, $\Box F$, $\Box(F \Rightarrow G)$
> PROVE $\Box G$

We use it to prove that if P is an invariant of a specification $Spec$ and P implies Q, then Q is an invariant of $Spec$. In other words, we deduce

> THEOREM $QInvar$ $\triangleq$ $Spec \Rightarrow \Box Q$

from

> THEOREM $PInvar$ $\triangleq$ $Spec \Rightarrow \Box P$
> THEOREM $PimpliesQ$ $\triangleq$ $P \Rightarrow Q$

This deduction follows from *BoxImplies* because, by the Necessitation Rule, theorem *PImpliesQ* implies the truth of THEOREM $\Box(P \Rightarrow Q)$.

We would like to use the Necessitation Rule inside a proof, to deduce $\Box F$ from a proof step that asserts F. However, that's not sound. For example, suppose we are proving that *Inv* is an invariant of a specification $Init \wedge \Box[Next]_v$. Our proof might begin

    ⟨1⟩1. SUFFICES ASSUME $Init,\ \Box[Next]_v$
                PROVE $\Box Inv$

    ⟨1⟩2. *Inv*

We can't apply the Necessitation Rule to deduce $\Box Inv$ from step ⟨1⟩2. The proof of ⟨1⟩2 could have used the assumption *Init*, so all we know is that *Inv* is true of the initial state. We can't conclude that it's true of all states in a behavior satisfying the specification.

We need a generalization of the Necessitation Rule so it can be applied to proof steps when it's valid. For that, define a formula F to be a $\Box$ formula iff THEOREM $F \equiv \Box F$ is true—that iff, F is true on behavior β iff $\Box F$ is, for any behavior β. Since $\Box F \Rightarrow F$ is true for any formula F, we can also define F to be a $\Box$ formula iff THEOREM $F \Rightarrow \Box F$ is true. For any formulas F and G, the following formulas are all $\Box$ formulas:

$$\Box F \qquad \Diamond \Box F \qquad \Box F \vee \Box G \qquad 1 + 1 = 3$$

It follows from these that weak fairness (WF) and strong fairness (SF) formulas are $\Box$ formulas. We can now write:

> **Generalized Necessitation Rule** A proof of a formula F using only assumptions that are $\Box$ formulas proves $\Box F$.

The Necessitation Rule follows from this because a THEOREM in a module must be proved using only ASSUME statements and previously proved theorems, and the formula of an ASSUME statement must be a constant, so it is a $\Box$ formula.

Currently, TLAPS can only perform propositional temporal reasoning; it cannot reason about quantified temporal formulas. Using *PTL* in a BY clause tells TLAPS to use a propositional temporal-logic prover. (*PTL* uses the Generalized Necessitation Rule.) We plan eventually to add a back-end prover that can handle temporal formulas with constant quantification— that is, ones with the operators $\forall$ and $\exists$. We would also like to add a way for TLAPS to apply the rule

    THEOREM ASSUME   NEW TEMPORAL $F(\_)$, NEW STATE e
             PROVE $F(e) \Rightarrow \exists x : F(x)$

that justifies proofs by refinement mappings. There is little incentive to implement more general reasoning about temporal quantification.