

Lecture Notes, Computer-Aided Formal Reasoning

Martin Hofmann

SoSe 2016

1 Introduction

Computer-aided theorem proving means to carry out mathematical proofs on a computer whose job it is to check steps, to perform bookkeeping tasks and to automate routine steps. Conducting a proof on a computer may be compared to and has a lot in common with implementing an informally given algorithm or model. For example, a number of details must be filled in and, more importantly, mistakes and shortcomings of the high-level model are brought to the surface.

Computer-aided theorem proving has numerous applications in program and hardware verification as well as prototype development. To a lesser, perhaps increasing, degree it is used to aid the development of genuine mathematical proofs. In particular, it is becoming popular in proofs about programming languages and type systems, e.g. type soundness properties.

1.1 Course outline

In this course, we will get to know the computer-based theorem prover PVS (pvs.csl.sri.com) along with its theoretical foundations and some ramifications thereof. We will also get a glimpse at the theorem prover Coq whose theoretical basis is type theory rather than set theory.

- Logical foundations: sequent calculus, predicate calculus, higher-order logic, set theory
- Automation of logical reasoning: resolution

- Automation of equational reasoning: rewriting and decision procedures
- Type theory: Modularisation, independent checking of proof certificates, computation within proofs.

Mostly in the tutorials we will apply this knowledge to a variety of problems from

- Solving logical puzzles
- Algorithms on lists and trees
- Hardware components such as adder, counter, multiplier
- Distributed algorithms using invariants and reasoning
- Experiments with other theorem provers.

1.2 Notions of proof

What exactly is a “proof”. When asked this question a typical mathematician would produce something like the following: A proof is a convincing, undebatable argument establishing the truth of a mathematical statement. Back to Euclid (300 B.C.) goes the following concretion of this definition: A proof is a derivation of a statement from axioms by means of logical rules.

This sound good, but it remains to say what “axioms” should be and what the “logical rules” are. In Euclid’s case (geometry) the axioms were truisms such as “for any two non-equal points there is exactly one line passing through them”. The logical rules were essentially the ones we still use today and will learn about later in the course. An example of such a rule: if “ A implies B ” holds and “ A ” holds then “ B ” holds, too (*modus ponens*).

Later on, more complicated concepts such as real numbers and limits were introduced which made it less clear what reasonable axioms should be. For example, even the famous 18th century mathematician Leonhard Euler struggled with the infinite series $1 - 1 + 1 - 1 + 1 - 1 + \dots$ and ended up ascribing the value $1/2$ to it on the basis of the same informal mathematical reasoning he used for his celebrated theorems.

The lack of solid logical foundations for mathematics, and in particular analysis led to an actual crisis in mathematics (*Grundlagenkrise*) which was settled early in the last century by the invention of set theory (and, following up on this,

by the formalisation of real numbers, limits, integrals based on work by Weierstrass, Riemann and others.)

1.2.1 Set theory

Set theory is a formalism which allows one to *define* all other mathematical concepts and to *prove* their axioms, thus enabling a rigorous proof of their *consistency*, i.e., sensibility. For instance, we can define points as triples of real numbers (which in turn are defined as certain sequences of rational numbers (which in turn are defined as certain pairs of integer numbers (which in turn are defined as certain pairs of natural numbers (which are defined as certain sets: $0 = \emptyset$, $1 = \{0\}$, $2 = \{0, 1\}$, etc.)))) and then *prove* that through any two distinct points goes one and only one line (a line being defined, for instance, as a set of points satisfying some linear relation).

The present formulation of set theory consists of approximately nine axioms¹ among them

- two sets having the same elements are equal
- for each set we can form the set of its subsets
- there exists an infinite set
- for each set we can form the subset consisting of those elements sharing a given property
- for each set of sets A there exists a set containing exactly one element of each nonempty set in A (“axiom of choice”)

Using a formalised language statements in set theory can be written as strings and recognised as such, for example the first four axioms are written as follows:

- $\forall a. \forall b. (\forall x. x \in a \iff x \in b) \Rightarrow a = b$
- $\forall a. \exists b. \forall x. x \in b \iff x \subseteq a$
(where $x \subseteq a \stackrel{\text{def}}{=} \forall y. y \in x \Rightarrow y \in a$)

¹the precise number depends on what we count as axiom and what as a logical rule.

- $\exists a. \neg \text{finite}(a)$
 (where $\text{finite}(a) \stackrel{\text{def}}{=} \forall b. (\emptyset \in b \wedge (\forall c. \forall x. c \in b \wedge x \in a \Rightarrow c \cup \{x\} \in b)) \Rightarrow a \in b$)

Here $c \cup \{x\}$ is a notation for a set whose existence is asserted by two other axioms (union and singleton).

- $\forall a. \exists b. \forall x. x \in b \Leftrightarrow (x \in a \wedge \phi(x))$ (where $\phi(x)$ is an arbitrary statement involving x)

We should remark at this point that despite the formal notation the axioms of set theory necessarily remain unproved and their justification relies on philosophical and pragmatic arguments.

The *logical rules* of set theory are precisely the ones of first-order logic which we are going to learn about in more detail later in the course.

1.2.2 Proofs as formal derivations

Once we have a formal concept of axioms and rules, we can define a *proof* of a statement ϕ as a sequence of statements

$$\phi_0, \phi_1, \phi_2, \phi_3, \dots, \phi_n = \phi$$

ending in ϕ such that each ϕ_i is either an axiom or follows from previous statements by a logical rule.

So, to check whether an alleged proof indeed is one is a matter of entirely mechanical symbol manipulation and does not require any creative skills or intelligence.

Rather than merely asserting the next formula ϕ_i one might tell by which logical rule it follows and which of the ϕ_j , $j < i$ were used as premises for that inference. In this way, one arrives at the notions of proof tree or proof-DAG. (DAG=directed acyclic graph = tree with shared nodes).

In practice, however, writing out proofs at this level of detail would be far too cumbersome and so checking whether a purported proof, say in a journal submission, indeed is one, does require considerable mathematical skill and devotion!

And while the vast majority of mathematicians agrees that any proof in mathematics can *theoretically* be formalised in set theory and hence mechanically checked, many of them believe that *in practice* such formalisation is impossible for all but the simplest toy examples. This belief might have remained unchallenged if there had not been the request for formalised proofs from informatics and the advent of sufficiently powerful machines.

1.3 Proof assistants

So, why do we need formalised proofs in informatics? Well, the correctness of software (or hardware) is nothing but a mathematical statement amenable to formalisation and mechanical checking.

Here are some examples of formal theorem proving occurring in informatics.

- Is program X correct? *very rare*
- Does method X satisfy invariant Y ?
- Does variable X always hold values with property Y ? E.g., Y =points to a sorted linked list, a balanced binary tree, data items consistent with store.
- Does protocol X guarantee property Y ? E.g., Y =cache coherence, sequential execution, absence of deadlock.
- Does circuit X implement function Y ? E.g. Y =FP multiplication, Fourier transform.
- Does algorithm X satisfy specification Y ? E.g. X =garbage collector, Y =absence of interference+liveness.
- Does theorem X about programming language Y hold? E.g. X = type safety, correctness of proof rules.
- Verification of certificates (“proof-carrying code”)
- Is the code generated by a compiler equivalent to the source code.

While in early stages of soft- and hardware development these proofs could be carried out by hand (possibly using some notation and intermediate calculations) the size of systems has reached a state where this has become impossible in many cases.

Prompted by these requirements systems called *proof assistants* or *theorem provers* have been developed which perform not only the task of checking sizeable formalised proofs but also help with coming up with formalised proofs in the first place by bookkeeping assumptions and variables, providing tactics and decision procedures (e.g., for propositional formulas, certain fragments of arithmetic, modal and temporal logic, equational theories, etc.) and by providing libraries of definitions and already proved theorems.

- Bookkeeping of assumptions and variables
- Tactics
- Type checking
- Decision procedures
- Libraries of definitions and theorems

Figure 1: Tasks of a proof assistant

- Expressive logic
- Powerful decision procedures
- Large body of basic notions

Figure 2: Strengths of PVS

1.3.1 The PVS System

In this course we will get to know one such proof assistant in some detail, namely the PVS system developed by Owre, Rushby, and Shankar at SRI, Menlo Park, California. As any system, PVS has a number of strengths and also weaknesses. As a partial compensation for the weaknesses we will later in the course take a look at complementary systems, in particular SPASS and Coq.

PVS has a very expressive underlying logic (classical higher-order logic), it comes equipped with a number of powerful decision procedures, e.g., for linear arithmetic and equational reasoning, and it has a large body of basic notions which allow one to start a formalisation on a relatively high level.

On the other hand, PVS has recurrent soundness problems, that is, from time to time someone finds out that a weird combination of tactics use and language features allows one to prove $0 = 1$! Moreover, there is no formal representation of proofs. One reason why these problems are not trivial to fix is precisely the large body of basic notions which here turns into a disadvantage.

- Soundness problems
- No formal representation of proofs
- Large body of basic notions
- (Bad heuristics for first-order instantiation)

Figure 3: Weaknesses of PVS

1.3.2 Soundness and proof objects

A problem with a proof assistant is that its correct behaviour is hard to verify. Whether a word processor provides decent looking output can be checked at a glance (correctness for all inputs notwithstanding), similarly a video game either is fun to play with or not.

On the other hand, correct behaviour of a proof assistant is rather hard to detect. After all it's because we don't want to do the proofs by hand that we use a proof assistant in the first place. The "output" of a proof assistant doesn't consist of a nice looking document or a thrilling sequence of images.

In PVS you can have a complicated looking subgoal to prove, you type in `(grind)` and PVS responds that this proves the statement. There is no way to check this proof independently; all that's being recorded is that the tactic `(grind)` has been invoked.

PVS stands for *prototype verification system* which is explained by the following quote from the PVS Prover Guide 2.3, see `pvs.csl.sri.com`.

The primary purpose of PVS is to provide formal support for conceptualization and debugging in early stages of the life cycle of a hardware or software system. In these stages, both the requirements and designs are expressed in abstract terms that are not necessarily executable. We find that the best way to analyse such an abstract specification is by attempting proofs of desirable consequences of the specification.

So, provided soundness problems occur rarely ² they do not really compromise the usability of the system.

²they do sometimes, see the PVS web site

```

|-----
[1]  FORALL (y: t, v_106: list[t]):
      (FORALL (x: t):
        occ(x, merge(null, v_106)) =
          occ(x, null) + occ(x, v_106))
      IMPLIES
      (FORALL (x: t):
        occ(x, merge(null, cons(y, v_106))) =
          occ(x, null) + occ(x, cons(y, v_106)))
Rule? (grind)

```

lots of rewrites etc. are printed

This completes the proof of merge1.4.

Figure 4: A quick proof in PVS

Proofs as guarantee In recent years researchers have proposed a use of proofs as a certificate not unlike the cryptographic certificates such as digital signatures, etc. While the latter certify authenticity of a datum, i.e., a relationship between the datum and the sender, a formal proof certifies a property of the datum itself which is independent of the sender.

For example, a third-party provider of a component of a safety-critical system might be required to provide a formal proof of correctness.

A referee of a paper in mathematics or theoretical informatics might not be willing to verify all details of a proof but would rather run the formalised proofs of the theorems in the paper through a proof checker.

Also it has been proposed under the name *proof-carrying code* that mobile code should be equipped with independently checkable proofs of certain safety properties, e.g. memory safety, type safety, etc.

The design of systems like Coq is such that independent verification is possible. These systems generate a formal representation of a proof (a *proof object* or *proof term*) amenable to separate verification by a *proof checker* which is simple and small. Even if tactics or decision procedures contain bugs these will always show up at the checking stage so the worst that can happen is that an attempted proof has to be redone.

At present it seems that none of the systems with explicit proof objects can

Atoms: A, B, C, D, \dots

Connectives:

$\phi \wedge \psi$: ϕ “and” ψ (conjunction)

$\phi \vee \psi$: ϕ “or” ψ (disjunction)

$\phi \Rightarrow \psi$: ϕ “implies” ψ (implication)

$\neg\phi$: “not” ϕ (negation)

Precedence: $\neg, \wedge, \vee, \Rightarrow$

Example: $(A \Rightarrow B) \wedge \neg A \Rightarrow \neg A$ reads $((A \Rightarrow B) \wedge (\neg A)) \Rightarrow (\neg A)$

Figure 5: Syntax of propositional formulas

compete with PVS or similar systems. However, I think that this is mainly a problem of organisation and manpower, not an inherent theoretical one. I believe that in the not too distant future we will see powerful proof assistants with (almost) bug-free proof-checkers inside.

2 Sequent calculus

In this section we will learn about the “logical rules” which underly the PVS system. In their present form they were introduced by the logician Gerhard Gentzen around 1940 as a means to analyse the proof-theoretic strength of formal arithmetic.

2.1 Formulas

We start with a set of *atoms*, aka *identifiers* or *symbols* A, B, C, D, \dots . *Formulas* are built up from atoms by the *connectives* $\vee, \wedge, \Rightarrow$ (binary) and \neg (unary), so $(A \Rightarrow B) \wedge \neg A \Rightarrow C$ is a formula. By convention the binding power is $\neg > \wedge > \vee > \Rightarrow$, so the above formula reads $((A \Rightarrow B) \wedge (\neg A)) \Rightarrow C$.

The *meaning* of a formula is given relative to an interpretation of the atoms as either true or false. For instance, if A is true, and B, C are both false, then our example formula will be true because then $A \Rightarrow B$ is false (the only way for an implication to be false is that its antecedent (here A) is true and its consequent (here B) is false).

Digression on semantics of implication Please notice that this so-called *classical interpretation of implication* is sometimes at odds with our intuitive understanding of implication. For instance, the sentence “if MH wears a tie during the lecture then he can turn lead into gold.” is actually true under this interpretation. It is possible to formalise the intuitive meaning of such sentence by implicitly quantifying over a set of *worlds* that describe possibilities. One could then imagine a world in which MH wears ties and also worlds in which he can turn lead into gold, but the latter would not form a superset of the former because there is no causal relationship whatsoever. On the other hand, for $\phi \Rightarrow \psi$ to be valid in this refined sense one requires that the set of worlds in which ψ holds forms a superset of the set of worlds in which ϕ holds.

Another paradox involving classical implication goes as follows: it is commonly agreed that in order to show that a recursively defined method $m()$ is correct it suffices to show that its body e is correct assuming that any recursive calls to the method already perform correctly. In other words:

$$(m() \text{ correct} \Rightarrow e \text{ correct}) \Rightarrow m() \text{ correct}$$

Were this rule valid in the sense of classical implication then any method would be correct: either it is correct in the first place or else it isn't in which case the premise to the above rule is trivially true whereby correctness of the method follows from the rule!

Formalisation of meaning Anyway, under the aforementioned classical interpretation the formula $\phi \stackrel{\text{def}}{=} \neg A \Rightarrow (A \Rightarrow B)$ always comes out true, no matter what A and B actually stand for.

Likewise, $((A \Rightarrow B) \Rightarrow A) \Rightarrow A$ always holds.

Formally, a partial function η mapping atoms to $\{\text{tt}, \text{ff}\}$ can be extended to formulas by interpreting the connectives in the obvious way.

A formula is a *tautology* if its meaning is true regardless of the interpretation of the atoms.

A formula is *satisfiable* if it is true for *some* interpretation of the atoms. Clearly, A is satisfiable if and only if $\neg A$ is not a tautology and therefore A is a tautology if and only if $\neg A$ is unsatisfiable.

2.2 Applications of propositional logic

Many naturally occurring problems admit encodings in propositional logic in the sense that to know whether a certain formula is satisfiable or a tautology amounts

If $\eta(A) = \mathbf{tt}$, $\eta(B) = \eta(C) = \mathbf{ff}$ then

$$\begin{aligned} & \eta((A \Rightarrow B) \wedge \neg A \Rightarrow C) \\ = & \eta(A \Rightarrow B) \wedge \eta(\neg A) \Rightarrow \mathbf{ff} \\ = & (\mathbf{ff} \Rightarrow \mathbf{tt}) \wedge \mathbf{ff} \Rightarrow \mathbf{ff} = \mathbf{tt} \end{aligned}$$

Figure 6: Formal meaning of a formula

$$A \Rightarrow A$$

$$((A \Rightarrow B) \Rightarrow A) \Rightarrow A$$

$$\begin{aligned} & \neg((A_{11} \vee A_{12}) \wedge (A_{21} \vee A_{22}) \wedge (A_{31} \vee A_{32}) \wedge \\ & \quad \neg(A_{11} \wedge A_{21}) \wedge \neg(A_{11} \wedge A_{31}) \wedge \neg(A_{21} \wedge A_{31}) \wedge \\ & \quad \neg(A_{12} \wedge A_{22}) \wedge \neg(A_{12} \wedge A_{32}) \wedge \neg(A_{22} \wedge A_{32})) \end{aligned}$$

Figure 7: Examples of tautologies

to having a solution to the problem at hand. Examples are summarised in Figure 2.1.

It is therefore an important practical problem to determine whether a given propositional formula is a tautology or not and (equivalently) whether or not it is satisfiable.

A formula with 100 atoms admits ca. 10^{33} different valuations so checking tautologies by examining truth tables may be unfeasible.

- if Mary likes champagne then either Bob or Alice like red wine...
- Planning problems: find sequence of actions for a robot to—say—remove a certain item from a stockpile
- Behaviour of digital hardware circuits
- Combinatorial optimisation (scheduling, routing,...)

Figure 8: Applications of propositional logic

Sequents: $\Gamma \Longrightarrow \Delta$ where $\Gamma = \phi_1, \dots, \phi_m$ and $\Delta = \psi_1, \dots, \psi_n$ are lists of formulas.

Meaning: $\phi_1 \wedge \dots \wedge \phi_m \Rightarrow \psi_1 \vee \dots \vee \psi_n$

Examples: $A \Rightarrow B, C \Rightarrow D \Longrightarrow U \wedge (\neg V \wedge B)$

$A \Rightarrow B, A \Longrightarrow B$

$\Longrightarrow A, A \Rightarrow B$

$A, \neg A \Longrightarrow$

\Longrightarrow

Figure 9: Syntax of sequents

While no method is known to date which would be inherently better than checking truth tables there has been considerable progress in the last years at solving instances arising from practical problems (SAT solvers). These solvers vastly outperform any human logician trying to attack propositional formulas by logical reasoning! So why should we look at axioms and logical rules for propositional calculus?

The answer is that in many situations the atoms will themselves be complex formulas, typically a defined predicate applied to some variables, such as $x \in a$ or $\text{sorted}(\text{list}_1)$ and we want to be able to break down the validity of a formula involving these atoms into basic implications between them.

This is precisely the goal of sequent calculus which we will now describe.

2.3 Sequents

A *sequent* is an expression of the form $\Gamma \Longrightarrow \Delta$ where Γ, Δ are (possibly) empty lists of formulas.

The *meaning* of a sequent $\Gamma \Longrightarrow \Delta$ is defined as the meaning of the formula $\bigwedge \Gamma \Rightarrow \bigvee \Delta$ where $\bigwedge \Gamma$ is the conjunction (“and”, \wedge) of the formulas in Γ and $\bigvee \Delta$ is the disjunction (“or”, \vee) of the formulas in Δ .

For example, our formula $(A \Rightarrow B) \wedge \neg A \Rightarrow C$ is equivalent to the sequent $A \Rightarrow B, \neg A \Longrightarrow C$.

A *proof* in sequent calculus is a tree labelled with sequents such that the leaves are labelled with *axioms* and the label of an internal node is the conclusion of a *rule* which has the labels of its immediate descendants as premises.

A sequent $\Gamma \Longrightarrow \Delta$ is an *axiom* if Γ and Δ have a formula in common. For example, the sequent $A, B \Longrightarrow A, C$ is an axiom.

The *rules* for sequent calculus are best read backwards, i.e. “what do I need to

$$\begin{array}{c}
\frac{\Gamma_1, \phi, \psi, \Gamma_2 \Longrightarrow \Delta}{\Gamma_1, \psi, \phi, \Gamma_2 \Longrightarrow \Delta} \quad (\text{PERM-L}) \\
\frac{\Gamma \Longrightarrow \Delta_1, \phi, \psi, \Delta_2}{\Gamma \Longrightarrow \Delta_1, \psi, \phi, \Delta_2} \quad (\text{PERM-R}) \\
\frac{\Gamma \Longrightarrow \Delta}{\Gamma, \phi \Longrightarrow \Delta} \quad (\text{WEAK-L}) \\
\frac{\Gamma \Longrightarrow \Delta}{\Gamma \Longrightarrow \Delta, \phi} \quad (\text{WEAK-R}) \\
\frac{\Gamma, \phi, \phi \Longrightarrow \Delta}{\Gamma, \phi \Longrightarrow \Delta} \quad (\text{CONTR-L}) \\
\frac{\Gamma \Longrightarrow \Delta, \phi, \phi}{\Gamma \Longrightarrow \Delta, \phi} \quad (\text{CONTR-R})
\end{array}$$

Figure 10: Structural rules

prove in order to establish a sequent $\Gamma \Longrightarrow \Delta$?”. First, we have *structural rules* allowing to permute, duplicate or remove formulas. In the literature sequents are often defined as pairs of *sets* rather than lists of formulas. This makes all the structural rules except WEAK redundant. For computer-aided formal reasoning it is, however, useful to have explicit access to formulas, say by their position in a list.

For each connective there are two *logical rules*, one when the connective appears on the left, and one when it appears on the right.

To understand, e.g., \vee -L try to think as follows: to prove Δ under the assumption $\phi \vee \psi$ (and some other stuff Γ) we must make a case distinction as to whether ϕ or ψ holds, hence we must prove Δ under assumptions Γ, ϕ and then again under assumptions Γ, ψ .

Rule \Rightarrow -R is probably the easiest of these: to prove $\phi \Rightarrow \psi$ we must prove ψ under the additional assumption ϕ (if we disregard the side formulas Γ, Δ, \dots). If we forget about Δ we can also explain \neg -R: to prove $\neg\phi$ we must derive a contradiction (empty Δ) from the assumption ϕ .

The \neg -L rule says: if $\neg\phi$ is among our assumptions then to prove anything (Δ) its enough to prove ϕ (or Δ straightaway, of course). This is known as *ex*

$$\frac{\Gamma, \phi, \psi \Longrightarrow \Delta}{\Gamma, \phi \wedge \psi \Longrightarrow \Delta} \quad (\wedge\text{-L})$$

$$\frac{\Gamma \Longrightarrow \Delta, \phi \quad \Gamma \Longrightarrow \Delta, \psi}{\Gamma \Longrightarrow \Delta, \phi \wedge \psi} \quad (\wedge\text{-R})$$

$$\frac{\Gamma, \phi \Longrightarrow \Delta \quad \Gamma, \psi \Longrightarrow \Delta}{\Gamma, \phi \vee \psi \Longrightarrow \Delta} \quad (\vee\text{-L})$$

$$\frac{\Gamma \Longrightarrow \Delta, \phi, \psi}{\Gamma \Longrightarrow \Delta, \phi \vee \psi} \quad (\vee\text{-R})$$

$$\frac{\Gamma \Longrightarrow \Delta, \phi}{\Gamma, \neg\phi \Longrightarrow \Delta} \quad (\neg\text{-L})$$

$$\frac{\Gamma, \phi \Longrightarrow \Delta}{\Gamma \Longrightarrow \Delta, \neg\phi} \quad (\neg\text{-R})$$

$$\frac{\Gamma \Longrightarrow \Delta, \phi \quad \Gamma, \psi \Longrightarrow \Delta}{\Gamma, \phi \Rightarrow \psi \Longrightarrow \Delta} \quad (\Rightarrow\text{-L})$$

$$\frac{\Gamma, \phi \Longrightarrow \psi, \Delta}{\Gamma \Longrightarrow \Delta, \phi \Rightarrow \psi} \quad (\Rightarrow\text{-R})$$

$$\begin{array}{c}
\frac{A \implies B, A \quad B, A \implies B}{A \implies B, A \implies B} \implies\text{-L} \\
\frac{A \implies B, A \implies B}{A \implies B \implies B, \neg A} \neg\text{-L} \\
\frac{A \implies B \implies B, \neg A}{A \implies B, \neg B \implies \neg A} \neg\text{-L}
\end{array}$$

Figure 11: Example proof

$$\frac{\frac{A, B \implies A, A \wedge C \quad A, B \implies B, A \wedge C}{A, B \implies A \wedge B, A \wedge C} \quad \frac{A, C \implies A \wedge B, A \quad A, C \implies A \wedge B, C}{A, C \implies A \wedge B, A \wedge C}}{A, (B \vee C) \implies A \wedge B, A \wedge C}$$

Figure 12: Example proof

falso quodlibet.

Rule $\implies\text{-L}$, finally, says: if we want to use the assumption $\phi \implies \psi$ then we can add its conclusion (ψ) to our assumptions provided we succeed (independently) in proving its antecedent (ϕ).

One can derive the implication rules from the encoding of $\phi \implies \psi$ as $\neg\phi \vee \psi$.

2.4 Soundness and completeness

Definition: A sequent is *derivable* if there exists a proof with it as root label.

Theorem: A sequent is derivable if and only if it is a tautology.

Proof: Let us call a *proof tree* a tree whose nodes (and leaves) are labelled with sequents in such a way that whenever a node labelled S has immediate ancestors S_1, \dots, S_n then there is a logical rule with S_1, \dots, S_n as assumptions as S as conclusion. For example, $S_1 = A \implies B$ and $S_2 = A \implies C$ and $S = A \implies B \wedge C$. Given the form of our rules we always have $n = 1$ or $n = 2$. The leaves may but do not need to be labelled with axioms. Let us write $S_1, \dots, S_\ell \vdash S$ to mean that there is a proof tree whose root is labelled S and whose leaves are labelled with S_1, \dots, S_ℓ .

Notice that a *proof* of sequent S is a proof tree all whose leaves are labelled with axioms.

By induction on (depth of) proof trees one easily shows that if $S_1, \dots, S_n \vdash S$

and S_1, \dots, S_n are all true under some valuation η then S , too, comes out true under η . Recall that the truth value of a sequent $S = \phi_1, \dots, \phi_m \Longrightarrow \psi_1, \dots, \psi_n$ under some valuation η is defined as $\bigwedge_i \eta(\phi_i) \Rightarrow \bigvee_j \eta(\psi_j)$, i.e., S comes out false precisely if all the ϕ_i come out true and all the ψ_j come out false.

The above proves that if a sequent has a proof, i.e., a proof tree with axioms labelling its leaves, then it is tautologous, i.e., true under all valuations.

For completeness we first notice (again by induction on proof trees) that if $S_1, \dots, S_n \vdash S$ is proved by a proof tree not involving rule WEAK then S is equivalent to the conjunction of the S_i , i.e., S comes out false under some valuation η as soon as one of the S_i is falsified by η . Now, for any sequent S we can always find a proof tree not involving rule WEAK whose root is labelled S and whose leaves are labelled with sequents consisting of atoms only. This is done by successively “breaking down” all the connectives in S . If S is *not derivable* then at least one of the atomic sequents labelling the leaves of this proof tree will not be an axiom (otherwise our proof tree would be a proof!). This sequent will thus be of the form $A_1, \dots, A_m \Longrightarrow B_1, \dots, B_n$ where the A_i and B_j are atoms and $\{A_1, \dots, A_m\} \cap \{B_1, \dots, B_n\} = \emptyset$. Any valuation η with $\eta(A_i) = \text{tt}$, $\eta(B_j) = \text{ff}$ will falsify this sequent, hence S . So S is not a tautology. \square

We remark that this also shows that if a sequent is derivable with rules CONTR and WEAK then it is provable without those rules; the “generic” proof tree obtained by breaking down the connectives must lead to a proof in this case, otherwise we would obtain a falsifying valuation.

2.4.1 Linear logic

The fact that rules WEAK and CONTR can be eliminated is due to their being built into the other rules and axioms. In *linear logic* weakening and contraction are removed and the side formulas in different premises of a rule are required to be disjoint. For instance, a linear version of \wedge -R would look thus:

$$\frac{\Gamma_1 \Longrightarrow \Delta_1, \phi \quad \Gamma_2 \Longrightarrow \Delta_2, \psi}{\Gamma_1, \Gamma_2 \Longrightarrow \Delta_1, \Delta_2, \phi \wedge \psi} \quad (\wedge\text{-R-LIN})$$

Moreover, the only axioms are $A \Longrightarrow A$ for A an atom. If we replace all rules and axioms by their linear versions and remove WEAK and CONTR (which now are no longer redundant) we obtain a (a fragment of) linear logic in which, e.g., the formula $A \Rightarrow A \wedge A$ is not provable.

2.4.2 Cut elimination

In spite of completeness it is useful to have yet another rule: the famous cut rule:

$$\frac{\Gamma_1 \Longrightarrow \Delta_1, \phi \quad \Gamma_2, \phi \Longrightarrow \Delta_2}{\Gamma_1, \Gamma_2 \Longrightarrow \Delta_1, \Delta_2} \quad (\text{CUT})$$

Usually, Γ_1 and Δ_1 are empty. In this case, CUT corresponds to invocation of a lemma: if we have proved ϕ “as a lemma” then we can add it to our list of assumptions whenever we want.

Of course, in view of completeness the cut rule is redundant and there is even a procedure that systematically eliminates cuts from a proof containing instances of CUT (“cut elimination”). The size of the proof may grow considerably upon elimination of cuts.

3 Introduction to PVS

As already mentioned, PVS is based on the sequent calculus. We will start by using PVS as a proof-assistant for that system.

To do that we call PVS from the (Linux!) command line with

```
pvs
```

This brings up an Emacs window entitled PVS. A possible question concerning context creation should be answered affirmatively (means: type in *yes*).

Now create a file named `sequent.pvs` containing something like

```
sequent_calculus: THEORY
BEGIN

A, B, C, D: boolean
A11, A12, A21, A22, A31, A32: boolean

K : PROPOSITION
  A IMPLIES B IMPLIES A

S : PROPOSITION
  (A IMPLIES B IMPLIES C) IMPLIES (A IMPLIES B)
  IMPLIES (A IMPLIES C)
```

```

Peirce : PROPOSITION
  ((A IMPLIES B) IMPLIES A) IMPLIES A

Contra : PROPOSITION
  (A IMPLIES B) IMPLIES NOT B IMPLIES NOT A

dist1 : PROPOSITION
  A AND (B OR C) IMPLIES (A AND B) OR (A AND C)

dist2 : PROPOSITION
  A AND B OR C AND D IMPLIES (A OR C) AND (B OR D)

schub : PROPOSITION
  NOT (
    (A11 OR A12) AND
    (A21 OR A22) AND
    (A31 OR A32) AND
    NOT (A11 AND A21) AND
    NOT (A11 AND A31) AND
    NOT (A21 AND A31) AND
    NOT (A12 AND A22) AND
    NOT (A12 AND A32) AND
    NOT (A22 AND A32) )

END sequent_calculus

```

Click on the first “proposition”. This will bring up a prover window and a prompt Rule?. The first rule you should enter is (skolem!). This will get rid of the FORALL quantifier which we’ll talk about later and basically display the formula as a sequent with empty premise list and one conclusion named [1]. Unfortunately, all the atoms are decorated with !1. ; Now you can

- apply a disjunctive rule (\wedge -L, \vee -R, \Rightarrow -R) by entering (flatten x) where x is the number of the formula you want to apply the rule to. (x must obviously be 1 at the beginning.
- apply a conjunctive rule (\wedge -R, \vee -L, \Rightarrow -L) by entering (split x) where, again, x is the number of the formula.

The rules for negation are applied automatically.

Try to prove all the propositions in this way but don't waste too much time on the last one (`schub`).

Rather give up after a few steps by typing (`quit`) and redo the proof (answering no when asked whether you want to rerun the existing proof) this time typing (`prop`) after the (`skolem!`) step or (`grind`) right at the start.

You can display the tree structure of your current proof with `M-x x-show-current-proof` and of a finished proof with `M-x x-show-proof`.

You may wish to do some more ad-hoc experiments with PVS. Extensive documentation is available at the PVS homepage pvs.csl.sri.com.

4 Resolution

Gentzen's sequent calculus provides a decision procedure for the validity of propositional formulas: construct the proof tree as in the completeness proof and check whether all leaves are labelled with axioms.

Unfortunately, the complexity of this procedure is exponential in the size of the formula to be proved. This is due to the duplication of "goals" in rules \vee -L, \wedge -R, \Rightarrow -L. Worse still, a lot of work is done twice: if we break down $\Gamma \Rightarrow \Delta, \phi \wedge \psi$ into $\Gamma \Rightarrow \Delta, \phi$ and $\Gamma \Rightarrow \Delta, \psi$ then the "breaking down" of Γ and Δ must be done in each branch individually.

Of course, unless $P = NP$ we cannot expect a really efficient (polynomial) method for deciding propositional formulas; however, there are algorithms that behave quite well in practice for moderately sized formulas. One of these is the method of *resolution* invented by ROBINSON which we will take a look at for one thing for its popularity and for another because its applicability to first-order logic which we will come to shortly.

Resolution is a method for deciding *satisfiability* of a propositional formula presented as a *set of clauses*. Recall that a formula ϕ is satisfiable if $\neg\phi$ is not a tautology. A *literal* is either an atom or a negated atom, e.g., A , $\neg B$, `door_closed`, `\neg alarm_on` are all literals. A *clause* is a set of literals whose meaning is their disjunction. Some people write a clause as an explicit disjunction using \vee , others use set notation. A *set of clauses* represents the conjunction of the individual clauses. Watch out for the empty clause (representing `ff`) and the empty set of clauses (representing `tt`).

To check whether a formula ϕ is a tautology we can represent $\neg\phi$ as a set of clauses and see whether it is not satisfiable.

- a *literal* is either an atom or a negated atom: $A, \neg B, \neg \text{door_closed}$.
- a *clause* is a set of literals understood as their *disjunction*: $\{\neg \text{lift_moves}, \text{door_closed}, \text{alarm_on}\}$
- a *set of clauses* is understood as the conjunction of the individual clauses.
- empty clause = ff, empty set of clauses = tt
- One is interested in *satisfiability* of sets of clauses.
- Validity (to be a tautology) is trivial for sets of clauses. Why?

Figure 13: Clauses

4.1 Representation of formulas as sets of clauses

We all know that any formula ϕ can be converted into conjunctive normal form, by “multiplying out” according to de Morgan’s rule. The problem with this is that in general the size of the conjunctive normal form will be exponential in the size of the formula to start with. Actually, if this blow up would not occur we had a simple method for checking whether ϕ is a tautology. Just bring ϕ into conjunctive normal form and see whether each clause is a tautology.

What we can do without exponential blowup, though, is to construct a set of clauses \mathcal{C} which is *satisfiable* if and only if ϕ is. To that end, we proceed as follows: first, we may assume that ϕ contains connectives \vee, \wedge, \neg only and that, moreover, \neg occurs in front of atoms only. One calls this the negation normal form of ϕ . Now, if ϕ happens to be a literal then there is nothing to do. If $\phi = \phi_1 \wedge \phi_2$ and ϕ_1, ϕ_2 are equi-satisfiable with $\mathcal{C}_1, \mathcal{C}_2$, respectively, then $\mathcal{C}_1 \cup \mathcal{C}_2$ is equi-satisfiable with ϕ . If, finally, $\phi = \phi_1 \vee \phi_2$ then $(\mathcal{C}_1 \vee P) \cup (\mathcal{C}_2 \vee \neg P)$ is equi-satisfiable with ϕ . Here P is a fresh atom and $\mathcal{C} \vee P$ means the addition of P to each clause in \mathcal{C} .

4.2 The method of resolution

The method of resolution decides whether a given set of clauses is satisfiable. It works as follows: given two clauses C_1 and C_2 such that C_1 contains some literal ℓ and C_2 contains its negation $\neg \ell$ (with the understanding that the negation of $\neg A$ is A) then the *rule of resolution* applied to C_1, C_2 yields the clause $C_1 \setminus \{\ell\} \cup C_2 \setminus \{\neg \ell\}$. For example, applying the rule of resolution to $\{A, \neg B, D\}$ and

Rule of resolution:

$$\frac{C_1 \cup \{\ell\} \quad C_2 \cup \{\neg\ell\}}{C_1 \cup C_2}$$

Example: $\{\neg A, B, D\}$ and $\{\neg B, X\}$ yield $\{\neg A, D, X\}$.

Figure 14: The rule of resolution

```

INPUT: a set of clauses  $\mathcal{C}$ 
WHILE  $\emptyset \notin \mathcal{C}$  OR  $\mathcal{C}$  still grows DO
  CHOOSE  $C_1, C_2 \in \mathcal{C}$  S. T.  $\ell \in C_1$  and  $\neg\ell \in C_2$  for some literal  $\ell$ .
   $\mathcal{C} := \mathcal{C} \cup \{C_1 \setminus \{\ell\} \cup C_2 \setminus \{\neg\ell\}\}$ 
IF  $\emptyset \in \mathcal{C}$  THEN OUTPUT “ $\mathcal{C}$  is not satisfiable”
ELSE OUTPUT “ $\mathcal{C}$  is satisfiable”

```

Figure 15: The method of resolution

$\{A, \neg D, E\}$ yields $\{A, \neg B, E\}$. The method of resolution consists of closing up a set of clauses under the rule of resolution and seeing whether the so closed-up set contains the empty clause or not. This is formalised in Fig.15. Note that if the initial clause set \mathcal{C} is finite then the algorithm terminates since there is only a finite number of possible clauses over any given (finite) set of variables.

4.3 Correctness of resolution

If the set of clauses—after this closure—contains the empty clause then it is unsatisfiable, otherwise we can find a satisfying valuation. This is the content of the correctness theorem for resolution.

A set of clauses \mathcal{C} is *closed under resolution* if the result of applying the rule of resolution to any two clauses in \mathcal{C} is already contained in \mathcal{C} . The above method precisely computes the closure under resolution of an arbitrary set of clauses.

Let \mathcal{C} be a set of clauses, A an atom. We define the clause set $\mathcal{C}[A \mapsto \text{tt}]$ by removing from \mathcal{C} every clause that contains the literal A and removing the literal $\neg A$ from every clause that contains it. Analogously, we define $\mathcal{C}[A \mapsto \text{ff}]$. It is easily seen by case distinction that if \mathcal{C} is closed under resolution so are these two sets.

If η is a valuation that satisfies $\mathcal{C}[A \mapsto v]$ then the valuation $\eta[A \mapsto v]$ which maps A to v and all other atoms according to η will satisfy \mathcal{C} .

Theorem: Let \mathcal{C} be a possibly infinite set of clauses closed under resolution,

i.e., Then \mathcal{C} is satisfiable if and only if $\emptyset \notin \mathcal{C}$.

Proof: If a set of clauses \mathcal{C}' has been obtained from a *satisfiable* set of clauses \mathcal{C} by a single application of the rule of resolution then \mathcal{C}' is satisfiable, too. One says that the rule of resolution preserves satisfiability. Thus, if we derive from \mathcal{C} a set of clauses containing the empty clause then \mathcal{C} must have been unsatisfiable in the first place.

For the converse, suppose that \mathcal{C} does not contain the empty clause yet is closed under the rule of resolution. We explicitly construct a valuation η that will satisfy all the clauses in \mathcal{C} : Enumerate the atoms as A_1, A_2, \dots . We define the values $\eta(A_1), \eta(A_2), \dots$ in order. Let us begin with the variable A_1 . Not both $\mathcal{C}[A_1 \mapsto \text{tt}]$ and $\mathcal{C}[A_1 \mapsto \text{ff}]$ can contain the empty clause for otherwise, \mathcal{C} would contain both $\{A_1\}$ and $\{\neg A_1\}$ and hence the empty clause by one resolution step. Thus, choose $\eta(A_1)$ such that $\mathcal{C}[A_1 \mapsto \eta(A_1)]$ does not contain the empty clause. We continue in this way replacing A_1 with A_2 and \mathcal{C} with $\mathcal{C}[A_1 \mapsto \eta(A_1)]$ yielding a value $\eta(A_2)$ such that $\mathcal{C}[A_1 \mapsto \eta(A_1)][A_2 \mapsto \eta(A_2)]$ does not contain the empty clause.

Continuing in this way, we obtain a valuation η which satisfies all the clauses in \mathcal{C} . This can be seen by noticing that as the variables are considered every clause in \mathcal{C} eventually disappears.

An important corollary is the compactness theorem for propositional logic:

Theorem: Let \mathcal{C} be a (possibly infinite) set of clauses or propositional formulas. If every finite subset of \mathcal{C} is satisfiable then the whole of \mathcal{C} is satisfiable.

Proof: If contrary to the conclusion the whole of \mathcal{C} is unsatisfiable then by the previous theorem it must be possible to deduce the empty clause from \mathcal{C} . But such a proof will only involve a finite portion of \mathcal{C} which would then already be unsatisfiable contradicting the assumption.

4.4 Long resolution proofs

While for many tautologies (or rather their negations) resolution works astonishingly fast there are other ones, e.g., `schub` above for which it is rather slow. Indeed, HAKEN has shown that the resolution method has exponential worst case complexity.

Theorem (Haken): There is a constant $c > 1$ and an infinite family of (unsatisfiable) sets of clauses P_1, P_2, \dots such that P_n consists of $O(n^2)$ clauses yet any derivation of \emptyset from P_n will involve $O(c^n)$ many clauses.

To wit, the set of clauses P_n expresses that $n + 1$ pigeons fit into n holes, e.g., $P_2 \Leftrightarrow \neg \text{schub}$.

The proof of Haken's theorem is elementary but fairly long and technical. Recently, WIGDERSON has presented an important simplification. Use Google to find his paper if you are interested.

4.5 Cook's programme

Haken's theorem is quite drastic evidence for exponential time complexity of the resolution procedure. Even if Haken's theorem would not hold as stated, resolution could still fail to be a polynomial time procedure: it might be difficult to find a polynomially sized derivation of \emptyset even if it exists and in the case of *satisfiable* sets of clauses the process of closing up might result in exponentially many clauses.

It is an important open complexity-theoretic question related to $P=NP$ as to whether there is a proof system for propositional logic with the property that any tautology ϕ has a proof of size polynomial in the size of ϕ . Refuting this for concrete proof systems (as Haken has done for resolution) is a popular research activity initiated by S. COOK. As far as I know, it is presently not known whether or not each tautology has a polynomially sized proof in sequent calculus with the CUT rule.

4.6 The DPLL procedure

In practical implementations of decision procedures for propositional logic (SAT-solvers) resolution has been superseded by a surprisingly naive search procedure known as DPLL algorithm (Davis-Putnam-Loveland-Logemann).

It also operates on clause sets and is based on three interleaved steps

- Unit propagation: if the clause set contains a unit clause, i.e., one containing a single literal, then it is possible to set the value of that atom and propagate it through the other clauses. I.e., if there is a clause consisting of A alone, remove that clause and all clauses containing the literal A . Remove $\neg A$ from all clauses containing it. Similarly for $\neg A$. Continue with further unit clauses so obtained.
- Branching: choose an arbitrary atom A and try to satisfy first $\mathcal{C}[A \mapsto \text{tt}]$ and $\mathcal{C}[A \mapsto \text{ff}]$. If either turns out to be satisfiable then so is \mathcal{C} . Otherwise, \mathcal{C} is unsatisfiable.

- Learning clauses: if during the branching it turns out that for some partial valuation η the clause set $\mathcal{C}[\eta]$ is unsatisfiable by unit propagation alone then we can identify those settings in η which lead to the empty clause and build a corresponding clause k that can be added to \mathcal{C} without affecting satisfiability. If e.g., we find that setting $A \mapsto \text{tt}, B \mapsto \text{ff}, C \mapsto \text{tt}$ leads to a contradiction (empty clause) then we can add the clause $\{\neg A, B, \neg C\}$ to \mathcal{C} . The hope is that in the future this addition will speed up unit propagation; in particular, if at a later stage we try to, again, set $A \mapsto \text{tt}, B \mapsto \text{ff}, C \mapsto \text{tt}$ then we find out immediately that this leads to a contradiction rather than having to re-perform the corresponding steps of unit propagation.

Slightly more formally DPLL can be viewed as a recursive procedure that uses a global variable containing a set of clauses \mathcal{C} . It is an invariant that clauses are never removed from \mathcal{C} and whenever a clause k is to be added to \mathcal{C} then it is a logical consequence of \mathcal{C} , i.e., whenever η satisfies \mathcal{C} then it also satisfies k .

The procedure takes as argument a partial valuation η and $DPLL(\eta)$ returns “satisfiable” if there exists a valuation extending η that satisfies \mathcal{C} ; it returns “unsatisfiable”, if no such valuation exists. Figure 4.6 contains pseudo code for DPLL. In a practical implementation a number of improvements are possible. Rather than using recursion one can maintain a stack of partial environments. Furthermore, it is not necessary to compute the clause set $\mathcal{C}[\eta]$ explicitly. For details, we refer to the literature.

5 First-order logic

Propositional calculus is nice, but in many applications we need a way of talking about elements, predicates, and operations. That is what first order logic is for. Figure 5 shows some examples of sentences that lend themselves to a formalisation in first-order logic. A formal system for writing down such statements is obtained by augmenting propositional logic with the *quantifiers* \forall (for all) and \exists (there exists). As to the range of these quantifiers one has two options which we consider in order.

Untyped first-order logic. This is the traditional and most common version of first-order logic. We let all quantifiers and variables range over one and the same implicit, a priori given, domain. In this case we must use special predicates to


```

UNITPROP( $\mathcal{C}, \eta$ ) =
  IF  $\mathcal{C}[\eta]$  contains the empty clause
    let  $\rho \subseteq \eta$  be the smallest sub-valuation of  $\eta$  such that  $\emptyset \in \mathcal{C}[\rho]$ 
    RETURN ( $\rho$ , "unsatisfiable")
  ELSE IF  $\mathcal{C}[\eta]$  is empty
    RETURN ( $\eta$ , "satisfiable")
  ELSE IF  $\mathcal{C}[\eta]$  contains a unit clause  $\{A\}$ 
    RETURN UNITPROP( $\mathcal{C}, \eta[A \mapsto \text{tt}]$ )
  ELSE IF  $\mathcal{C}[\eta]$  contains a unit clause  $\{\neg A\}$ 
    RETURN UNITPROP( $\mathcal{C}, \eta[A \mapsto \text{ff}]$ )
  ELSE RETURN ( $\eta$ , "undecided")

DPLL( $\eta$ ) =
  ( $\eta', v$ ) := UNITPROP( $\mathcal{C}, \eta$ )
  IF  $v = \text{"satisfiable"}$ 
    RETURN "satisfiable"
  ELSE IF  $v = \text{"unsatisfiable"}$ 
     $\mathcal{C} := \mathcal{C} \cup \{k\}$  where  $k$  asserts that at least one atom is valued different from  $\eta'$ .
    RETURN "unsatisfiable"
  ELSE
    choose an atom  $A \notin \text{dom}(\eta')$ 
    IF DPLL( $\eta'[A \mapsto \text{tt}]$ ) = "satisfiable"
      RETURN "satisfiable".
    ELSE
      RETURN DPLL( $\eta'[A \mapsto \text{ff}]$ )

INPUT a set of clauses  $\mathcal{X}$ 
 $\mathcal{C} := \mathcal{X}$ 
OUTPUT DPLL( $\emptyset$ )

```

Figure 16: DPLL-algorithm with clause learning

1. Every student has a matric number.
2. If a student fails to matriculate she will be expelled.
3. Every human being is a philosopher.
4. There always is one student who complains about every course.
5. There is a set with no elements.
6. For every natural number n there exists a natural number d such that $2^d \leq n$ and $n < 2^{d+1}$.
7. The knirp of each bilg is a prugl but does not bebelf any quist.

Figure 17: First-order formulas (informal)

restrict the range of quantifiers. Figure 18 shows how the example sentences are formalised in untyped first-order logic.

Typed first-order logic. Alternatively, we fix a collection of *types* and require that each quantifier is annotated with a type determining its range. Assuming that we have fixed types

human, student, number, course, set, bilg, quist

we could write the example sentences as in Figure 5. Untyped first-order logic has the advantage of being slightly simpler to formulate and it suffices for many applications, especially in mathematics. Typed first-order logic has the advantage of being more readable and, more importantly, that it allows the user to distinguish between actual properties he wants to prove and typing judgments which should follow automatically in most cases. For example, if we were to prove formula 6 above in the untyped setting then we would at some point come up with a d and would then have to prove $\text{number}(d)$ as well as the actually interesting property of d .

In the typed setting the first one falls under typing and can often be discharged automatically.

Of course, the distinction between types and predicates is a subjective one and can be “misused”.

1. $\forall x.\text{student}(x) \Rightarrow \exists n.\text{number}(x) \wedge \text{has_matric_no}(x, n)$
2. $\forall x.\text{student}(x) \Rightarrow \neg\text{has_matriculated}(x) \Rightarrow \text{will_be_expelled}(x)$
3. $\forall x.\text{human}(x) \Rightarrow \text{philosopher}(x)$
4. $\exists x.\text{student}(x) \wedge \forall c.\text{course}(c) \Rightarrow \text{complains_about}(x, c)$
5. $\exists x.\forall y.\neg(y \in x)$
6. $\forall n.\text{number}(n) \Rightarrow \exists d.\text{number}(d) \wedge \text{leq}(\text{power}(2, d), n) \wedge \text{lt}(n, \text{power}(2, \text{plus}(d, 1)))$
7. $\forall x.\text{bilg}(x) \Rightarrow \text{prugl}(\text{knirp}(x)) \wedge \forall y.\text{quist}(y) \Rightarrow \neg\text{bebelfs}(\text{knirp}(x), y)$

Figure 18: Formalisation in untyped first-order logic

1. $\forall x:\text{student}.\exists n:\text{number}.\text{has_matric_no}(x, n)$
2. $\forall x:\text{student} \Rightarrow \neg\text{has_matriculated}(x) \Rightarrow \text{will_be_expelled}(x)$
3. $\forall x:\text{human}.\text{philosopher}(x)$
4. $\exists x:\text{student}.\forall c:\text{course}.\text{complains_about}(x, c)$
5. $\exists x:\text{set}.\forall y:\text{set}.\neg(y \in x)$
6. $\forall n:\text{number}.\exists d:\text{number}.\text{leq}(\text{power}(2, d), n) \wedge \text{lt}(\text{power}(2, \text{plus}(d, 1)))$
7. $\forall x:\text{bilg}.\text{prugl}(\text{knirp}(x)) \wedge \forall y:\text{quist}.\neg\text{bebelfs}(\text{knirp}(x), y)$

Figure 19: Typed first-order logic

Typechecking can be difficult. For example, if D is the “type” consisting of quadruples of integers (x, y, z, n) such that $n > 2$ and $x^n + y^n = z^n$ then proving

$$\exists p: D.p = p$$

is tantamount to proving Wiles’ theorem!

If used reasonably then types can considerably simplify (formal) proofs and the appearance of statements.

5.1 Typed first-order language

A typed first-order language is specified by the following data:

1. a collection \mathcal{T} of *types*
2. a collection \mathcal{P} of *predicate constants*, each endowed with an *arity* $[\tau_1, \dots, \tau_n \rightarrow \text{boolean}]$ where $\tau_1, \dots, \tau_n \in \mathcal{T}$
3. a collection \mathcal{F} of *function constants* each endowed with an *arity* $[\tau_1, \dots, \tau_n \rightarrow \tau_{n+1}]$ where $\tau_1, \dots, \tau_{n+1} \in \mathcal{T}$

The arity of a predicate constant is nothing but a list of types; the brackets, the arrow, and “boolean” are merely notation. We use the notation $P : [\tau_1, \dots, \tau_n \rightarrow \text{boolean}]$ and $f : [\tau_1, \dots, \tau_n \rightarrow \tau_{n+1}]$ to indicate the arities of predicate, resp. function constants.

Examples: For the formulas 1, ..., 4 an appropriate language is as follows:

$$\begin{aligned} \mathcal{T} &= \{\text{student, number, course}\} \\ \mathcal{P} &= \{\text{has_matric_no} : [\text{student, number} \rightarrow \text{boolean}] \\ &\quad \text{has_matriculated} : [\text{student} \rightarrow \text{boolean}] \\ &\quad \text{will_be_expelled} : [\text{student} \rightarrow \text{boolean}]] \\ &\quad \text{complains_about} : [\text{student, course} \rightarrow \text{boolean}]\} \\ \mathcal{F} &= \emptyset \end{aligned}$$

For the formula 7 an appropriate first-order language would be

$$\begin{aligned} \mathcal{T} &= \{\text{bilg, quist, knirp_t}\} \\ \mathcal{P} &= \{\text{prugl} : [\text{knirp_t} \rightarrow \text{boolean}], \text{bebelfs} : [\text{knirp_t, quist} \rightarrow \text{boolean}]\} \\ \mathcal{F} &= \{\text{knirp} : [\text{bilg} \rightarrow \text{knirp_t}]\} \end{aligned}$$

Typing contexts: $K = x_1:\tau_1, \dots, x_n:\tau_n$; formally: finite function from variables to types.

Typing rules:

$$\frac{K(x) = \tau}{K \triangleright_{\mathcal{L}} x : \tau} \quad (\text{VAR})$$

$$\frac{f : [\tau_1, \dots, \tau_n \rightarrow \tau_{n+1}] \in \mathcal{F} \quad K \triangleright_{\mathcal{L}} t_1 : \tau_1, \dots, K \triangleright_{\mathcal{L}} t_n : \tau_n}{K \triangleright_{\mathcal{L}} f(t_1, \dots, t_n) : \tau_{n+1}} \quad (\text{FUN})$$

Figure 20: Well formed terms

The type `knirp_t` arises only as the range of a function symbol not as the range of a quantifier. When translating this formula from natural language to formal language we could also have opted for a predicate constant `knirp` : [`bilg`, `knirp_t` \rightarrow boolean] denoting its graph. The content of the definite article “*the* `knirp` of . . .” could be rendered by another formula stating unique existence. This would require the notion of equality which we will come to later.

5.2 Syntax

In order to define properly what formulas are we have to talk about terms and formulas possibly involving variables as those occur in scopes of quantifiers. We thus assume an infinite set \mathcal{V} of variables distinct from the other symbols and fix a first order language \mathcal{L} . A *typing context* is a finite partial function K mapping variables to types. If K is a typing context and $x \notin \text{dom}(K)$ and $\tau \in \mathcal{T}$ then $K, x:\tau$ is the typing context K extended with $x \mapsto \tau$.

We write $K \triangleright_{\mathcal{L}} t : \tau$ to mean that t is a well formed term in \mathcal{L} of type τ possibly involving the variables declared and typed as given by K . This judgment is inductively defined by the following *typing rules*.

Well-formed formulas are built up from atomic formulas (predicate constants applied to appropriately typed terms) by propositional connectives and quantifiers which *bind* variables. Formally, we introduce the judgment $K \triangleright_{\mathcal{L}} \phi : \text{boolean}$ to mean that ϕ is a well formed formula in \mathcal{L} possibly involving the free variables declared and typed in K . This judgment is defined by the following rules:

$$\frac{P : [\tau_1, \dots, \tau_n \rightarrow \text{boolean}] \in \mathcal{P} \quad K \triangleright_{\mathcal{L}} t_1 : \tau_1, \dots, K \triangleright_{\mathcal{L}} t_n : \tau_n}{K \triangleright_{\mathcal{L}} P(t_1, \dots, t_n) : \text{boolean}} \quad (\text{ATOM})$$

$$\begin{array}{c}
\frac{K \triangleright_{\mathcal{L}} \phi : \text{boolean}}{K \triangleright_{\mathcal{L}} \neg \phi : \text{boolean}} \quad (\text{NEG}) \\
\frac{K \triangleright_{\mathcal{L}} \phi : \text{boolean} \quad K \triangleright_{\mathcal{L}} \psi : \text{boolean} \quad \star \in \{\vee, \wedge, \Rightarrow\}}{K \triangleright_{\mathcal{L}} \phi \star \psi : \text{boolean}} \quad (\text{CONN}) \\
\frac{K, x:\tau \triangleright_{\mathcal{L}} \phi : \text{boolean} \quad Q \in \{\forall, \exists\}}{K \triangleright_{\mathcal{L}} Qx:\tau.\phi : \text{boolean}} \quad (\text{QUANT})
\end{array}$$

These rules define *abstract syntax* together with *typing* (there latter is also known and misnamed as “semantic analysis”). For concrete syntax one needs to specify precedence rules and use parentheses to disambiguate otherwise. The connectives take precedence as before; quantifiers always extend as far to the right as possible, i.e., until an unmatched closing parenthesis is encountered.

5.3 Semantics

A formula ϕ is *closed* if it contains no free variables, i.e., if $\emptyset \triangleright_{\mathcal{L}} \phi : \text{boolean}$. These are the ones we are really interested in; the open formulas are introduced only as an auxiliary device for the definition of the closed ones.

The meaning of a closed first-order formula is given as a truth value relative to an interpretation of the types, the predicate constants, and the function constants. To specify the meaning of a non closed formula we also need a valuation of the variables.

Interpretation An interpretation \mathcal{I} of a first-order language $(\mathcal{T}, \mathcal{P}, \mathcal{F})$ is given by

1. a set $\llbracket \tau \rrbracket_{\mathcal{I}}$ for each $\tau \in \mathcal{T}$.
2. a function $\llbracket P \rrbracket_{\mathcal{I}} : \llbracket \tau_1 \rrbracket_{\mathcal{I}} \times \cdots \times \llbracket \tau_n \rrbracket_{\mathcal{I}} \rightarrow \{\text{tt}, \text{ff}\}$ for each $P : [\tau_1, \dots, \tau_n \rightarrow \text{boolean}]$.
3. a function $\llbracket f \rrbracket_{\mathcal{I}} : \llbracket \tau_1 \rrbracket_{\mathcal{I}} \times \cdots \times \llbracket \tau_n \rrbracket_{\mathcal{I}} \rightarrow \llbracket \tau_{n+1} \rrbracket_{\mathcal{I}}$ for each function constant $f : [\tau_1, \dots, \tau_n \rightarrow \tau_{n+1}]$.

Such an interpretation associates a truth value with every closed formula and more generally, a function mapping valuations of variables to truth values with every open formula. For pedants we give here a formal definition.

$$\begin{aligned}
\llbracket x \rrbracket_{\rho, \mathcal{I}} &= \rho(x) \\
\llbracket f(t_1, \dots, t_n) \rrbracket_{\rho, \mathcal{I}} &= \llbracket f \rrbracket_{\mathcal{I}}(\llbracket t_1 \rrbracket_{\rho, \mathcal{I}}, \dots, \llbracket t_n \rrbracket_{\rho, \mathcal{I}}) \\
\llbracket P(t_1, \dots, t_n) \rrbracket_{\rho, \mathcal{I}} &= \llbracket P \rrbracket_{\mathcal{I}}(\llbracket t_1 \rrbracket_{\rho, \mathcal{I}}, \dots, \llbracket t_n \rrbracket_{\rho, \mathcal{I}}) \\
\llbracket \neg \phi \rrbracket_{\rho, \mathcal{I}} &= \neg \llbracket \phi \rrbracket_{\rho, \mathcal{I}} \\
\llbracket \phi \star \psi \rrbracket_{\rho, \mathcal{I}} &= \llbracket \phi \rrbracket_{\rho, \mathcal{I}} \star \llbracket \psi \rrbracket_{\rho, \mathcal{I}} \\
\llbracket \forall x:\tau. \phi \rrbracket_{\rho, \mathcal{I}} &= \begin{cases} \mathbf{tt}, & \text{if } \llbracket \phi \rrbracket_{\rho[x \mapsto v], \mathcal{I}} = \mathbf{tt} \text{ for all } v \in \llbracket \tau \rrbracket_{\mathcal{I}} \\ \mathbf{ff}, & \text{otherwise} \end{cases} \\
\llbracket \exists x:\tau. \phi \rrbracket_{\rho, \mathcal{I}} &= \begin{cases} \mathbf{tt}, & \text{if } \llbracket \phi \rrbracket_{\rho[x \mapsto v], \mathcal{I}} = \mathbf{tt} \text{ for some } v \in \llbracket \tau \rrbracket_{\mathcal{I}} \\ \mathbf{ff}, & \text{otherwise} \end{cases}
\end{aligned}$$

Here ρ is a partial function on variables.

We note that if ρ is compatible with typing context K in the sense that $\rho(x) \in \llbracket K(x) \rrbracket_{\mathcal{I}}$ for all $x \in \text{dom}(K)$, in particular, $\rho(x)$ is defined in this case, then $\llbracket t \rrbracket_{\rho, \mathcal{I}} \in \llbracket \tau \rrbracket_{\mathcal{I}}$ whenever $K \triangleright_{\mathcal{L}} t : \tau$ and $\llbracket \phi \rrbracket_{\rho, \mathcal{I}} \in \{\mathbf{tt}, \mathbf{ff}\}$ whenever $K \triangleright_{\mathcal{L}} t : \text{boolean}$.

A closed formula is *valid* if its meaning comes out as true under all possible interpretations of the language it is based on. Examples of such valid formulas are as follows.

- $\forall x:\tau. P(x) \Rightarrow \exists y:\tau. P(y)$ (recall that quantifiers always extend to the left as far as possible),
- $(\forall x:\tau_1. R(x, f(x))) \Rightarrow \forall x:\tau_1. \exists y:\tau_2. R(x, y)$ when $f : [\tau_1 \rightarrow \tau_2]$,
- $(\forall x:\tau_1. \forall y:\tau_2. P(x) \vee Q(y)) \Rightarrow (\forall x:\tau_1. P(x)) \vee (\forall x:\tau_2. Q(x))$,
- $(Q \vee \exists x:\tau. P(x)) \Rightarrow \exists x:\tau. Q \vee P(x)$ where $Q : [\rightarrow \text{boolean}]$ is a constant and, moreover, we have a constant $c : [\rightarrow \tau]$,
- $\exists x:\tau. P(x) \Rightarrow \forall y:\tau. P(y)$ again in the presence of a constant $c : [\rightarrow \tau]$.

Nullary predicate and function constants are propositional, resp., “ordinary” constants. We may write $Q : \text{boolean}$ and $c:\tau$ instead of $Q : [\rightarrow \text{boolean}]$ and $c : [\rightarrow \tau]$ to declare them and omit empty parentheses (as done above) when using them.

5.4 First-order sequent calculus

As before we form sequents $\Gamma \Longrightarrow_{\mathcal{L}} \Delta$ from lists of *closed* formulas Γ, Δ over some language \mathcal{L} . The meaning of such a sequent is that the conjunction of the formulas in Γ implies the disjunction of the formulas in Δ .

Notice that the presence of constants in \mathcal{L} can affect the meaning of a formula hence of a sequent even if these constants do not occur explicitly. This explains the explicit mentioning of \mathcal{L} .

We introduce the notation $\mathcal{L}, c:\tau$ for the extension of \mathcal{L} with a new constant c of type τ . We keep all the rules for the propositional connectives and add four rules to deal with the quantifiers which we will now explain.

To prove an existential statement we have the rule

$$\frac{\Gamma \Longrightarrow_{\mathcal{L}} \Delta, \phi[t/x]}{\Gamma \Longrightarrow_{\mathcal{L}} \Delta, \exists x:\tau.\phi} \quad (\exists\text{-R})$$

where $\emptyset \triangleright_{\mathcal{L}} t : \tau$.

Here $\phi[t/x]$ denotes the substitution of *closed* term t for variable x in ϕ .

The rule says that to prove an existential statement we must come up with a witness. It corresponds to phrases like “*The desired value x is therefore given by t ...*”. Next, we have the following rule to use a universally quantified statement.

$$\frac{\Gamma, \phi[t/x] \Longrightarrow_{\mathcal{L}} \Delta}{\Gamma, \forall x:\tau.\phi \Longrightarrow_{\mathcal{L}} \Delta} \quad (\forall\text{-L})$$

where $\emptyset \triangleright_{\mathcal{L}} t : \tau$.

To use a universally quantified statement we must instantiate it with some concrete term. The rule corresponds to phrases like “*We apply Lemma xxx / the above assumption to $x = t$...*” or “*Applying Lemma yyy / the above assumption in this situation yields...*”.

One should note that these two rules preserve but do not always reflect validity, i.e., it may be that the conclusion of a rule is valid, yet the premise is not. After all, one might have chosen the wrong instantiation. Moreover, it is possible that a universal assumption must be instantiated more than once (consider e.g. an assumption asserting that some relation is transitive), so sometimes one has to keep the quantified formula for later use by prior invocation of rule CONTR.

Next, we have a rule for proving a universal statement:

$$\frac{\Gamma \Longrightarrow_{\mathcal{L}, c:\tau} \Delta, \phi[c/x]}{\Gamma \Longrightarrow_{\mathcal{L}} \Delta, \forall x:\tau.\phi} \quad (\forall\text{-R})$$

Here $c : \tau$ is a fresh constant symbol not occurring in \mathcal{L} hence in Γ, Δ, ϕ .

To prove $\forall x:\tau.\phi$ we must prove ϕ for a fixed but arbitrary $c : \tau$.

“*Fix an arbitrary $c : \tau$... this proves $\forall x:\tau.\phi$* ”

Finally, we need a rule to use an existential statement:

$$\frac{\Gamma, \phi[c/x] \Longrightarrow_{\mathcal{L}, c:\tau} \Delta}{\Gamma, \exists x:\tau.\phi \Longrightarrow_{\mathcal{L}} \Delta} \quad (\exists\text{-L})$$

To use an existential statement we introduce a fresh name for its witness. We know nothing about the witness except that it satisfies ϕ .

“*Lemma xxx provides us with a c such that $\phi[c/x]$ ”*

“*Let c be the x provided by (13) above”*

We notice that in the latter two rules no formulas with free variables arise as the bound variable is immediately replaced with a fresh constant. There are alternative presentations in which free variables are used for the “fixed but arbitrary constants” occurring in those rules. In a typed setting admitting empty domains of quantification this seems less appropriate as we then would have to annotate each sequent with the set of variables it depends on. Moreover, a variable is supposed to vary, whereas these constants are fixed.

Let’s take a look at a couple of representative examples.

$$\begin{array}{c}
\frac{}{P(c) \Longrightarrow_{\mathcal{L}, c:\tau} P(c)} \text{AXIOM} \\
\frac{P(c) \Longrightarrow_{\mathcal{L}, c:\tau} P(c)}{P(c) \Longrightarrow_{\mathcal{L}, c:\tau} \exists x:\tau.P(x)} \exists\text{-R} \\
\frac{P(c) \Longrightarrow_{\mathcal{L}, c:\tau} \exists x:\tau.P(x)}{\Longrightarrow_{\mathcal{L}, c:\tau} P(c) \Rightarrow \exists x:\tau.P(x)} \Rightarrow\text{-R} \\
\frac{\Longrightarrow_{\mathcal{L}, c:\tau} P(c) \Rightarrow \exists x:\tau.P(x)}{\Longrightarrow_{\mathcal{L}} \forall x:\tau.P(x) \Rightarrow \exists x:\tau.P(x)} \forall\text{-R} \\
\frac{}{P(c_1) \vee Q(c_2) \Longrightarrow_{\mathcal{L}, c_1:\tau_1, c_2:\tau_2} P(c_1), Q(c_2)} \text{PROP} \\
\frac{P(c_1) \vee Q(c_2) \Longrightarrow_{\mathcal{L}, c_1:\tau_1, c_2:\tau_2} P(c_1), Q(c_2)}{\forall y:\tau_2.P(c_1) \vee Q(y) \Longrightarrow_{\mathcal{L}, c_1:\tau_1, c_2:\tau_2} P(c_1), Q(c_2)} \forall\text{-L} \\
\frac{\forall y:\tau_2.P(c_1) \vee Q(y) \Longrightarrow_{\mathcal{L}, c_1:\tau_1, c_2:\tau_2} P(c_1), Q(c_2)}{\forall x:\tau_1.\forall y:\tau_2.P(x) \vee Q(y) \Longrightarrow_{\mathcal{L}, c_1:\tau_1, c_2:\tau_2} P(c_1), Q(c_2)} \forall\text{-L} \\
\frac{\forall x:\tau_1.\forall y:\tau_2.P(x) \vee Q(y) \Longrightarrow_{\mathcal{L}, c_1:\tau_1, c_2:\tau_2} P(c_1), Q(c_2)}{\forall x:\tau_1.\forall y:\tau_2.P(x) \vee Q(y) \Longrightarrow_{\mathcal{L}} \forall x:\tau_1.P(x), \forall x:\tau_2.Q(x)} \forall\text{-R} \\
\frac{\forall x:\tau_1.\forall y:\tau_2.P(x) \vee Q(y) \Longrightarrow_{\mathcal{L}} \forall x:\tau_1.P(x), \forall x:\tau_2.Q(x)}{\forall x:\tau_1.\forall y:\tau_2.P(x) \vee Q(y) \Longrightarrow_{\mathcal{L}} (\forall x:\tau_1.P(x)) \vee (\forall x:\tau_2.Q(x))} \forall\text{-R}
\end{array}$$

5.5 Soundness and completeness of first-order sequent calculus

As before we have that a sequent is valid if and only if it is derivable in the sequent calculus, i.e., if there is a proof tree whose leaves are labelled with axioms. Unlike in the propositional case, the contraction rule CONTR is not redundant corresponding to the fact that universal premises may need to be used more than once. This thwarts a naive decision procedure for validity based on constructing a generic proof tree and, indeed, as was shown by TURING validity in first order logic is undecidable. Actually, this result is not very surprising if we consider that basically all of mathematics can be formalised in first-order logic.

Theorem: A sequent is valid if and only if it is derivable in the sequent calculus.

Proof: The “if” direction of the correctness theorem (“soundness”) is proved as before by induction on derivations; we simply have to check that all the rules *preserve* validity.

For the “only if” direction (“completeness”) we construct a generic proof tree as in the propositional case by breaking down connectives and if nothing else helps instantiating quantifiers \forall -L, \exists -R. We must make sure that we keep those quantified statements around using CONTR prior to instantiating. If we arrange things in such a way that eventually a quantified formula will be instantiated with every possible term we are sure to find a proof if one exists.

If no proof exists our generic proof tree contains a leaf that is not an axiom or has an infinite path.

From the infinite path we will construct a counter interpretation by taking terms (also containing the constants newly introduced along the path) to interpret the types, function constants interpreting themselves, and interpreting predicate constants according to how atomic formulas involving them occur along the path. This will ensure that the interpretation falsifies all the sequents along the path hence the root sequent which by assumption has no proof.

Let us look at this in some more detail. Firstly, to counter the information loss in the instantiating rule \forall -L and \exists -R we replace them by the following combinations with rule CONTR:

$$\frac{\Gamma, \forall x:\tau.\phi, \phi[t/x] \Longrightarrow_{\mathcal{L}} \Delta}{\Gamma, \forall x:\tau.\phi \Longrightarrow_{\mathcal{L}} \Delta} \quad (\forall\text{-L}')$$

$$\frac{\Gamma \Longrightarrow_{\mathcal{L}} \Delta, \exists x:\tau.\phi, \phi[t/x]}{\Gamma \Longrightarrow_{\mathcal{L}} \Delta, \exists x:\tau.\phi} \quad (\exists\text{-R}')$$

It is clear that there is a proof with the primed rules if and only if there is one in the original system.

We also note that the new rules preserve and reflect validity: the conclusion of any one of the new rules is valid iff its premise is.

A *generic proof tree* is a possibly infinite tree labelled with sequents which has the following properties:

1. each internal node is the conclusion of its immediate ancestors by some proof rule.
2. rule WEAK is not used,

3. rules \forall -L' and \exists -R' are used only with conclusion $\Gamma \Longrightarrow_{\mathcal{L}} \Delta$ where Γ contains atoms and universally quantified formulas only and Δ contains atoms and existentially quantified formulas only. Otherwise we could use one of the validity-reflecting rules.
4. no internal node is labelled with an axiom, i.e., we stop once we have found an axiom
5. on every infinite path starting from $\Gamma \Longrightarrow_{\mathcal{L}} \Delta, \exists x:\phi$ the formula $\exists x:\phi$ is instantiated with all (closed) terms of type τ in \mathcal{L}
6. Ditto for infinite paths starting from $\Gamma, \forall x:\tau \Longrightarrow_{\mathcal{L}} \Delta$

These properties basically dictate a strategy for obtaining a generic proof tree starting from any sequent. Simply apply the rules backwards with the mentioned restriction on the rules that instantiate quantifiers. When selecting instantiations make sure that every possible instantiation will be eventually chosen unless of course a path ends with an axiom leaf. Please note, that as soon as we make a language extension we must instantiate our quantified formulas with all the terms in the new language as well.

Now suppose that a sequent S has no proof. The generic proof tree constructed from S might have a finite path ending in a non axiom consisting of atomic formulas only. In this case, we can argue as in the case of propositional logic that the root sequent is not valid.

Alternatively, and this is the interesting case, the generic proof tree will contain an infinite path π (starting from the root). This is ‘‘König’s Lemma’’: a finitely branching tree with infinitely many nodes has an infinite path. Along this infinite path π we encounter an increasing (by constants) sequence of languages $\mathcal{L}_1 \subseteq \mathcal{L}_2 \subseteq \dots$ whose union we call \mathcal{L}_∞ . So, a term in \mathcal{L}_∞ will be a term of one of the \mathcal{L}_i .

To construct our desired counterinterpretation \mathcal{I} we interpret types by

$$\llbracket \tau \rrbracket_{\mathcal{I}} = \{t \mid \emptyset \triangleright_{\mathcal{L}_\infty} t : \tau\}$$

We interpret function constants by

$$\llbracket f \rrbracket_{\mathcal{I}}(t_1, \dots, t_n) = f(t_1, \dots, t_n)$$

We interpret predicate constants by

$$\llbracket P \rrbracket_{\mathcal{I}}(t_1, \dots, t_n) = \begin{cases} \text{tt}, & \text{if } P(t_1, \dots, t_n) \text{ is among the antecedents} \\ & \text{(left of } \Longrightarrow) \text{ of a sequent in } \pi, \\ \text{ff}, & \text{in all other cases.} \end{cases}$$

Now we show by induction on the size of formulas that whenever a formula ϕ appears as an antecedent of a sequent in π then $\llbracket \phi \rrbracket_{\mathcal{I}} = \text{tt}$ and whenever a formula ψ appears as a succedent (to the right of the \implies) of a sequent in π then $\llbracket \psi \rrbracket_{\mathcal{I}} = \text{ff}$, so that in particular all the sequents along π including the root will be falsified by \mathcal{I} . So, \mathcal{I} shows that the root is not a valid sequent.

Atomic formulas are true under \mathcal{I} precisely if they appear as an antecedent. If an atomic formula appears as a succedent then—since atomic formulas never disappear along the path—it cannot also appear as an antecedent for otherwise we would have an axiom sequent on π contrary to the construction of the generic proof tree. Thus, atomic formulas appearing as succedents are falsified by \mathcal{I} . A formula which is not an existentially quantified succedent or a universally quantified antecedent will eventually be broken down by a validity reflecting rule into its subformulas to which the induction hypothesis applies. Consider for example a succedent of the form $\forall x:\tau.\phi$. At some point rule \forall -L will be applied, so $\phi[c/x]$ also occurs as a succedent on π . By the induction hypothesis $\llbracket \phi[c/x] \rrbracket_{\mathcal{I}} = \text{ff}$, but then $\llbracket \forall x:\tau.\phi \rrbracket_{\mathcal{I}} = \text{ff}$, too.

If, finally, we have an existentially quantified succedent, e.g., $\exists x:\phi$ then by the “round robin” policy used for instantiating all formulas, all the formulas $\phi[t/x]$ with $\emptyset \triangleright_{\mathcal{L}_\infty} t : \tau$ will occur as succedents along π hence are falsified by \mathcal{I} . Since $\llbracket \tau \rrbracket_{\mathcal{I}}$ comprises precisely all those terms we conclude that $\llbracket \exists x:\tau.\phi \rrbracket_{\mathcal{I}} = \text{ff}$, as well. The case of a universally quantified antecedent is analogous. This completes the proof.

One should not underestimate the power of first-order logic. Even without function constants the counter interpretation may be infinite due to infinitely many newly introduced constants. Consider for example the formula

$$(\forall x, y, z:\tau.R(x, y) \wedge R(y, z) \implies R(x, z)) \wedge (\forall x:\tau.\exists y:\tau.R(x, y)) \implies \exists x:\tau.R(x, x)$$

It is not valid but holds in all finite interpretations. You may find it instructive to form the generic proof tree for this formula.

5.6 First-order logic in PVS

Language concepts are declared anywhere in a theory, but before being used

```
D, T1, T2 : TYPE+
c : T1
P : [D -> boolean]
```

```
Q : boolean
f : [T1->T2]
```

Here TYPE+ stands for *nonempty* type. There is also the declaration `T : TYPE` which stands for a possibly empty type. In this case it would not be allowed to declare a constant of type T.

This design decision of PVS is open to debate. By declaring a constant of a type we explicitly state that it is nonempty so why say it twice?

Quantifiers are written `FORALL (x:T) :` and `EXISTS (x:T) :` Do not forget the colon after the parenthesis.

```
exI : THEOREM
      FORALL (x:D) : P (x) IMPLIES EXISTS (x:D) : P (x)
orex : THEOREM
      (Q OR EXISTS (x:D) : P (x)) IMPLIES EXISTS (x:D) : Q OR P (x)
depp : THEOREM
      EXISTS (x:D) : FORALL (y:D) : P (x) IMPLIES P (y)
gen : THEOREM
      (EXISTS (x:D) : P (x)) IMPLIES FORALL (x:D) : P (x)
```

The rules \forall -L and \exists -R are invoked with the command `inst` (instantiation). The rules \forall -L and \exists -R are invoked with the command `skolem` (after TH. SKOLEM).

The `inst` command takes as argument a formula number (the formula to be instantiated) and a term to instantiate with. For example, in the situation

```
{-1} P (d)
    |-----
{1}  EXISTS (x:T1) : P (x)
```

the command

```
(inst 1 "d")
```

leads to

```
{-1} P (d)
    |-----
{1}  P (d)
```

which is an axiom.

The `skolem` command takes as argument a formula number and the name of a new constant. If it isn't fresh then PVS complains. For example, in the situation

```

|-----
{1}   FORALL (x:D) : P (x)  IMPLIES  EXISTS (x:D) : P (x)

```

The command `(skolem 1 "c")` is no good because we have already used `c` for a constant above. However, `(skolem 1 "d")` succeeds and gives

```

|-----
{1}   P (d)  IMPLIES  EXISTS (x:D) : P (x)

```

The command `M-x show-skolem-constants` displays all the constants introduced in the course of the proof.

The commands `inst` and `skolem` allow the treatment of several variables at once. There are also the derived form `inst?` which guesses an appropriate instantiation heuristically (alas often quite badly) and `skolem!` which automatically introduces as many constants as possible (making up fresh names for them). Furthermore, `skosimp` is a combination of `skolem!` and simplification. See the PVS prover guide for details.

As an exercise try to prove all of the “theorems” below.

```

fol : THEORY
  BEGIN
    D, T1, T2 : TYPE+
    c : T1
    d : D
    P : [D -> boolean]
    Q : boolean

    alle : THEOREM
      FORALL (x:D) : (FORALL (y:D) : P (y)) IMPLIES P (x)

    andall : THEOREM
      (Q AND (FORALL (x:D) : P (x))) IMPLIES FORALL (x:D) : Q AND P (x)

    exI : THEOREM
      FORALL (x:D) : P (x) IMPLIES EXISTS (x:D) : P (x)

    andex : THEOREM

```

```

(Q AND EXISTS(x:D):P(x)) IMPLIES EXISTS(x:D):Q AND P(x)

orex : THEOREM
      (Q OR EXISTS(x:D):P(x)) IMPLIES EXISTS(x:D):Q OR P(x)

depp : THEOREM
      EXISTS(x:D): FORALL(y:D): P(x) IMPLIES P(y)

doub : THEOREM
      FORALL(x,y:D) : EXISTS (z:D) : P(z) IMPLIES P(x) & P(y)

P1 : [T1->boolean]
P2 : [T2 -> boolean]

por : THEOREM
      (FORALL(x:T1,y:T2):P1(x) OR P2(y)) IMPLIES
      (FORALL(x:T1):P1(x)) OR (FORALL(x:T2):P2(x))

R : [D,D->boolean]

per: THEOREM
      (FORALL (x,y:D):R(x, y) IMPLIES R(y, x)) AND
      (FORALL (x,y,z:D): R(x, y) AND R(y, z) IMPLIES R(x, z)) IMPLIES
      (FORALL (x:D): (EXISTS (y:D): R(x,y)) IMPLIES R(x,x))
END fol

```

6 First-order resolution

As in the propositional case the method of resolution provides a generally more efficient way to decide validity of formulas than proof search in Gentzen's sequent calculus. In the first-order case we may instantiate universally quantified variables prior to resolving so as to achieve agreement of literals. For example, we may resolve the clauses $\{P(f(x, g(y))), Q(x)\}$ which denotes $\forall x, y. P(f(x, g(y))) \vee Q(x)$ (types omitted) and $\{\neg P(f(g(z), w)), R(w)\}$ which denotes $\forall w, z. \neg P(f(g(z), w)) \vee R(w)$ to form $\{R(g(y)), Q(g(z))\}$ which denotes $\forall y, z. R(g(y)) \vee Q(g(z))$.

Notice that again a satisfying interpretation for the former two clauses will also satisfy the latter. In this example, we could also have instantiated x with

$$\begin{aligned}
C_1 &= \{P(f(x, g(y))), Q(x)\} && , \text{ i.e., } \forall x, y. P(f(x, g(y))) \vee Q(x) \\
C_2 &= \{\neg P(f(g(z), w)) \vee R(w)\} && , \text{ i.e., } \forall w, z. \neg P(f(g(z), w)) \vee R(w) \\
&&& \text{resolve to} \\
C_3 &= \{R(g(y)) \vee Q(g(z))\} && , \text{ i.e., } \forall y, z. R(g(y)) \vee Q(g(z))
\end{aligned}$$

Figure 21: Example of resolution

something like $g(f(h(c())))$ and accordingly y with $f(h(c()))$. However, in order to maximise future success it is advisable to choose the instantiation which makes the least possible commitment or, in formal terms, the *most general unifier*. While this is in practice always done, it is, for the purpose of establishing completeness, easier to allow arbitrary instantiations.

Let us look at the details. First-order resolution operates on clauses which are sets of first-order literals, i.e., negated or non-negated atomic formulas, which are understood as being universally quantified over the variables they contain. In order to avoid problems with empty types we assume that our language is such that every type contains at least one closed term.

- A first-order literal is a negated or non-negated atomic formula.
- A first-order clause is a set of first-order literals
- It denotes the disjunction of the literals universally quantified over the variables

Given a set of first-order clauses \mathcal{C} we can use first-order resolution to decide whether it is satisfiable, i.e., whether there exists an interpretation which makes it true. If we can derive the empty clause from \mathcal{C} by successive application of rules INST and RES then surely \mathcal{C} is unsatisfiable. Conversely, if \mathcal{C} is unsatisfiable then it is possible to derive the empty clause. The proof of this result is based on correctness of propositional resolution and Herbrand's theorem which asserts that a set of formulas of the form $\forall \vec{x}:\vec{\tau}.\phi$ with ϕ quantifier-free is satisfiable if and only if the set of its closed instantiations is propositionally satisfiable:

Theorem ("Herbrand's theorem"): Let \mathcal{S} be a set of formulas of the form $\forall \vec{x}:\vec{\tau}.\phi$ with ϕ quantifier-free.

Define

$$\bar{\mathcal{S}} := \{\phi[t_1/x_1, \dots, t_n/x_n] \mid \forall x_1:\tau_1, \dots, \forall x_n:\tau_n.\phi \in \mathcal{S} \text{ and } \emptyset \triangleright_{\mathcal{L}} t_i:\tau_i\}$$

- Instantiation rule:

$$\frac{C}{C[t_1/x_1, \dots, t_n/x_n]} \quad (\text{INST})$$

where x_i are the variables mentioned in C and the t_i are *possibly open* terms (of the right type!).

- Resolution rule

$$\frac{C_1 \cup \{A\} \quad C_2 \cup \{\neg A\}}{C_1 \cup C_2} \quad (\text{RES})$$

- Side condition: there is a closed term of each type.
- Aim: try to derive empty clause from initial set so as to show unsatisfiability.

Figure 22: First-order resolution

as the set of closed instantiations of formulas in \mathcal{S} .

There exists an interpretation \mathcal{I} validating all formulas in \mathcal{S} if and only if there exists a propositional valuation η of the atomic formulas (viewed as propositional atoms) validating all formulas in $\bar{\mathcal{S}}$.

Proof: Given \mathcal{I} define η by $\eta(P(t_1, \dots, t_n)) = \llbracket P(t_1, \dots, t_n) \rrbracket_{\mathcal{I}}$. Given η interpret types as sets of (closed) terms, function symbols by themselves, and predicates as given by η . \square

Thus to establish unsatisfiability of a set of first-order clauses it is enough to establish propositional unsatisfiability of their closed instantiations, but that's precisely what rule RES can do as shown in Section 4.

Rule INST on the other hand, allows us to generate the set of closed instantiations. Performing resolution on clauses containing (universally quantified) variables certainly does no harm, but may of course speed up success.

6.1 Most general unifiers

As already mentioned, in practice one resolves clauses by instantiating with the *most general unifier*. The most general unifier of two open terms $u(x_1, \dots, x_m)$ and $v(y_1, \dots, y_n)$ consists of two sequences of open terms t_1, \dots, t_m and s_1, \dots, s_n involving variables z_1, \dots, z_k , such that $u(t_1, \dots, t_m) = v(s_1, \dots, s_n)$ and, moreover, any instantiation making u equal to v arises from this one by instantiation,

- $u(x_1, \dots, x_m), v(y_1, \dots, y_n)$ two term with free variables \vec{x} and \vec{y} .
- Most general unifier consists of $\vec{s}(\vec{z})$ and $\vec{t}(\vec{z})$ such that $u(\vec{s}(\vec{z})) = v(\vec{t}(\vec{z}))$ and
- whenever $u(\vec{s}') = v(\vec{t}')$ then $s' = \vec{s}(\vec{a}), t' = \vec{t}(\vec{b})$.
- **Example:** $u = f(x, g(y)), v = f(h(y), x)$:
 $\vec{s} = [h(z)/x, x/y], \vec{t} = [z/y, g(x)/x]$
- **Note:** the most general unifier might not exist, e.g., $s = f(x), t = g(y)$.

Figure 23: Most general unifier

$$\begin{array}{ll}
\{ \neg S(r), \neg A(r, x, y), A(r, y, x) \}, & \text{(C1)} \\
\{ A(r, f(r), g(r)), S(r) \}, & \text{(C2)} \\
\{ \neg A(r, g(r), f(r)), S(r) \}, & \text{(C3)} \\
\{ S(s()) \}, & \text{(C4)} \\
\{ A(s(), a(), b()) \}, & \text{(C5)} \\
\{ \neg A(s(), b(), a()) \} & \text{(C6)}
\end{array}$$

Figure 24: Example of first-order resolution

i.e., whenever $u(\vec{t}') = v(\vec{s}')$ then $\vec{t}' = \vec{t}[\vec{a}/\vec{z}]$ and $\vec{s}' = \vec{s}[\vec{b}/\vec{z}]$ for some \vec{a}, \vec{b} .

For example, the most general unifier of $f(x, g(y))$ and $f(h(y), x)$ is $\vec{s} = [h(z)/x, x/y], \vec{t} = [z/y, g(x)/x]$ because we have $f(x, g(y))[h(z)/x, x/y] = f(h(z), g(x)) = f(h(y), x)[z/y, g(x)/x]$. Notice that the variable names in u, v as well as the common ones in the s, t are rather arbitrary. In particular, if the same variable happens to occur in both u and v , we can instantiate it differently in both.

The most general unifier is effectively found by comparing the terms in question in a top down fashion starting from the outermost function constant. We omit the details of the unification algorithm and also a formal proof that resolution with most general unifiers is complete. The idea is to map any proof using RES and INST to a proof using only the following combined rule.

$$\frac{C_1 \cup \{A_1\} \quad C_2 \cup \{\neg A_2\} \quad \vec{s}, \vec{t} \text{ m.g.u. of } A_1, A_2}{C_1[\vec{s}] \cup C_2[\vec{t}]} \text{ (RES-UNIF)}$$

Of course it goes without saying that all the instantiations made in the course

of resolution must be type correct, i.e., the resulting terms and atomic formulas must be well-formed.

At this point it is worth reiterating the point made earlier in Section 5 about types separating interesting and potentially difficult facts from uninteresting obvious facts. The number of clauses hence the search space for resolution becomes smaller the more we make use of types.

6.2 Skolemisation

We now discuss how to translate arbitrary first-order formulas into first-order clauses; somewhat surprisingly, any first-order formula is equivalent to set of first-order clauses albeit in a richer language.

In order to do that we use the following fact known as *skolemisation*, again after TH. SKOLEM.

Fact: Let $\psi := \forall x_1:\tau_1 \dots \forall x_n:\tau_n. \exists y:\tau_{n+1}. \phi(\vec{x}, y)$ in some language \mathcal{L} and let \mathcal{L}' be the language \mathcal{L} extended with a new function constant $f : [\tau_1, \dots, \tau_n \rightarrow \tau_{n+1}]$. The formula ψ is satisfiable if and only if the formula $\forall x_1:\tau_1 \dots \forall x_n:\tau_n. \phi(\vec{x}, f(\vec{x}))$ is satisfiable.

The function constant f is called a *Skolem function*.

Now consider an arbitrary first-order formula ϕ . Using the following tautologies

$$\begin{aligned} \alpha \star (Qx:\tau. \beta(x)) &\Leftrightarrow Qx:\tau. \alpha \vee \beta(x) & \star \in \{\vee, \wedge, \Rightarrow\}, Q \in \{\forall, \exists\} \\ (Qx:\tau. \alpha(x)) \Rightarrow \beta &\Leftrightarrow \bar{Q}x:\tau. \alpha(x) \Rightarrow \beta \\ (\neg Qx:\tau. \alpha(x)) &\Leftrightarrow \bar{Q}x:\tau. \neg \alpha(x) \end{aligned}$$

where $\bar{\exists} = \forall, \bar{\forall} = \exists$ we can bring each formula (up to equivalence) into the form

$$Q_1 x_1:\tau_1. Q_2 x_2:\tau_2. \dots. Q_n x_n:\tau_n. \phi_0$$

with $Q_1, \dots, Q_n \in \{\forall, \exists\}$ and ϕ_0 quantifier-free. The latter formula is by definition in *prenex form*.

Thereafter, using the “fact” we can successively replace existential quantifiers with new function constants that take all the previous universally quantified variables as arguments so as to obtain a universally quantified boolean combination of atomic formulas which in turn is equivalent to a set of first-order clauses. Summing up, we have the following result.

Theorem: For every first-order formula ϕ one can effectively find a set of first-order clauses \mathcal{C} such that ϕ is satisfiable if and only if \mathcal{C} is.

Proof: Bring ϕ into prenex form (move all quantifiers to the front exchanging \forall and \exists when moving out of a negative position). Replace existential quantifiers by *Skolem functions* using language extension by function constants. Bring the resulting universally quantified formula into clausal form as in propositional case. \square

In my view, the reason why resolution is superior to proof search in sequent calculus is that the choice of instantiations is made after looking at two clauses (and performing unification) which, when informally translated back to sequent calculus, means that the form of the side formulas Γ, Δ in a sequent, say, $\Gamma \Longrightarrow_{\mathcal{L}} \Delta, \exists x:\tau.\phi$ helps in finding an appropriate instantiation for x . I am not aware of a precisation of this argument in the form of a unification-based strategy for finding instantiations in sequent calculus proof search. In this context one should note that the flattening of nested quantifications using Skolemisation is crucial for the success of unification.

6.3 Some puzzles

Here are some small examples that should be brought into clausal form and proved by hand, using PVS, or automatically using SPASS.

The mislabelled boxes (from <http://www.cs.miami.edu/~tptp/>): There are three boxes a, b, and c on a table. Each box contains apples or bananas or oranges. No two boxes contain the same thing. Each box has a label that says it contains apples or says it contains bananas or says it contains oranges. No box contains what it says on its label. The label on box a says "apples". The label on box b says "oranges". The label on box c says "bananas". You pick up box b and it contains apples. What do the other two boxes contain?

Barber's problem (from <http://www.cs.miami.edu/~tptp/>): There is a barbers' club that obeys the following three conditions:

1. If any member has shaved any other member – whether himself or another – then all members have shaved him, though not necessarily at the same time.
2. Four of the members are named Guido, Lorenzo, Petrucio, and Cesare.
3. Guido has shaved Cesare. Prove Petrucio has shaved Lorenzo

Continuity of composition

$$\begin{aligned}\mathcal{T} &= \{\rho, \iota\} \\ \mathcal{F} &= \{f : [\rho \rightarrow \rho]\} \\ \mathcal{P} &= \{\in : [\rho, \iota \rightarrow \text{boolean}]\}\end{aligned}$$

$$\begin{aligned}\forall x:\rho \forall U:\iota. \in (f(x), U) &\Rightarrow \exists V:\iota. \in(x, V) \wedge \forall y:\rho. \in(y, V) \Rightarrow \in(f(y), U) \\ &\Rightarrow \\ \forall x:\rho \forall U:\iota. \in(f(f(x)), U) &\Rightarrow \exists V:\iota. \in(x, V) \wedge \forall y:\rho. \in(y, V) \Rightarrow \in(f(f(y)), U)\end{aligned}$$

6.4 Compactness of first-order logic

Theorem: Let Φ be a set of first-order formulas over some signature. If every finite subset of Φ has a model then Φ itself has a model, too.

Proof: Using skolemisation we may assume without loss of generality that Φ consists of formulas of the form $\forall \vec{x}:\vec{\tau}. \phi$ with ϕ quantifier-free.

Let us form the propositional theory Π consisting of closed-instantiations of the formulas ϕ as in Herbrand's theorem. If every finite subset of Φ has a model then every finite subset of Π will be satisfiable, since a finite subset of Π can only involve a finite subset of Φ . By compactness of propositional logic therefore the whole of Π is satisfiable and by Herbrand's theorem Φ has a model.

The compactness theorem has a number of perhaps surprising consequences. Consider, for example, the set $\Theta_{\mathbb{N}}$ of closed formulas (over the signature $(\text{nat}, +, \times, 0, 1, \geq)$) that are true in the standard interpretation that interprets nat as the natural numbers etc. This set of formulas, the *first-order theory* of the natural numbers contains in particular all the instances of the Peano axioms but much more, e.g., those formulas that are true but not provable from the Peano axioms.

Now let us extend the signature by a special constant $c : [\rightarrow \text{nat}]$ and the formulas $\phi_n := \exists y:\text{nat}. c \geq 1 + \dots + 1$ (n summands).

Every finite subset of this extended set has a model, namely the natural numbers with c interpreted as a large enough number. By compactness therefore the whole set has a model which is a structure validating the same first-order formulas as the natural numbers themselves, yet contains an infinitely large number—the interpretation of c .

Let us explore the structure of such a *non-standard model* of arithmetic. As we have seen, it contains a number c that is greater than any standard number. Since every number (including c has a successor there are more infinite numbers $c + 1, c + 2, c + 3$, etc. Since every non-zero number has a predecessor (that is

a valid first-order sentence!) there must also be $c - 1, c - 2, c - 3$, etc. So the numbers around c form a structure isomorphic to \mathbb{Z} . Since every number can be doubled and halved (in the floor-sense) there must be another such \mathbb{Z} block above the one surrounding c and one below. In between any two distinct \mathbb{Z} -blocks there must be another one, etc. So any countable non-standard model has an order type isomorphic to $\mathbb{N} + \mathbb{Z} \cdot \mathbb{Q}$.

In a similar way, we can use compactness to show consistency (with respect to first-order logic!) of infinitesimal numbers. Add to the theory of the real numbers the infinitely many axioms $0 < c < 2^{-n}$. Every finite subset is consistent so the whole set is and it thus has a model in which there is a constant c that is arbitrarily close to zero. Of course, then, say, $\pi + c$ is arbitrarily close to π , etc. We can then define a derivative as something like $f(x + c)/c$. Notice that if we prove some first-order statement from this extended theory then, again by compactness, only finitely many of the assumptions $0 < c < 2^{-n}$ will have been used, so in this case, an ordinary c will do.

7 Equality

Most mathematical statements of interest involve equality. It is in principle possible to treat equality just as a predicate constant and to assume axioms stating that equality is an equivalence relation compatible with all function and predicate constants. For practical purposes it is, however, more convenient to introduce equality as a special primitive concept.

So, in first order logic with equality atomic formulas can be formed in two ways:

- $P(t_1, \dots, t_n)$ where t_1, \dots, t_n are terms of types τ_1, \dots, τ_n and $P : [\tau_1, \dots, \tau_n \rightarrow \text{boolean}]$. That's as before.
- $t_1 = t_2$ where t_1, t_2 are terms of some common type τ .

Such an equality formula $t_1 = t_2$ is true if and only if under the interpretation at hand the two terms t_1, t_2 have equal meaning.

The sequent calculus can be extended so as to cope with equality by adding the following rules:

$$\frac{\emptyset \triangleright_{\mathcal{L}} t : \tau \text{ for some } \tau}{\Gamma \Longrightarrow_{\mathcal{L}} \Delta, t = t} \quad (\text{REFL})$$

$$\begin{array}{c}
\frac{\Gamma, t_1 = t_2, \phi[t_2/x] \Longrightarrow_{\mathcal{L}} \Delta}{\Gamma, t_1 = t_2, \Longrightarrow_{\mathcal{L}} \Delta, \neg\phi[t_2/x]} \neg\text{-R} \\
\frac{\Gamma, t_1 = t_2, \Longrightarrow_{\mathcal{L}} \Delta, \neg\phi[t_2/x]}{\Gamma, t_1 = t_2, \Longrightarrow_{\mathcal{L}} \Delta, \neg\phi[t_1/x]} \text{SUBST-R} \\
\frac{\Gamma, t_1 = t_2, \Longrightarrow_{\mathcal{L}} \Delta, \neg\phi[t_1/x]}{\Gamma, t_1 = t_2, \neg\neg\phi[t_1/x] \Longrightarrow_{\mathcal{L}} \Delta} \neg\text{-L} \\
\frac{\phi[t_1/x] \Longrightarrow_{\mathcal{L}} \neg\neg\phi[t_1/x] \quad \Gamma, t_1 = t_2, \neg\neg\phi[t_1/x] \Longrightarrow_{\mathcal{L}} \Delta}{\Gamma, t_1 = t_2, \phi[t_1/x] \Longrightarrow_{\mathcal{L}} \Delta} \text{CUT} \\
\\
\frac{}{t_1 = t_2 \Longrightarrow_{\mathcal{L}} t_2 = t_2} \text{REFL} \\
\frac{t_1 = t_2 \Longrightarrow_{\mathcal{L}} t_2 = t_2}{t_1 = t_2 \Longrightarrow_{\mathcal{L}} t_2 = t_1} \text{SUBST-R} \\
\text{CUTting with the conclusion gives rules SUBST-L-RL, SUBST-R-RL.} \\
\frac{}{t_2 = t_3, t_1 = t_2 \Longrightarrow_{\mathcal{L}} t_2 = t_3} \text{AXIOM} \\
\frac{t_2 = t_3, t_1 = t_2 \Longrightarrow_{\mathcal{L}} t_2 = t_3}{t_2 = t_3, t_1 = t_2 \Longrightarrow_{\mathcal{L}} t_1 = t_3} \text{SUBST-R}
\end{array}$$

Figure 25: Example derivations

$$\frac{\Gamma, t_1 = t_2 \Longrightarrow_{\mathcal{L}} \Delta, \phi[t_2/x]}{\Gamma, t_1 = t_2 \Longrightarrow_{\mathcal{L}} \Delta, \phi[t_1/x]} \quad (\text{SUBST-R})$$

Rule REFL says that $t = t$ is vacuously true; if this is among our conclusions then we're done.

Rule SUBST-R says that if we have an equality $t_1 = t_2$ among our assumptions and we need to prove a formula $\phi[t_1/x]$ which contains t_1 as a subexpression then we can replace t_1 by t_2 and therefore prove $\phi[t_2/x]$ instead.

We immediately have the following derived rules:

$$\frac{\Gamma, t_1 = t_2, \phi[t_2/x] \Longrightarrow_{\mathcal{L}} \Delta}{\Gamma, t_1 = t_2, \phi[t_1/x] \Longrightarrow_{\mathcal{L}} \Delta} \quad (\text{SUBST-L})$$

$$\frac{\Gamma, t_1 = t_2 \Longrightarrow_{\mathcal{L}} \Delta, \phi[t_1/x]}{\Gamma, t_1 = t_2 \Longrightarrow_{\mathcal{L}} \Delta, \phi[t_2/x]} \quad (\text{SUBST-R-RL})$$

$$\frac{\Gamma, t_1 = t_2, \phi[t_1/x] \Longrightarrow_{\mathcal{L}} \Delta}{\Gamma, t_1 = t_2, \phi[t_2/x] \Longrightarrow_{\mathcal{L}} \Delta} \quad (\text{SUBST-L-RL})$$

Figure 7 contains a derivation of rule SUBST-L.

It is possible to show that sequent calculus with the equality rules derives all valid formulas; to that end one considers the quotient of the term model by the congruence relation generated by the equations appearing as antecedents of sequents on the infinite path in the generic proof tree.

Similarly, one can extend resolution with rules that allow one to replace equals with equals within clauses.

7.1 Equality in PVS

The rule REFL is treated like an axiom: as soon as PVS encounters an instance of reflexivity the corresponding subgoal is discarded (“This completes the proof of ...”) or, if it was the last open branch of the proof, the proof is completed (“Q.E.D.”).

To invoke either SUBST-R or SUBST-L we use the command

$$(\text{replace } \langle \textit{what_with} \rangle \langle \textit{where} \rangle)$$

Here $\langle \textit{what_with} \rangle$ must be the (negative) number of an equation among the antecedents; $\langle \textit{where} \rangle$ must be the number of any formula either in the antecedents (that’s SUBST-L) or in the succedents (that’s SUBST-R) which contains the left-hand-side of $\langle \textit{what_with} \rangle$.

To invoke either rule SUBST-L-RL or SUBST-R-RL we use the command

$$(\text{replace } \langle \textit{what_with} \rangle \langle \textit{where} \rangle : \text{dir RL})$$

7.2 Extended example: monoids

We assume a nonempty set M with an associative operation $*$ in infix notation.

```
M : TYPE+
* : [M, M->M]
assoc : AXIOM
  FORALL (x, y, z:M) : (x*y)*z = x*(y*z)
```

Of course, all this must be placed in a .pvs file and within something like
monoids : THEORY BEGIN...END.

Generalised associativity We want to prove an extended law of associativity:

```
assoc4 : THEOREM
  FORALL (x, y, z, w:M) : ((x*y)*z)*w = x*(y*(z*w))
```


We will first do it the basic way and then using some more advanced commands.

After starting the prover we introduce constants for the universally quantified variables with the command `(skolem!)`:

```
|-----
{1} ((x!1 * y!1) * z!1) * w!1 = x!1 * (y!1 * (z!1 * w!1))
```

where `x!1`, `y!1`, `z!1`, `w!1` are fresh constants of type `M`.

Recall, that in this case `(skolem!)` is equivalent to `(skolem 1 "x!1" "y!1" "z!1" "w!1")` which in turn is equivalent to `(skolem 1 "x!1")` followed by `(skolem 1 "y!1")` followed by `(skolem 1 "z!1")` followed by `(skolem 1 "w!1")`.

We first want to rewrite the subterm `((x!1 * y!1) * z!1)` using `assoc`. To that end we add `assoc` to our antecedents with `(lemma "assoc")`.

```
{-1} FORALL (x, y, z: M): (x * y) * z = x * (y * z)
|-----
[1] ((x!1 * y!1) * z!1) * w!1 = x!1 * (y!1 * (z!1 * w!1))
```

Normally, the `lemma` command applies to something already proved, a “lemma”. In that case it corresponds to the `CUT` rule. With axioms it’s a bit different. We can think of them as being implicitly added to the antecedents. In that case the `lemma` command simply highlights them. It is also possible to extend the sequent calculus by real axioms.

At any rate, we will need our axiom more than once, so we start by copying it corresponding to `CONTR: copy -1`.

```
{-1} FORALL (x, y, z: M): (x * y) * z = x * (y * z)
[-2] FORALL (x, y, z: M): (x * y) * z = x * (y * z)
|-----
[1] ((x!1 * y!1) * z!1) * w!1 = x!1 * (y!1 * (z!1 * w!1))
```

The `-1` formula must now be instantiated with the `inst` command: `(inst -1 "x!1" "y!1" "z!1")`.

```
{-1} (x!1 * y!1) * z!1 = x!1 * (y!1 * z!1)
[-2] FORALL (x, y, z: M): (x * y) * z = x * (y * z)
|-----
[1] ((x!1 * y!1) * z!1) * w!1 = x!1 * (y!1 * (z!1 * w!1))
```

Now we can use an equality rule: (replace -1 1):

```
[-1] (x!1 * y!1) * z!1 = x!1 * (y!1 * z!1)
[-2] FORALL (x, y, z: M): (x * y) * z = x * (y * z)
    |-----
{1} (x!1 * (y!1 * z!1)) * w!1 = x!1 * (y!1 * (z!1 * w!1))
```

The parentheses around the just replaced subterm are not displayed which is irritating. Next, we must apply associativity to the whole left hand side, the middle term being this time not just a constant, but $y!1 * z!1$. This time we use a slightly more powerful command: `inst-cp` which works like `inst` but copies the formula to be instantiated beforehand. So,

```
(inst-cp -2 "x!1" "y!1*z!1" "w!1")
```

brings us to

```
-1] (x!1 * y!1) * z!1 = x!1 * (y!1 * z!1)
[-2] FORALL (x, y, z: M): (x * y) * z = x * (y * z)
[-3] (x!1 * (y!1 * z!1)) * w!1 = x!1 * (y!1 * z!1 * w!1)
    |-----
[1] x!1 * (y!1 * z!1) * w!1 = x!1 * (y!1 * (z!1 * w!1))
```

Now (replace -3 1) gives

```
[-1] (x!1 * y!1) * z!1 = x!1 * (y!1 * z!1)
[-2] FORALL (x, y, z: M): (x * y) * z = x * (y * z)
[-3] (x!1 * (y!1 * z!1)) * w!1 = x!1 * (y!1 * z!1 * w!1)
    |-----
{1} x!1 * ((y!1 * z!1) * w!1) = x!1 * (y!1 * (z!1 * w!1))
```

where again, I've inserted some parens. This is getting close; instantiating the remaining copy of associativity with

```
(inst -2 "y!1" "z!1" "w!1")
```

followed by (replace -2 1) completes the proof.

High-level proof Now let's do the same proof again with more powerful commands: After `(skolem!)` we get as before

```
|-----  
{1} ((x!1 * y!1) * z!1) * w!1 = x!1 * (y!1 * (z!1 * w!1))
```

Now rather than bringing in `assoc`, instantiating, and then rewriting our goal with it, we can use the command `rewrite-lemma` (p. 65 of `prover-guide.ps`) which does these two steps in one go:

```
(rewrite-lemma "assoc" ("x" "x!1" "y" "y!1" "z" "z!1"))
```

results in

Rewriting using `assoc` where

x gets `x!1`,

y gets `y!1`,

z gets `z!1`,

this simplifies to:

```
assoc4 :
```

```
|-----  
{1} x!1 * (y!1 * z!1) * w!1 = x!1 * (y!1 * (z!1 * w!1))
```

The command `rewrite-lemma` takes as second argument a substitution which is an even length list providing the required values for all the bound variables in the lemma. After performing this instantiation it must become an equation with which rewriting then takes place. Admittedly, this syntax is somewhat inconsistent with the syntax of the `(inst)` command.

Now, we want to perform the same procedure again, but with a different substitution:

```
(rewrite-lemma "assoc" ("x" "x!1" "y" "y!1 * z!1" "z" "w!1"))
```

Fortunately, we don't need to type in again from scratch: the keystroke `M-p`, that is the `[Alt]` key and the `[P]` key together, brings up the last command entered. We only need to edit the substitution. Further `M-p` bring up even earlier commands. If we've gone too far, we can use `M-n` to go back again. There's another way to ease typing: If we start to type a command like so

```
(rewri
```

and then type `M-s` it will be completed to the last command typed with the same beginning, i.e., in our case the last `rewrite-lemma` command which again can then be edited.

However we enter the command, it brings us to

```
|-----
{1}  x!1 * (y!1 * z!1 * w!1) = x!1 * (y!1 * (z!1 * w!1))
```

at which point

```
(rewrite-lemma "assoc" ("x" "y!1" "y" "z!1" "z" "w!1"))
```

completes the job.

The `rewrite-lemma` command can also be given the `:dir RL` optional argument, so we could have worked on the right hand side instead like so:

```
|-----
{1}  ((x!1 * y!1) * z!1) * w!1 = x!1 * (y!1 * (z!1 * w!1))
```

```
Rule? (rewrite-lemma "assoc" ("x" "y!1" "y" "z!1" "z" "w!1") :dir RL)
```

```
|-----
{1}  ((x!1 * y!1) * z!1) * w!1 = x!1 * ((y!1 * z!1) * w!1)
```

Even quicker proofs: Filling in the instantiations is tedious and can partly be automated. That's what the command `rewrite` (p.64 of `prover-guide.ps`) does for us. Unfortunately, not always successfully, which is why it's good to know the more basic commands. In the example at hand it works, however, and we can do the entire proof by issuing the following four commands:

```
(skolem!)
(rewrite "assoc")
(rewrite "assoc")
(rewrite "assoc")
```

Even quicker is the following approach: using the command (p. 89 f. of `prover-guide.ps`)

```
(auto-rewrite "assoc")
```

we tell PVS that it should consider all instances of `assoc` as automatic rewrite rules. After that command, `(grind)` completes the task.

Uniqueness of neutral elements We postulate a neutral element by adding

```
e : M
neutral_left : AXIOM
  FORALL (x:M) : e*x=x
neutral_right : AXIOM
  FORALL (x:M) : x*e=x
```

Our goal is

```
neutral_unique : THEOREM
  FORALL (e1:M) :
    (FORALL (x:M) : e1*x=x) AND
    (FORALL (x:M) : x*e1=x) IMPLIES e=e1
```

Here it is useful to first get an idea of how this proof should be done informally:

If e_1 is also a neutral element then $e = e * e_1$. By neutrality of e the right hand side equals e_1 and we're done.

In PVS after (skolem!) and (flatten) or, more compactly, (skosimp), we get

```
{-1}  FORALL (x: M) : e1!1 * x = x
{-2}  FORALL (x: M) : x * e1!1 = x
      |-----
{1}   e = e1!1
```

We want to expand e as $e * e_1!1$ using -2. We instantiate...

```
(inst -2 "e")
```

```
{-1}  FORALL (x: M) : e1!1 * x = x
{-2}  e * e1!1 = e
      |-----
{1}   e = e1!1
```

and replace

```
(replace -2 1 :dir RL)
```

bringing us to

```

[-1]  FORALL (x: M) : e!1 * x = x
[-2]  e * e!1 = e
      |-----
{1}   e * e!1 = e!1

```

which is an instance of `neutral_left`. The way to convince PVS of this is

```
(use "neutral_left")
```

A more pedestrian way would be to use `lemma` and `inst`.

A slightly quicker proof After `(skosimp)` we can use `(rewrite-with-fnum -2 ("x" "e") :dir RL)` to achieve the expansion of the left hand side. The command `rewrite-with-fnum`, is like `rewrite`, so doesn't normally require a substitution (instantiation). In this case, we have to give it because otherwise the replacement is applied to the right hand side, too!.

This brings us to

```

[-1]  FORALL (x: M) : e!1 * x = x
[-2]  FORALL (x: M) : x * e!1 = x
      |-----
{1}   e * e!1 = e!1

```

At which point we conclude using `(use "neutral_left")`.

I couldn't find a more efficient proof of that one. Can you?

Invertible elements An element of `M` is invertible if it has an inverse:

```
Invertible(x:M) : boolean = EXISTS(y:M) : x*y=e AND y*x = e
```

We can prove that the neutral element is invertible:

```

inv_neutral : THEOREM
  Invertible(e)

```

We see here, how, abbreviations a.k.a. definitions are introduced.

After invoking the prover the first command must be `expand "Invertible"` to open the definition. This brings us to

```

|-----
{1}  EXISTS (y: M) : e * y = e AND y * e = e

```

Now we have to come up with an alleged inverse to e . Surprise, it's going to be e itself.

```
(inst 1 "e")

|-----
{1}   e * e = e AND e * e = e
```

Here we could also have used `(inst? 1)` which would leave it to PVS to find the correct instantiation. It sometimes does....

We conclude with `(rewrite "neutral_left")` followed by `(split)`.
The last theorem in this series is that invertibles are closed under product:

```
|-----
{1}   FORALL (x, y: M):
      Invertible(x) AND Invertible(y) IMPLIES Invertible(x * y)
```

`(skosimp)` then `(expand "Invertible")` brings us to

```
{-1}  EXISTS (y: M): x!1 * y = e AND y * x!1 = e
{-2}  EXISTS (y: M): y!1 * y = e AND y * y!1 = e
|-----
{1}   EXISTS (y: M): x!1 * y!1 * y = e AND y * (x!1 * y!1) = e
```

Notice the then “strategy” (p. 111 of `prover-guide.ps`). It sequences commands.

Before being able to instantiate 1 , i.e., come up with an alleged inverse to $(x!1 * y!1)$ we must “open” the assumptions, i.e., introduce fresh constants for the inverses of $x!1$ and $y!1$, respectively, which are guaranteed by -1 and -2 .

I find it better to give suggestive names to these, so we do

```
(skolem -1 "xinv") then (skolem -2 "yinv") then (flatten)
```

to get

```
{-1}  x!1 * xinv = e
{-2}  xinv * x!1 = e
{-3}  y!1 * yinv = e
{-4}  yinv * y!1 = e
|-----
[1]   EXISTS (y: M): x!1 * y!1 * y = e AND y * (x!1 * y!1) = e
```

Now, we have to think a bit as to what the inverse to $x!1 * y!1$ should be. Well, thinking of $*$ as sequencing of “actions” it becomes clear that the inverse ought to be $yinv * xinv$. That’s the “rule of sock and shoe”. Therefore, `(inst 1 "yinv*xinv")` is the command of choice.

```
{-1} x!1 * xinv = e
{-2} xinv * x!1 = e
{-3} y!1 * yinv = e
{-4} yinv * y!1 = e
  |-----
[1] x!1 * y!1 * (yinv * xinv) = e AND (yinv * xinv) * (x!1 * y!1) =
```

Relying on PVS’ cleverness and doing `(inst? 1)` isn’t a good idea here.

Splitting (\wedge -R) brings us two subgoals of which we’ll only treat the first here:

```
{-1} x!1 * xinv = e
{-2} xinv * x!1 = e
{-3} y!1 * yinv = e
{-4} yinv * y!1 = e
  |-----
[1] x!1 * y!1 * (yinv * xinv) = e
```

Now we must first “rebracket” our goal to

$$x!1 * (y!1 * yinv) * xinv = e$$

While this can certainly be done by successive application of associativity, it is easier to just claim this and prove it separately. To do this, we issue the command

```
(case "x!1 * (y!1 * yinv) * xinv = e")
```

This presents us with two subgoals. One asking us to prove our goal under the extra assumption of the “claim”:

```
[-1] x!1 * (y!1 * yinv) * xinv = e
[-2] x!1 * xinv = e
[-3] xinv * x!1 = e
[-4] y!1 * yinv = e
[-5] yinv * y!1 = e
  |-----
[1] x!1 * y!1 * (yinv * xinv) = e
```


This follows from associativity, so

```
(auto-rewrite "assoc") then (grind)
```

does the job. Next, we must prove our claim:

```
[-1]  x!1 * xinv = e
[-2]  xinv * x!1 = e
[-3]  y!1 * yinv = e
[-4]  yinv * y!1 = e
      |-----
[1]   x!1 * (y!1 * yinv) * xinv = e
[2]   x!1 * y!1 * (yinv * xinv) = e
```

The old goal is still there, we can remove it with `(delete 2)` corresponding to rule `WEAK-R`. The rest is a rewriting consequence of `neutral_left` and `neutral_right`, so we install these and `grind`.

We could have turned off associativity with the command `(stop-rewrite "assoc")`, but this wasn't even necessary here.

8 Recursive functions

Many function definitions in mathematics, programming, and more so program specification are recursive.

Even if—for the sake of efficiency—the actual program uses an iterative solution, for specification and verification a recursive definition is usually more convenient.

Examples of recursive definitions Sum in pattern-matching notation:

$$\begin{aligned}\sum_{i=0}^0 a_i &= a_0 \\ \sum_{i=0}^{n+1} a_i &= a_{n+1} + \sum_{i=0}^n a_i\end{aligned}$$

Sum in fixpoint notation:

$$\begin{aligned}\sum_{i=0}^n a_i &= \text{if } n = 0 \\ &\text{then } a_0 \\ &\text{else } a_n + \sum_{i=0}^{n-1} a_i \text{ endif}\end{aligned}$$

Binary search in pattern-matching notation:

$$\begin{aligned} \mathit{find}(a, \text{null}) &= \text{ff} \\ \mathit{find}(a, \text{cons}(b, l)) &= (a=b) \vee (a \leq b \wedge \mathit{find}(a, \text{left}(l))) \vee (a > b \wedge \mathit{find}(a, \text{right}(l))) \end{aligned}$$

Binary search in fixpoint notation:

$$\begin{aligned} \mathit{find}(a, l) &= \text{if } l = [] \\ &\text{then ff} \\ &\text{else } a = \text{car}(l) \vee \\ &\quad (a \leq b \wedge \mathit{find}(a, \text{left}(\text{cdr}(l)))) \vee \\ &\quad (a > b \wedge \mathit{find}(a, \text{right}(\text{cdr}(l)))) \text{ endif} \end{aligned}$$

These clauses define honest-to-goodness functions on natural numbers and lists (or arrays). As you probably know this need not always be the case. For one thing, recursively defined functions may be partial ($f(n) = f(n)$), for another, some equations may not define a function at all.

$$\begin{aligned} f(n) &= f(n) \\ g(0) &= 0 \\ g(n+2) &= g(n) \\ g(1) &= \min\{g(2n) \mid n \in \mathbb{N}\} \\ h(n) &= 0 \\ h(n) &= 1 \end{aligned}$$

Fixpoint form Fix a_0, a_1, \dots and let $\mathit{sum} : \mathbb{N} \rightarrow \mathbb{R}$ be the function defined by $\mathit{sum}(n) = \sum_{i=0}^n a_i$.

We have $\mathit{sum}(n) = F(\mathit{sum}, n)$ where

$$\begin{aligned} F(f, n) &= \text{if } n = 0 \\ &\text{then } a_0 \\ &\text{else } a_n + f(n-1) \text{ endif} \end{aligned}$$

Exercise: define $\mathit{fact}(n) = n!$. Give F such that $\mathit{fact}(n) = F(\mathit{fact}, n)$.

In PVS all functions are total and therefore, general recursive function definitions are not permitted. Rather, an explicit measure must be provided ensuring that the definition terminates.

Theorem (Well-founded recursion): Let A, B be nonempty sets, let $F : (A \rightarrow B) \rightarrow (A \rightarrow B)$ be a functional, $w : A \rightarrow \mathbb{N}$ be a function (the “measure”).

Suppose that for each $f : A \rightarrow B$ and $a \in A$ the value $F(f, a)$ depends only on those values $f(a')$ for which $w(a') < w(a)$, that is to say

$$\forall a:A. \forall f, g:A \rightarrow B. (\forall x:A. w(x) < w(a) \Rightarrow f(x)=g(x)) \Rightarrow F(f, a)=F(g, a)$$

Then there exists a uniquely determined function $f_F : A \rightarrow B$ such that

$$\forall a:A. f_F(a) = F(f_F, a)$$

Proof. Let b_0 be a fixed element of B .

We define $f_F(a)$ by induction on $w(a)$. Suppose $w(a) = 0$. Then $F(f, a)$ is independent of f , so we can put $f_F(a) = F(f, a)$ where f is an arbitrary function from A to B , e.g. a constant one.

Suppose that $f_F(x)$ has already been defined for all x with $w(x) < n$ and that $w(a) = n$. Then we define a function $f : A \rightarrow B$ by

$$f(x) = \begin{cases} f_F(x), & \text{if } w(x) < n \\ b_0, & \text{if } w(x) \geq n \end{cases}$$

We then put $f_F(a) \stackrel{\text{def}}{=} F(f, a)$.

This procedure defines $f_F(a)$ for all values a .

Next, we show that $F(f_F, a) = f_F(a)$ for all a . Well, given a fixed but arbitrary element $a \in A$ (PVS would call it $a!1$) we see that $f_F(a)$ has been defined as $F(f, a)$ where f is the function which agrees with f_F on values x with $w(x) < n$ and is b_0 elsewhere.

But we have assumed that $F(f_F, a) = F(f, a)$ in this case.

For uniqueness we argue as follows. Suppose that $F(g, a) = g(a)$ for some function $g : A \rightarrow B$. We show by induction on $w(a)$ that $f(a) = g(a)$. The details are left to the reader. \square

In many examples the evaluation of $F(f, a)$ proceeds by evaluating the function f on a fixed number of arguments a_1, \dots, a_n depending only on a and having measure smaller than a , i.e., $w(a_i) < w(a)$. This was in particular the case for the *sum* and *fact* example.

Questions What would be an appropriate measure for the definition of *find*?

What is an appropriate measure for

```

$$F_{merge}(f, l_1, l_2) =$$

$$\begin{aligned} & \text{if } l_1 = \text{null} \\ & \quad \text{then } l_2 \\ & \text{elseif } l_2 = \text{null} \\ & \quad \text{then } l_1 \\ & \text{else } \text{cons}(\text{car}(l_1), \text{cons}(\text{car}(l_2), f(\text{cdr}(l_1), \text{cdr}(l_2)))) \end{aligned}$$

```

where $A = \text{list}[\text{nat}] \times \text{list}[\text{nat}]$ and $B = \text{list}[\text{nat}]$.

Hint: you may assume a function $\text{length} : \text{list}[\text{nat}] \rightarrow \text{nat}$.

8.1 Defining functions in PVS

We have already seen the definition of a predicate, namely `Invertible`. For PVS such a predicate is nothing but a function to the type `boolean`.

Using the same syntax we can define other functions like so:

```
f(x, y: nat) : nat = (x+y) * (x-y)
```

and we can prove

```
a: THEOREM f(5, 3) = 16
```

using `(grind)`. This method also does simple algebra:

```
b: THEOREM FORALL(x, y: nat) : f(x, y) = x^2 - y^2
```

Anyway, I'm getting distracted from today's topic: recursive definitions. Here is how we define `sum` in PVS provided `a : [nat->real]` has been defined or declared:

```
sum(n: nat) : RECURSIVE real =
  IF n=0 THEN a(0) ELSE a(n) + sum(n-1) ENDIF
MEASURE n
```

Try to memorise the slightly awkward syntax: the keyword `RECURSIVE` goes between the colon and the result type. And don't forget the measure either. It's supposed to go down as you unfold the recursion.

8.2 TCCs

When PVS typechecks such a definition (and this takes place before you enter the prover) it attempts to show that this is the case (the measure going down, that is). If it doesn't succeed a typechecking condition (TCC) is generated which you would then have to prove interactively using the prover.

You can display the TCC with the command `M-x show-tccs`. In the example at hand the TCCs are simple enough

```
sum_TCC1: OBLIGATION FORALL (n: nat): NOT n = 0 IMPLIES n - 1 >= 0;
sum_TCC2: OBLIGATION FORALL (n: nat): NOT n = 0 IMPLIES n - 1 < n;
```

The first one comes from the use of `n-1`. The type `nat` is in fact a subtype of the integers which is a subtype of the rationals, etc. A priori the minus function returns an integer. In the situation at hand, we know that `n` is not zero, so `n-1` is in fact a natural number. PVS was able to “prove” that by itself.

The other TCC comes from the recursion. We must show that the measure of the argument of the recursive call (here `n-1`) is smaller than the measure of the current argument (here `n`). Again, PVS can prove that itself.

We can now prove (using `grind`) simple goals which follow directly from the recursive definition like

```
c: THEOREM
sum(5) = a(0) + a(1) + a(4) + a(3) + a(2) + a(5)
```

We come to more interesting goals below.

8.2.1 Higher-order functions

Function types in PVS are like any other type. We can use this feature to pass the sequence `a: [nat->nat]` as an extra argument to `sum`:

```
sum(a: [nat->real], n: nat) : RECURSIVE real =
  IF n=0 THEN a(0) ELSE a(n)+sum(a, n-1) ENDIF
MEASURE n
```

The old definition of `sum` applies when the first argument isn't a function, this is an instance of *overloading*.

We don't even have to delete the previous definition of `sum`. PVS can tell the two apart by their types (this is known as *overloading*).

Now, we can apply `sum` to concrete functions, e.g., we might define

8.2.2 Examples

```
id(x:nat):nat = x
d : THEOREM
  sum(id,5) = 25
e : THEOREM
  sum(LAMBDA(x:nat):x*x,4) = 30
```

```
id(x:nat):nat = x
```

and then prove

```
d : THEOREM
  sum(id,5) = 25
```

If we don't want to sacrifice a name for the argument function we must use a lambda abstraction

```
e : THEOREM
  sum(LAMBDA(x:nat):x*x,4) = 30
```

Summary:

- PVS allows for definition of functions by well-founded recursion
- Such definitions generate proof obligations known as typechecking conditions (TCCs)
- TCCs also arise in conjunction with subtypes. More later.
- Within one and the same theory you can have several functions of the same name if their argument types are distinct (overloading)
- Functions can be arguments as well as results of functions. The LAMBDA notation allows one to construct function terms on the fly to be passed as argument to another function.
- PVS knows that a recursively defined function satisfies its defining equations

9 Proof by induction and higher-order logic

So far we have proved simple consequences of the recursive equations in which the recursive argument was a concrete value. If we want to prove more interesting universally quantified statements then we need a more powerful principle: proof by induction.

You probably have seen induction already: to prove a statement $\phi(n)$ for all natural numbers n you must prove it for 0 and then—assuming a fixed but arbitrary n' —you must prove it for $n' + 1$ under the extra assumption that it $\phi(n')$ holds.

In first-order logic:

$$\phi(0) \wedge (\forall n:\mathbb{N}.\phi(n) \Rightarrow \phi(n + 1)) \Rightarrow \forall n:\mathbb{N}.\phi(n)$$

In higher-order logic:

$$\forall \phi:[\mathbb{N} \rightarrow \text{boolean}].\phi(0) \wedge (\forall n:\mathbb{N}.\phi(n) \Rightarrow \phi(n + 1)) \Rightarrow \forall n:\mathbb{N}.\phi(n)$$

As a formula this *induction scheme* looks as follows:

$$\phi(0) \wedge (\forall n:\mathbb{N}.\phi(n) \Rightarrow \phi(n + 1)) \Rightarrow \forall n:\mathbb{N}.\phi(n)$$

In first-order logic we need one such formula for every predicate ϕ . In *higher-order logic* we can quantify over ϕ just as we quantify over individuals:

$$\forall \phi:[\mathbb{N} \rightarrow \text{boolean}].\phi(0) \wedge (\forall n:\mathbb{N}.\phi(n) \Rightarrow \phi(n + 1)) \Rightarrow \forall n:\mathbb{N}.\phi(n)$$

The proof rules for higher-order logic are essentially the same as those for first-order logic. Only the ways to form formulas are extended. A formula is just a term of type `boolean` and these can be formed using the connectives and quantifiers as well as by function application.

Semantically, the type `boolean` is interpreted as the set of truth values $\{\text{tt}, \text{ff}\}$; function types are interpreted as sets of all functions. Unlike pure first-order logic, higher-order logic do not admit complete proof systems. The reason is Gödel's incompleteness theorem which you may have come across in popular science books.

Higher-order logic can thus be defined as first-order logic with

- Type `boolean` and types closed under $[A_1, \dots, A_n \rightarrow B]$ (n ary function space)
- Unary predicate on type `boolean`, i.e., every term of type `boolean` can be seen / is a proposition

- $n + 1$ ary function symbols for application of functions to arguments.

Intended semantics: interpret `bool` as $\{\text{tt}, \text{ff}\}$, function spaces as sets of functions. No complete axiomatisation exists.

Approximations: λ -terms witnessing the existence of certain functions and predicates, β -equations, extensionality, comprehension axioms (or constants for quantifiers, connectives, functions), choice axioms.

Complete for non-standard models: Henkin models, toposes.

Rather than formally defining higher-order logic with its syntax and proof rules we will introduce it in PVS and get to know it by example.

Here is how the induction axiom is formulated in PVS:

```
nat_induction : LEMMA
  FORALL (p: [nat->boolean]) :
    (p(0) AND (FORALL j: p(j) IMPLIES p(j+1)))
      IMPLIES (FORALL i: p(i))
```

This is proved from a slightly more general axiom (well-foundedness of `<` on `nat`) in `lib/prelude.pvs`.

We are not so much interested in how to prove induction, but how to use it in order to prove other things.

Consider the classic $\sum_{i=0}^n i = n(n+1)/2$. We write $\phi(n) \equiv \sum_{i=0}^n i = n(n+1)/2$. We have $\phi(0) \equiv 0 = 0(0+1)/2$. True. We have $\phi(n_0+1) \equiv (n_0+1) + \sum_{i=0}^{n_0} i = (n_0+1)(n_0+2)/2$.

Using the induction hypothesis $\phi(n_0)$ this rewrites to $(n_0+1) + n_0(n_0+1)/2 = (n_0+1)(n_0+2)/2$ which is true by simple arithmetic.

Now we want to do the same thing in PVS:

We start with

```
|-----
{1}  FORALL (n: nat): sum(LAMBDA (i: nat): i, n) = n * (n + 1) / 2
```

Detailed proof We bring in `nat_induction`

```
Rule? (lemma "nat_induction")
```

```
{-1}  FORALL (p: pred[nat]):
      (p(0) AND (FORALL j: p(j) IMPLIES p(j + 1)))
        IMPLIES (FORALL i: p(i))
```



```

|-----
[1]  FORALL (n: nat):
      sum(LAMBDA (i: nat): i, n) = n * (n + 1) / 2

```

Now `nat_induction` works for all predicates `p`, but we need it for a particular one, so we must instantiate. But what with? We could have introduced an abbreviation for the predicate we're interested in but it's too late for that now, so we must resort to the lambda notation:

```

Rule? (inst -1 "LAMBDA(n:nat):
      sum(LAMBDA (i: nat): i, n) = n * (n + 1) / 2")

{-1}  ((LAMBDA (n: nat): sum(LAMBDA (i: nat): i, n) =
      n * (n + 1) / 2)(0) AND
      (FORALL j:
        (LAMBDA (n: nat): sum(LAMBDA (i: nat): i, n) =
          n * (n + 1) / 2)(j) IMPLIES
          (LAMBDA (n: nat): sum(LAMBDA (i: nat): i, n) =
            n * (n + 1) / 2)(j + 1)))
      IMPLIES
      (FORALL (i_49: nat): (LAMBDA (n: nat): sum(LAMBDA (i: nat): i,
      n) =
        n * (n + 1) / 2)(i_49))
|-----
[1]  FORALL (n: nat): sum(LAMBDA (i: nat): i, n) = n * (n + 1) / 2

```

That looks daunting. What has happened is that every occurrence of `p` in `-1` has been literally replaced by the instantiation we provided. But we would like a bit more than just that: we would like to see arguments to `p` such as `0` in the base case and `i+1` in the induction step to be plugged in for the lambda-bound variable `n`. This is known as beta reduction and we perform it in PVS by issuing the command `(beta)`

```

Rule? (beta)
{-1}  (sum(LAMBDA (i: nat): i, 0) = 0 * (0 + 1) / 2 AND
      (FORALL j:
        sum(LAMBDA (i: nat): i, j) = j * (j + 1) / 2 IMPLIES
          sum(LAMBDA (i: nat): i, j + 1) = (j + 1) * (j + 1 + 1) / 2
      IMPLIES (FORALL (i_49: nat):

```

```

          sum(LAMBDA (i: nat): i, i_49) = i_49 * (i_49 + 1) / 2)
    |-----
[1]   FORALL (n: nat): sum(LAMBDA (i: nat): i, n) = n * (n + 1) / 2

```

What we have achieved so far is that the required instance of the induction scheme is among our antecedents. We now want to use it which shouldn't be difficult as it ends exactly with what we want to prove. So we use rule \Rightarrow -L which is of the `split` kind:

```
Rule? (split)
```

This yields two subgoals, one asking us to deduce 1 from the conclusion of $\neg 1$ —that's a propositional axiom, so PVS won't bother presenting us with it—, and another one asking us to prove the premise of $\neg 1$ (or 1 again). Since the premise of $\neg 1$ is a conjunction (of base case and induction step) we must use \wedge -R which is again of the `split` kind and in fact PVS has already performed this at the last `split` command so that we don't need to enter it again. We are thus presented with two actual subgoals:

```

    |-----
{1}   sum(LAMBDA (i: nat): i, 0) = 0 * (0 + 1) / 2
[2]   FORALL (n: nat):
      sum(LAMBDA (i: nat): i, n) = n * (n + 1) / 2

```

and

```

    |-----
{1}   sum(LAMBDA (i: nat): i, 0) = 0 * (0 + 1) / 2
[2]   FORALL (n: nat):
      sum(LAMBDA (i: nat): i, n) = n * (n + 1) / 2

```

The first of these (the base case) follows by simple arithmetic: (`grind`) disposes of it.

We could also have proved this by hand using elementary properties of real arithmetic summarised in `prelude.pvs`.

The induction step is more interesting. We first delete 2 and then introduce a “fixed but arbitrary” name, say `n!0` by

```
Rule? (skolem 1 "n!0")
```

```

    |-----
{1}   sum(LAMBDA (i: nat): i, n!0) = n!0 * (n!0 + 1) / 2 IMPLIES
      sum(LAMBDA (i: nat): i, n!0 + 1) = (n!0 + 1) * (n!0 + 1 + 1) /

```

This being an implication we use `⇒-L` or `(flatten)` to give

```
{-1} sum(LAMBDA (i: nat): i, n!0) = n!0 * (n!0 + 1) / 2
|-----
{1}  sum(LAMBDA (i: nat):
      i, n!0 + 1) = (n!0 + 1) * (n!0 + 1 + 1) / 2
```

This looks like what we had expected. We may assume that what we want to show holds for a fixed but arbitrary `n!0` and from that we must show it for `n!0+1`. Let's expand the sum in the succedent:

Rule? `(expand "sum" 1)`

```
[-1] sum(LAMBDA (i: nat): i, n!0) = n!0 * (n!0 + 1) / 2
|-----
{1}  1 + sum(LAMBDA (i: nat): i, n!0) + n!0 =
      (2 + n!0 + (n!0 * n!0 + 2 * n!0)) / 2
```

Notice that with recursive definitions the `expand` command performs one recursive unfolding rather than replacing `sum` with its definition using the keyword `RECURSIVE`.

We now recognise the left-hand-side of `-1` as a subterm, it's therefore a good idea to replace it with the right hand side:

Rule? `(replace -1 1)`

```
[-1] sum(LAMBDA (i: nat): i, n!0) = n!0 * (n!0 + 1) / 2
|-----
{1}  1 + n!0 * (n!0 + 1) / 2 + n!0 =
      (2 + n!0 + (n!0 * n!0 + 2 * n!0)) / 2
```

The conclusion is an arithmetic identity so `(grind)` can establish it.

Quicker proof The first few steps were rather awkward and independent of the particular goal at hand. For this reason the command `induct` has been provided which performs them all in one go. In the situation

```
|-----
{1}  FORALL (n: nat): sum(LAMBDA (i: nat): i, n) = n * (n + 1) / 2
```

the command `(induct "n")` produces two subgoals:

```
|-----
{1}  sum(LAMBDA (i: nat): i, 0) = 0 * (0 + 1) / 2
```

and

```
|-----
{1}  FORALL j:
      sum(LAMBDA (i: nat): i, j) = j * (j + 1) / 2 IMPLIES
      sum(LAMBDA (i: nat): i, j + 1) =
      (j + 1) * (j + 1 + 1) / 2
```

which we deal with as before. Actually, the second one can be proved simply with the command `(skosimp!) then (grind)`.

An even quicker proof goes with the single command `(induct-and-simplify "n")`.

PVS is surprisingly good at doing inductive proofs almost automatically. For example, Cassini's identity

$$F_{n+2}F_n - F_{n+1}^2 = (-1)^n$$

can be proved with a single "induct-and-simplify". One should define $(-1)^n$ recursively.

The following is known as Abel's lemma and plays a role in number theory, more specifically, Dirichlet series.

Let (a_n) and (b_n) be two sequences. Put:

$$A_{m,p} = \sum_{n=m}^{n=p} a_n \text{ and } S_{m,m'} = \sum_{n=m}^{n=m'} a_n b_n$$

Then one has:

$$S_{m,m'} = A_{m,m'} b_{m'} + \sum_{n=m}^{n=m'-1} A_{m,n} (b_n - b_{n+1})$$

Here one must induct not on a single quantity, but rather the difference $m' - m$. PVS provides the command `measure-induct+` for that purpose.

10 Lists

A datatype of finite list over arbitrary type of entries is predefined.

If t is a type then `list [t]` is the type of finite lists with entries over t .

Semantically, elements of `list [t]` take the form $[x_1, x_2, \dots, x_n]$ where the x_i are elements of type t .

The constant `null` denotes the empty list, the function

```
cons : [t, list [t] -> list [t]]
```

tacks an element on to the beginning of a list. For example, if l is the list $[3, 4, 1, 2]$ then `cons(5, l)` is $[5, 3, 4, 1, 2]$.

The type of lists has the subtype `(cons? [t])`, consisting of all non-empty lists. The function `car` takes a nonempty list and returns its first element, i.e., we have

```
car : [(cons? [t]) -> t]
```

the function `cdr` takes a nonempty list and returns its “tail”, i.e., the (possibly empty) list obtained by stripping off its first element.

```
cdr : [(cons? [t]) -> list [t]]
```

For example, if $l = [5, 4, 3, 2]$ then

$$\begin{aligned}\text{car}(l) &= 5 \\ \text{cdr}(l) &= [4, 3, 2]\end{aligned}$$

We also have a predicate `null?` which tells whether a list is empty

```
null? : [list [t] -> boolean]
```

and another—less used—predicate `cons?` which tells whether a list is nonempty. In fact the subtype `(cons? [t])` is derived from that predicate and more generally, if $p : [t \rightarrow \text{boolean}]$ then `(p)` is the subtype of t consisting of those elements for which p holds.

In practice, the functions `car` and `cdr` are applied to arguments of type `list [t]` rather than `(cons? [t])`.

Such application generates a typechecking condition (TCC) which either we or PVS has to prove. A typical case, when such a TCC can be proved automatically, is when `car(l)` or `cdr(l)` is used in the `ELSE` branch of a conditional `IF null?(l)`.

Another typical such case is a usage like `cdr(cons(x, l))` which automatically simplifies to `l`. Similarly, `car(cons(x, l))` automatically simplifies to `x`.

Subtypes If $P : [t \rightarrow \text{bool}]$ then (P) is a type.

If $a : t$ and $P(a)$ holds then also $a : (P)$.

If we implicitly assert $a : (P)$ then a proof obligation (TCC) $P(a)$ arises.

10.1 Recursion on lists

There is a predefined function

```
length : [list[t] -> nat]
```

allowing us to define functions on lists by recursion (using) `length` as a measure.

Here is a definition of the function which appends two lists.

```
t :TYPE
```

```
append(l1:list[t], l2:list[t]) : RECURSIVE list[t] =  
IF null?(l1)  
  THEN l2  
  ELSE cons(car(l1), append(cdr(l1), l2))  
ENDIF  
MEASURE length(l1)
```

This means that the function `append` satisfies the following two defining equations:

```
append(null, l2) = l2  
append(cons(x, l1), l2) = cons(x, append(l1, l2))
```

In fact, PVS provides a cases-construct allowing us to write `append` in this slightly more perspicuous form. See `prelude.pvs` or the documentation.

Here are a few more function definitions:

```
occ(x:t, l:list[t]) : RECURSIVE nat =  
  IF null?(l)  
    THEN 0  
    ELSIF x=car(l) THEN occ(x, cdr(l))+1 ELSE occ(x, cdr(l)) ENDIF  
  MEASURE length(l)
```

```
filter(p:[t->boolean], l:list[t]) : RECURSIVE list[t] =  
  IF null?(l)
```

```

        THEN null
        ELSIF p(car(l)) THEN cons(car(l), filter(p, cdr(l)))
        ELSE filter(p, cdr(l)) ENDIF
    MEASURE length(l)

rev(l:list[t]) : RECURSIVE list[t] =
    IF null?(l) THEN null
    ELSE append(rev(cdr(l)), cons(car(l), null))
    ENDIF
    MEASURE length(l)

revl(l:list[t], acc:list[t]) : RECURSIVE list[t] =
    IF null?(l) THEN acc
    ELSE revl(cdr(l), cons(car(l), acc))
    ENDIF
    MEASURE length(l)

```

occ(x, l) returns the number of occurrences of x in the list l; filter(p, l) returns the list consisting of those elements of l which satisfy the predicate p. rev(l) returns the reversal of list l and revl(l, acc), finally, returns the reversal of l followed by acc.

10.2 Reduce

The length function itself admits a recursive definition:

```

length(l:list[t]) : RECURSIVE nat =
    IF null?(l)
    THEN 0
    ELSE 1 + length(cdr(l))
    ENDIF
    MEASURE ???

```

The trouble is that the only reasonable measure function is length itself. As I said, length is fortunately predefined (in prelude.pvs), but nevertheless it's worth knowing how, namely using the reduce_nat functional, which now really is basic. If

```

null_case : nat
cons_case : [t, nat -> nat]

```

are given terms of the indicated types then

```
reduce_nat (null_case, cons_case) : [list[t] -> nat]
```

is (as indicated) a function from lists to natural numbers, namely the function f defined recursively by

```
f(null) = null_case  
f(cons(x, l)) = cons_case(x, f(l))
```

It is clear that whatever `null_case` and `cons_case` are this defines a unique total function so that `reduce_nat` is justified.

The length function can now be defined as

```
length : [list[t] -> nat] =  
  reduce_nat(0, LAMBDA(x:t, prev:nat): prev+1)
```

More generally, we have a construct `reduce` which allows for the definition of functions on lists with result type other than `nat` by such structural recursion; however, these are more easily defined using well founded recursion with `length` as the measure, and indeed, if you find `reduce` confusing, just take `length` for granted and define all your functions using `length` or derived forms as measure.

11 Proof by list induction

Also predefined is the following induction principle for lists

```
list_induction: AXIOM  
  FORALL (p: [list -> boolean]):  
    (p(null) AND  
     (FORALL (cons1_var: T, cons2_var: list):  
       p(cons2_var) IMPLIES p(cons(cons1_var, cons2_var))))  
    IMPLIES (FORALL (list_var: list): p(list_var))
```

which states that a property p which

1. holds for the empty list
2. holds for an arbitrary list of the form `cons(x, l)` provided it holds for `l`

holds for all lists.

This principle can be invoked just as `nat_induction` using `lemma`, `inst`, `beta`, etc. or using the `induct` command, or, indeed, using the powerful `induct-and-simplify` command.

11.1 Associativity of `append`

Let's do an example proof: associativity of `append`

```
|-----  
{1}  FORALL (l1, l2, l3: list[t]):  
      append(append(l1, l2), l3) = append(l1, append(l2, l3))
```

We need to do induction on one of `l1`, `l2`, `l3`.

Induction on `l1` will make a defining equation for `append` applicable which (in the `cons` case) brings even the outermost `append` into a rewritable form, so that looks promising: We invoke `(induct "l1")` and get two subgoals, the first of which is

```
|-----  
{1}  FORALL (l2, l3: list[t]):  
      append(append(null, l2), l3) = append(null, append(l2, l3))
```

We could either try to induct on `l2` or `l3` here, or solve it directly which should intuitively be possible, as by virtue of the recursive equations both sides equal `append(l2, l3)`. The way to convince PVS of this is either to use `(grind)` or to introduce fresh names for the universal quantifier: `(skolem 1 ("l2!1" "l3!1"))`

```
|-----  
{1}  append(append(null, l2!1), l3!1) = append(null, append(l2!1, l3!
```

Now we want to replace the second and third occurrence of `append` by their definitions (as this will bring about a simplification), but not the first and fourth (as this will make things more complicated). The command `(expand "append" 1 2)` expands the second occurrence of `append` in formula 1. Using this again on occurrence 3 achieves our goal. Alternatively, we can use the command

```
(expand "append" :if-simplifies T)
```

which only expands those occurrences of `append` which result in a simplification (formally: those whose definition contains an if-then-else whose guard is equal to either `TRUE` or `FALSE`).

At any rate, we get

```
|-----
{1}   append(l2!1, l3!1) = append(l2!1, l3!1)
```

which is an instance of reflexivity and thus discharged.

Now, we get to see the second subgoal:

```
|-----
{1}   FORALL (cons1_var: t, cons2_var: list[t]):
      (FORALL (l2, l3: list[t]):
        append(append(cons2_var, l2), l3) =
          append(cons2_var, append(l2, l3)))
      IMPLIES
      (FORALL (l2, l3: list[t]):
        append(append(cons(cons1_var, cons2_var), l2), l3) =
          append(cons(cons1_var, cons2_var), append(l2, l3)))
```

We introduce fresh names for the universally quantified constants:

```
(skolem 1 ("x!1" "l!1"))
```

and `(flatten)` giving us

```
{-1}  FORALL (l2, l3: list[t]):
      append(append(l!1, l2), l3) = append(l!1, append(l2, l3))
|-----
{1}   FORALL (l2, l3: list[t]):
      append(append(cons(x!1, l!1), l2), l3) =
        append(cons(x!1, l!1), append(l2, l3))
```

Here `-1` is the induction hypothesis. We must now introduce skolem constants for the universal quantifier as in the base case with `(skolem 1 ("l2!1" "l3!1"))` and we can in fact at this point instantiate the induction hypothesis with these values: `(inst -1 "l2!1" "l3!1")`. This means that we cannot use the induction hypothesis with values other than these two. In general, this is risky as there are cases in which we have to use the induction hypothesis with other cleverly chosen instantiations, see Section 11.3 below. Here, however, it is safe. If in doubt, postpone instantiations as long as possible.

```

{-1}  append(append(l!1, l2!1), l3!1) = append(l!1, append(l2!1, l3!1))
      |-----
[1]   append(append(cons(x!1, l!1), l2!1), l3!1) =
      append(cons(x!1, l!1), append(l2!1, l3!1))

```

Now we see that we have a couple of instances of `append` which can be simplified by way of the recursive equations.

```
(expand "append" :if-simplifies T)
```

Remember that typing `(exp` followed by `M-s`, i.e., `Alt` `s` produces this command.

```

{-1}  append(append(l!1, l2!1), l3!1) = append(l!1, append(l2!1, l3!1))
      |-----
{1}   cons(x!1, append(append(l!1, l2!1), l3!1)) =
      cons(x!1, append(l!1, append(l2!1, l3!1)))

```

Now we discover the lhs of the induction hypothesis (-1) as a subterm and in fact the current sequent follows by mere equational reasoning. Therefore, it can be dispatched with `(grind)`. If we insist on doing it by hand we do `(replace -1 1)` leading to an instance of reflexivity.

This entire proof could also be done with the single command `(induct-and-simplify "l1")`.

11.2 Occurrences

Recall the function `occ` which counts the number of occurrences of a given element in a list. We want to prove:

```

      |-----
{1}   FORALL (x: t, l1, l2: list[t]):
      occ(x, append(l1, l2)) = occ(x, l1) + occ(x, l2)

```

Again, `(induct-and-simplify "l1")` will do the job, but for the sake of it we go for something more elementary. We start with `(induct "l1")`. As said before, there is no general rule as to whether one should just skolemise or use induction and if yes on which argument. Rule of thumb is that theorems involving recursively defined functions require proof by induction and that the induction should be on the argument which promises the most simplifications to happen.

Here this is `l1` so we do `(induct "l1")` giving us two subgoals, first the base case:

```

|-----
{1}  FORALL (x: t, l2: list[t]):
      occ(x, append(null, l2)) = occ(x, null) + occ(x, l2)

```

which after rewriting the append-term and the second occurrence of `occ` becomes an arithmetic identity. We dispatch the whole subgoal with `(grind)` bringing us to the second subgoal, the inductive step which after skolemising and flattening looks like so:

```

{-1}  FORALL (x: t, l2: list[t]):
      occ(x, append(l!1, l2)) = occ(x, l!1) + occ(x, l2)
|-----
{1}  FORALL (x: t, l2: list[t]):
      occ(x, append(cons(x!1, l!1), l2)) = occ(x, cons(x!1, l!1)) +

```

As before we introduce fresh names and instantiate our induction hypothesis with them giving us

```

{-1}  occ(x!2, append(l!1, l2!1)) = occ(x!2, l!1) + occ(x!2, l2!1)
|-----
[1]  occ(x!2, append(cons(x!1, l!1), l2!1)) =
      occ(x!2, cons(x!1, l!1)) + occ(x!2, l2!1)

```

Again, we find a number of occurrences of recursively defined functions which now admit a simplification:

```

(expand "occ" :if-simplifies T)
(expand "append" :if-simplifies T)
(expand "occ" :if-simplifies T)

{-1}  occ(x!2, append(l!1, l2!1)) = occ(x!2, l!1) + occ(x!2, l2!1)
|-----
{1}  IF x!2 = x!1
      THEN 1 + occ(x!2, append(l!1, l2!1))
      ELSE occ(x!2, append(l!1, l2!1))
      ENDIF
      =
      IF x!2 = x!1 THEN 1 + occ(x!2, l!1) ELSE occ(x!2, l!1) ENDIF +
      occ(x!2, l2!1)

```

We can now directly replace the lhs of the induction hypothesis with its rhs and end up with an instance of reflexivity thus completing the proof. For the sake of the example let's move the if-then-else constructs to the surface and proceed by case distinction: the command `(lift-if)` followed by `(split)` brings us two subgoals

```
[-1]  occ(x!2, append(l!1, l2!1)) = occ(x!2, l!1) + occ(x!2, l2!1)
      |-----
{1}   x!2 = x!1 IMPLIES
      1 + occ(x!2, append(l!1, l2!1)) = 1 + occ(x!2, l!1) + occ(x!2,
```

and

```
[-1]  occ(x!2, append(l!1, l2!1)) = occ(x!2, l!1) + occ(x!2, l2!1)
      |-----
{1}   NOT x!2 = x!1 IMPLIES
      occ(x!2, append(l!1, l2!1)) = occ(x!2, l!1) + occ(x!2, l2!1)
```

corresponding to the two branches of the conditional. The first subgoal is first flattened and then dispatched with `(grind)` as it is an equational consequence of the induction hypothesis. The second even becomes a propositional axiom after flattening.

11.2.1 Completeness of filtering

Of a similar kind is

```
filter_complete : THEOREM
  FORALL(x:t, l:list[t], p:[t->boolean]):
    p(x) IMPLIES occ(x, filter(p, l)) = occ(x, l)
```

Invoking induction on `l`, grinding away the base case, skolemising and instantiating brings us to

```
{-1}  p!1(x!2) IMPLIES occ(x!2, filter(p!1, l!1)) = occ(x!2, l!1)
[-2]  p!1(x!2)
      |-----
[1]   occ(x!2, filter(p!1, cons(x!1, l!1))) = occ(x!2, cons(x!1, l!1))
```

Simplifying the recursive function calls gives

```

[-1] p!1(x!2) IMPLIES occ(x!2, filter(p!1, l!1)) = occ(x!2, l!1)
[-2] p!1(x!2)
    |-----
{1}  occ(x!2,
      IF p!1(x!1)
      THEN cons(x!1, filter(p!1, l!1))
      ELSE filter(p!1, l!1)
      ENDIF)
      = IF x!2 = x!1 THEN 1 + occ(x!2, l!1) ELSE occ(x!2, l!1) ENDIF

```

(lift-if) followed by (split 1) followed by (flatten) gives two sub-goals the first of which is

```

[-1] p!1(x!1)
[-2] p!1(x!2) IMPLIES occ(x!2, filter(p!1, l!1)) = occ(x!2, l!1)
[-3] p!1(x!2)
    |-----
{1}  occ(x!2, cons(x!1, filter(p!1, l!1))) =
      IF x!2 = x!1 THEN 1 + occ(x!2, l!1) ELSE occ(x!2, l!1) ENDIF

```

This gives now the opportunity for another simplification:

```

[-1] p!1(x!1)
[-2] p!1(x!2) IMPLIES occ(x!2, filter(p!1, l!1)) = occ(x!2, l!1)
[-3] p!1(x!2)
    |-----
{1}  IF x!2 = x!1
      THEN 1 + occ(x!2, filter(p!1, l!1))
      ELSE occ(x!2, filter(p!1, l!1))
      ENDIF
      = IF x!2 = x!1 THEN 1 + occ(x!2, l!1) ELSE occ(x!2, l!1) ENDIF

```

and again we must lift the conditionals and split: to deal with this we need to invoke lift-if again and split giving us

```

[-1] x!2 = x!1
[-2] p!1(x!1)
[-3] p!1(x!2) IMPLIES occ(x!2, filter(p!1, l!1)) = occ(x!2, l!1)
[-4] p!1(x!2)

```

```

|-----
{1}  1 + occ(x!2, filter(p!1, l!1)) = 1 + occ(x!2, l!1)

```

Now we finally are in a position to use the induction hypothesis: `(split -3)` followed by `(replace -1 1)` dispatches this branch of the proof. The other ones are dealt with similarly. Of course at the displayed point and even before we could have used `(grind)` as well.

11.3 List reversal

Here is one way to reverse a list:

```

rev(l:list[t]) : RECURSIVE list[t] =
  IF null?(l) THEN null
    ELSE append(rev(cdr(l)), cons(car(l), null))
  ENDIF
  MEASURE length(l)

```

This is considered inefficient because the recursive implementation of the `append` function takes linear time hence the overall runtime is quadratic.

Better is the following tail recursive formulation of reversal:

```

rev1(l:list[t], acc:list[t]) : RECURSIVE list[t] =
  IF null?(l) THEN acc
    ELSE rev1(cdr(l), cons(car(l), acc))
  ENDIF
  MEASURE length(l)

```

The idea is that `rev1(l, acc)` equals the reversal of `l` followed by `acc`, so that we obtain the reversal of `l` as `rev1(l, null)`. Let's prove that this is true:

```

rev_rev1 : THEOREM
  FORALL(l, acc:list[t]):
    rev1(l, acc) = append(rev(l), acc)

```

Induction on `l`, grinding away the base case, skolemising and flattening brings us to

```

[-1]  FORALL (acc: list[t]): rev1(l!1, acc) = append(rev(l!1), acc)
|-----
{1}  rev1(cons(x!1, l!1), acc!1) = append(rev(cons(x!1, l!1)), acc!1)

```

Now in this case, it is wrong to instantiate the induction hypothesis with `acc!1` and that's the reason why `induct-and-simplify` doesn't work here. We just leave the induction hypothesis and simplify our recursive functions.

```
[-1]  FORALL (acc: list[t]): rev1(l!1, acc) = append(rev(l!1), acc)
      |-----
{1}   rev1(l!1, cons(x!1, acc!1)) =
      append(append(rev(l!1), cons(x!1, null)), acc!1)
```

Now we see that the lhs is in fact an instance of the lhs of the induction hypothesis, however with `acc` set to `cons(x!1, acc!1)`. This suggests to instantiate the induction hypothesis accordingly:

```
(inst -1 "cons(x!1, acc!1)")

[-1]  rev1(l!1, cons(x!1, acc!1)) = append(rev(l!1), cons(x!1, acc!1))
      |-----
{1}   rev1(l!1, cons(x!1, acc!1)) =
      append(append(rev(l!1), cons(x!1, null)), acc!1)
```

If the induction hypothesis holds for all `acc`, then in particular for the one we just gave ...

Now we rewrite with the induction hypothesis (`replace -1 1`) and get

```
[-1]  rev1(l!1, cons(x!1, acc!1)) = append(rev(l!1), cons(x!1, acc!1))
      |-----
{1}   append(rev(l!1), cons(x!1, acc!1)) =
      append(append(rev(l!1), cons(x!1, null)), acc!1)
```

This looks like an instance of associativity of `append` albeit slightly hidden. We could now use (`lemma "append_assoc"`) etc. but it is easier to tell PVS to consider it as a rewrite rule: (`auto-rewrite "append_assoc"`) after which (`grind`) can complete the proof.

11.4 Summary

To prove theorems involving recursively defined functions we usually need induction. The first decision to make is which variable to induct on. This should be the one leading to the most simplifications of recursively defined functions.

Once this decision has been made one can always try `induct-and-simplify`. If that doesn't help then we follow the following steps:

1. Invoke induction with the `induct` command
2. Skolemize and flatten
3. Simplify instances of recursively defined functions using `(expand ... :if-simplifies T)`
4. make case distinctions on conditionals using `lift-if` and `split`
5. try to massage your goal so that the induction hypothesis becomes applicable
6. when the induction hypothesis is universally quantified you must decide on the right instantiation. Often but not always it consists of the constants obtained from skolemising the current goal.

When you get stuck:

- Try to figure out what your current goal really says and what it could be a consequence of.
- Try not to get into a “symbol-pushing mode” and just blindly enter commands
- At least sketch an informal proof before you attempt a proof with PVS
- If your current goal seems true, but you can’t get PVS to prove it you may try to isolate it as a separate lemma, i.e., abandon the proof, type in the lemma, prove it separately (perhaps again using induction) and then retry.

12 General datatypes

Inductive datatypes other than lists can be defined in PVS using the datatype construct. Rather than going into formalities let’s look at two concrete examples:

12.1 Labelled binary trees

If T is a set then the set of binary trees `tree(T)` with labels in T is inductively defined as follows:

- `leaf` is a tree,

- if $label \in T$ and $left, right \in \text{tree}(T)$ then $\text{node}(label, left, right) \in \text{tree}(T)$.

E.g. $\text{node}(3, \text{node}(2, \text{leaf}, \text{leaf}), \text{leaf}) \in \text{tree}(\mathbb{N})$.

One may ask in what sense this *prima facie* self-referential explanation is at all a valid definition.

Some people just take it for granted, others prefer to explain it in terms of set theory. For instance, we can say that a tree is a finite prefix-closed set of paths, or we can define trees by induction on their depth level. Then the only tree of level 0 would be a leaf, encoded e.g. as the empty set; a tree of level $n + 1$ is either a tree of level n or a triple $(label, left, right)$ where $label \in T$ and $left, right$ are trees of level n . We then *define* the constructor `leaf` as being the empty set and `node` as the function which groups three things into a triple.

Since trees are inductively generated, we have the following principle of tree induction:

Let P be a property of binary T -labelled trees.

If

- $P(\text{leaf})$ and
- whenever $P(left)$ and $P(right)$ for some $left, right \in \text{tree}(T)$ then also $P(\text{node}(label, left, right))$ for all $label \in T$

then P holds for all binary T -labelled trees.

If we take inductive definitions for granted then we must also take this principle on board; if we define trees in terms of more primitive concepts then tree induction becomes a theorem, e.g., provable by course-of-values induction on the level.

Another principle is that nodes are different from leaves:

$$\text{node}(label, left, right) \neq \text{leaf}$$

- $\text{depth}(\text{leaf}) = 0$
- $\text{depth}(label, left, right) = \max(\text{depth}(left), \text{depth}(right)) + 1$

To define functions on trees we need some measure on them; the most primitive such is the depth given as on the slide. Again, we can either take the depth for granted or define it as the least level containing the tree. In that case the above equation could be proved.

Once we've got the depth we can define other functions by well-founded recursion:

- $\text{no_leaves}(\text{leaf}) = 1$
- $\text{no_leaves}(\text{node}(\text{label}, \text{left}, \text{right})) = \text{no_leaves}(\text{left}) + \text{no_leaves}(\text{right})$
- $\text{no_nodes}(\text{leaf}) = 0$
- $\text{no_nodes}(\text{node}(\text{label}, \text{left}, \text{right})) = \text{no_nodes}(\text{left}) + \text{no_nodes}(\text{right}) + 1$

This was the pattern-matching form. In order to get the fixpoint form we need (partial) destructor functions:

Define $\text{leaf?}(t) \iff t = \text{leaf}$,

Define $\text{node?}(t) \iff \exists \text{label}, \text{left}, \text{right}. t = \text{node}(\text{label}, \text{left}, \text{right})$,

Define $(\text{leaf?}) = \{\text{leaf}\}$

Define $(\text{node?}) = \{t \mid \text{node?}(t)\}$.

Notice: $\forall t:\text{tree}(T). \text{leaf?}(t) \vee \text{node?}(t)$

We have

$$\begin{aligned} \text{label} &: (\text{node?}) \rightarrow T \\ \text{left} &: (\text{node?}) \rightarrow \text{tree}(T) \\ \text{right} &: (\text{node?}) \rightarrow \text{tree}(T) \end{aligned}$$

defined by

$$\begin{aligned} \text{label}(\text{node}(\text{label}, \text{left}, \text{right})) &= \text{label} \\ \text{left}(\text{node}(\text{label}, \text{left}, \text{right})) &= \text{left} \\ \text{right}(\text{node}(\text{label}, \text{left}, \text{right})) &= \text{right} \end{aligned}$$

$$\text{no_leaves}(t) = \begin{cases} 1, & \text{if } \text{leaf?}(t) \\ \text{no_leaves}(\text{left}(t)) + \text{no_leaves}(\text{right}(t)), & \text{o/w} \end{cases}$$

$$\text{no_nodes}(t) = \begin{cases} 0, & \text{if } \text{leaf?}(t) \\ \text{no_nodes}(\text{left}(t)) + \text{no_nodes}(\text{right}(t)) + 1, & \text{o/w} \end{cases}$$

Both recursive definitions are well-founded using the depth as measure.

$\forall t:\text{tree}(T). \text{no_leaves}(t) = \text{no_nodes}(t) + 1$.

Proof by tree induction:

- $\text{no_leaves}(\text{leaf}) = 1 = 0 + 1 = \text{no_nodes}(\text{leaf}) + 1$

- $\text{no_leaves}(\text{node}(\text{label}, \text{left}, \text{right})) =$
 $\text{no_leaves}(\text{left}) + \text{no_leaves}(\text{right}) \stackrel{IH}{=}$
 $\text{no_nodes}(\text{left}) + 1 + \text{no_nodes}(\text{right}) + 1 =$
 $(\text{no_nodes}(\text{left}) + \text{no_nodes}(\text{right}) + 1) + 1 =$
 $\text{no_nodes}(\text{node}(\text{label}, \text{left}, \text{right})) + 1$

```
tree[t:TYPE]: DATATYPE
BEGIN
  leaf : leaf?
  node(label:t, left, right:tree) : node?
END tree
```

To introduce trees into PVS we use the above declaration. It will automatically generate a file `tree_adt.pvs` when we type check the file containing the declaration. The contents of `tree_adt.pvs` are as follows:

```
tree_adt[t: TYPE]: THEORY
BEGIN

  tree: TYPE
  leaf?, node?: [tree -> boolean]
  leaf: (leaf?)
  node: [[t, tree, tree] -> (node?)]
  label: [(node?) -> t]
  left: [(node?) -> tree]
  right: [(node?) -> tree]

  tree_label_node: AXIOM
    FORALL (node1_var: t, node2_var: tree, node3_var: tree):
      label(node(node1_var, node2_var, node3_var)) = node1_var;

  tree_left_node: AXIOM
    FORALL (node1_var: t, node2_var: tree, node3_var: tree):
      left(node(node1_var, node2_var, node3_var)) = node2_var;

  tree_right_node: AXIOM
    FORALL (node1_var: t, node2_var: tree, node3_var: tree):
      right(node(node1_var, node2_var, node3_var)) = node3_var;
```

```

tree_inclusive: AXIOM
  FORALL (tree_var: tree): leaf?(tree_var) OR node?(tree_var);

tree_induction: AXIOM
  FORALL (p: [tree -> boolean]):
    (p(leaf) AND
     (FORALL (node1_var: t, node2_var: tree, node3_var: tree):
       p(node2_var) AND p(node3_var) IMPLIES
       p(node(node1_var, node2_var, node3_var))))
    IMPLIES (FORALL (tree_var: tree): p(tree_var));

```

The file contains other useful stuff such as the definition of a subtree relation, a mapping functional, as well as properties describing these. Take a look yourselves!

Unfortunately, the depth isn't defined for us, so we do it ourselves using the `reduce_nat` functional which also works for trees with different typing, though. Can you work out the typing of `reduce_nat` from the example? If not take a look at `tree_adt.pvs` where it's defined.

```

tree_depth[t:TYPE]: THEORY
BEGIN
IMPORTING tree_adt

  depth: [tree[t] -> nat] =
    reduce_nat(0, LAMBDA(x:t, l, r:nat):max(l, r)+1)

END tree_depth

IMPORTING tree_adt

```

gives us the type former `tree[]`, e.g. `tree[nat]` is the type of `nat` labelled trees.

`leaf[t]`, `node[t]`, `leaf?[t]`, `node?[t]` etc.

can usually be abbreviated by

`leaf`, `node`, `leaf?`, `node?` etc.

With

```

IMPORTING tree_adt, tree_depth

```

we also get the `depth` function.

```
|-----  
{1}  FORALL (t: tree[t]): no_leaves(t) = no_nodes(t) + 1
```

Rule? `(induct-and-simplify "t")`

Alternatively, `induct "t", skosimp, etc.`

12.2 The option datatype

We saw that partial functions can be turned into total functions by defining them on a subset. Another possibility is to change the result type of the function so as to contain a special “error element” which—when taken on—flags that we are outwith the domain of the function. An example: the predecessor function on natural numbers can be defined on the set $\{n \mid n > 0\}$, then returning a natural number. Alternatively, we can define it on the whole of \mathbb{N} and then return a value in $\mathbb{N} \cup \{\text{none}\}$ with the understanding that $\text{pred}(0) = \text{none}$ and $\text{pred}(n) = n - 1$ otherwise.

Since this situation occurs sufficiently often, it is handy to have a new type former which tacks on a special element to any other type. Since this error element might already have been present it’s more convenient to also flag the other elements which is achieved with the `option` datatype. If T is a set so is `option(T)` and its members are `none` and `some(x)` when $x \in T$.

A property holds for all elements in `option(T)` provided it holds for `none` and for all elements of the form `some(x)`. That’s the induction principle for the type `option`. We also have the subsets `(none?)`, `(some?)` consisting of `none` and the `some(x)`, respectively.

So, everything is as before, except that this time the constructors don’t take arguments from the inductively defined set. In this case, the “induction principle” is equivalent to a first-order formula (no quantification over predicates). Do you see, which one?

```
option[t:TYPE]: DATATYPE  
BEGIN  
  none : none?  
  some(content:t) : some?  
END option
```

Again, a file `option_adt.pvs` is created; look at it and try to understand its contents.

13 Case study: n -bit ripple-carry adder

In this and the following section we look at two case studies that are part of the PVS documentation. They demonstrate many of the features that we have seen up to now and also give some idea of how to model real-world systems in PVS.

The first case study is about an n -bit ripple adder.

A bitvector of size n is a sequence of n bits. The bitvector $\mathbf{b} = (b_{n-1}, \dots, b_0)$ denotes the integer $\sum_{i=0}^{n-1} b_i 2^i$. Here and in the sequel we use the implicit conversion from booleans to numbers that maps `FALSE` to 0 and `TRUE` to 1.

An n -bit adder takes two bitvectors \mathbf{b} and \mathbf{c} as well as a single bit cin (“carry in”) and computes a bitvector \mathbf{s} —the sum—and a single bit $cout$ (“carry out”) characterised as follows:

$$\sum_{i=0}^{n-1} s_i 2^i + cout 2^n = \sum_{i=0}^{n-1} b_i 2^i + \sum_{i=0}^{n-1} c_i 2^i + cin$$

13.1 Full adder

For $n = 1$ such a device is known as a “full adder”; it can be realised by the following boolean circuitry where \oplus is *exclusive or* (“xor”).

$$\begin{aligned} s &= b \oplus c \oplus cin \\ cout &= (b \wedge c) \vee (b \wedge cin) \vee (c \wedge cin) \end{aligned}$$

Characteristic property:

$$s + 2cout = b + c + cin$$

Let us define the full adder in PVS:

```
full_adder: THEORY
BEGIN
  IMPORTING bitvectors@bit

  FA(x, y, cin): [# carry, sum: bit #] =
    (# carry := (x AND y) OR (x AND cin) OR (y AND cin),
```

```

bit: THEORY
BEGIN

  bit  : TYPE = bool

  b: VAR bit

  nbit : TYPE = below(2)  % could be {n: nat | n = 0 OR n = 1}
  b2n(b:bool)      : nbit = IF b THEN 1 ELSE 0 ENDIF
  n2b(nb: nbit)    : bool = (nb = 1)

  CONVERSION b2n
END bit

```

Figure 26: The theory of bits

```

upto(i):  NONEMPTY_TYPE = {s: nat | s <= i} CONTAINING i
below(i): TYPE = {s: nat | s < i}  % may be empty

```

Figure 27: Below and upto—simple dependent types

```

sum := (x XOR y) XOR cin #)

FA_char: LEMMA
  sum(FA(x, y, cin)) =
    x + y + cin - 2 * carry(FA(x, y, cin))
END full_adder

```

This definition contains a few new concepts. The `IMPORTING` clause

```
IMPORTING bitvectors@bit
```

imports the theory `bit` residing in the PVS context `bitvectors`, see Fig. 26. Here `below(2)` is the subtype of `nat` consisting of just 0,1. More generally, `below(n)` is the subtype consisting of $0, \dots, n-1$. This is defined in the prelude. They form an example of *dependent types*, i.e., families of types depending on values. The `CONVERSION` keyword specifies that the conversion function `b2n` can be left implicit, i.e., if a bit is used where a natural number is expected then the conversion is automatically inserted.

Records: The type

```
[#carry: bit, sum: bit #]
```

is a record type with two fields: `carry` and `sum` both of type `bit`. Elements of the record type are formed using the syntax

```
(#carry := ..., sum := ... #)
```

If `e: [#carry: bit, sum: bit#]` then `e`carry` and `e`sum` are its components. Alternative syntax is `carry(e)` and `sum(e)`. We can (functionally) update `e` using

```
e WITH [ `carry := ... ]
```

This denotes the record with `carry` field equal to `...` and the rest as in `e`.

Proof of `FA_char` The characteristic property of the full adder can be proved by distinguishing eight cases (the possible values of the three input variables) and arithmetic simplification. Therefore, `(grind)` can do it. Of course, we can manually do the case distinction by invoking `bit_cases`.

13.2 Formalisation of the ripple adder

We formulate a theory `ripple_adder` parametrised by the size `N`:

```
ripple_adder[N: posnat] : THEORY
```

```
BEGIN
```

```
  IMPORTING full_adder, bitvectors@bv[N], bitvectors@bv_nat
```

We import the theories `bv[N]` defining a type of bitvectors of size `N` and a conversion function from that type to the natural numbers.

```
bv[N: nat]: THEORY
```

```
BEGIN
```

```
  IMPORTING bit
```

```
  bvec : TYPE = [below(N) -> bit]
```

```
  ^ (bv: bvec, (i: below(N))): bit = bv(i)
```

```
END bv
```

```

bv_nat [N: nat]: THEORY
BEGIN
  IMPORTING bv, exp2
  bv2nat_rec (n: upto(N), bv:bvec[N]): RECURSIVE nat =
    IF n = 0 THEN 0
    ELSE exp2 (n-1) * bv^(n-1) + bv2nat_rec (n - 1, bv)
    ENDIF
  MEASURE n
  bv2nat (bv:bvec[N]): below(exp2(N)) = bv2nat_rec (N, bv)

```

This defines a bitvector of size N as a function from $\text{below}(N)$. If b is a bitvector and $i:\text{below}(N)$ then we can use the notation b^i to access its i -th component.

The conversion to natural numbers is defined recursively.

We continue with the theory `ripple_adder`:

```

nth_cin (j:below(N), bv1, bv2:bvec[N]): RECURSIVE bit =
  IF j = 0 THEN FALSE
  ELSE carry (FA (bv1 (j-1), bv2 (j-1),
    nth_cin (j-1, bv1, bv2)))
  ENDIF MEASURE j

adder (bv1, bv2): [# carry: bit, sum :bvec[N] #] =
  (# carry := nth_cin (N, bv1, bv2),
  sum := (LAMBDA (n:below(N)): sum (FA (bv1 (n), bv2 (n),
    nth_cin (n, bv1, bv2)))) #)

```

The function `nth_cin` computes the carry input to the j -th full adder.

Typechecking these definitions generates a few typechecking conditions (TCCs) due to the use the subtype `below` (all except `nth_cin_TCC3`) and recursive definition:

```

nth_cin_TCC1: OBLIGATION
  FORALL (j: upto[N]): NOT j = 0 IMPLIES
    j - 1 >= 0 AND j - 1 < N;
nth_cin_TCC2: OBLIGATION
  FORALL (j: upto[N]): NOT j = 0 IMPLIES
    j - 1 >= 0 AND j - 1 <= N;
nth_cin_TCC3: OBLIGATION

```

```

FORALL (j: upto[N]): NOT j = 0 IMPLIES j - 1 < j;
adder_TCC1: OBLIGATION N <= N;
adder_TCC2: OBLIGATION FORALL (n: below(N)): n <= N;

```

These are all trivial and can be proved with `M-x tcp`. We'll nevertheless prove the first one by hand:

```

|-----
{1}  FORALL (j: upto[N]): NOT j = 0 IMPLIES j - 1 >= 0 AND j - 1 < N

```

Now `(skosimp)` introduces a Skolem constant `j!1` of type `upto(n)`. We can use `M-x show-skolem-constants` to display it. With the command `(typepred "j!1")` (see prover guide p. 96) we can add the type constraint for `j!1` to our current sequent:

```

{-1}  j!1 <= N
      |-----
[1]   j!1 = 0
[2]   j!1 - 1 >= 0 AND j!1 - 1 < N

```

Implicitly, since `j!1` is a natural number we also know that `j!1 >= 0`. Therefore, we have now an arithmetic truth which we can discharge with `(grind)`.

13.3 Specification of the ripple adder

Our goal is to prove the following.

```

adder_correct: THEOREM
  FORALL (bv1, bv2:bvec[N]):
    bv2nat (sum (adder (bv1, bv2))) =
      bv2nat (bv1) + bv2nat (bv2) -
        carry (adder (bv1, bv2)) * exp2 (N)

```

We want to prove this by induction on `N`, but `N` is fixed. So, we introduce a new parameter `j` allowing us to reformulate this goal for initial segments of the adder:

```

adder_invariant: LEMMA
  FORALL (j:upto(N), bv1, bv2:bvec[N]):
    bv2nat_rec (j, sum (adder (bv1, bv2)))

```

$$\begin{aligned} &= \text{bv2nat_rec}(j, \text{bv1}) + \\ &\text{bv2nat_rec}(j, \text{bv2}) \\ &\quad - \text{nth_cin}(j, \text{bv1}, \text{bv2}) * \text{exp2}(j) \end{aligned}$$

The following command proves the lemma:

```
(induct-and-simplify "j" :theories "full_adder" :exclude "FA")
```

The `:theories` directive says that the lemma `FA_char` may be used; the `:exclude` directive stipulates that the definition of `FA` should not be expanded. Without the directive `induct-and-simplify` will eagerly expand `FA` and then not see anymore where to apply `FA_char`.

Now the main result is a direct consequence.