



The Isabelle/Isar Reference Manual

Makarius Wenzel

With Contributions by Clemens Ballarin, Stefan Berghofer,
Timothy Bourke, Lucas Dixon, Florian Haftmann,
Gerwin Klein, Alexander Krauss, Tobias Nipkow,
David von Oheimb, Larry Paulson, and Sebastian Skalberg

January 30, 2011

Contents

I	Basic Concepts	1
1	Introduction	2
1.1	Overview	2
2	The Isabelle/Isar Framework	4
2.1	The Pure framework	6
2.1.1	Primitive inferences	7
2.1.2	Reasoning with rules	8
2.2	The Isar proof language	10
2.2.1	Context elements	11
2.2.2	Structured statements	13
2.2.3	Structured proof refinement	14
2.2.4	Calculational reasoning	16
2.3	Example: First-Order Logic	17
2.3.1	Equational reasoning	17
2.3.2	Basic group theory	18
2.3.3	Propositional logic	20
2.3.4	Classical logic	22
2.3.5	Quantifiers	23
2.3.6	Canonical reasoning patterns	23
II	General Language Elements	26
3	Outer syntax	27
3.1	Lexical matters	28
3.2	Common syntax entities	30
3.2.1	Names	30
3.2.2	Numbers	31
3.2.3	Comments	31
3.2.4	Type classes, sorts and arities	31
3.2.5	Types and terms	32
3.2.6	Term patterns and declarations	34
3.2.7	Attributes and theorems	35

4	Document preparation	38
4.1	Markup commands	39
4.2	Document Antiquotations	41
4.3	Markup via command tags	46
4.4	Draft presentation	47
5	Theory specifications	49
5.1	Defining theories	49
5.2	Local theory targets	50
5.3	Basic specification elements	51
5.4	Generic declarations	53
5.5	Locales	54
5.5.1	Locale expressions	55
5.5.2	Locale declarations	56
5.5.3	Locale interpretations	59
5.6	Classes	62
5.6.1	The class target	64
5.6.2	Co-regularity of type classes and arities	65
5.7	Unrestricted overloading	65
5.8	Incorporating ML code	66
5.9	Primitive specification elements	68
5.9.1	Type classes and sorts	68
5.9.2	Types and type abbreviations	69
5.9.3	Constants and definitions	70
5.10	Axioms and theorems	72
5.11	Oracles	72
5.12	Name spaces	73
6	Proofs	75
6.1	Proof structure	75
6.1.1	Formal notepad	75
6.1.2	Blocks	76
6.1.3	Omitting proofs	76
6.2	Statements	77
6.2.1	Context elements	77
6.2.2	Term abbreviations	79
6.2.3	Facts and forward chaining	80
6.2.4	Goals	82
6.3	Refinement steps	85
6.3.1	Proof method expressions	85
6.3.2	Initial and terminal proof steps	86

6.3.3	Fundamental methods and attributes	88
6.3.4	Emulating tactic scripts	91
6.3.5	Defining proof methods	93
6.4	Generalized elimination	93
6.5	Calculational reasoning	95
6.6	Proof by cases and induction	97
6.6.1	Rule contexts	97
6.6.2	Proof methods	100
6.6.3	Declaring rules	104
7	Inner syntax — the term language	106
7.1	Printing logical entities	106
7.1.1	Diagnostic commands	106
7.1.2	Details of printed content	108
7.1.3	Printing limits	110
7.2	Mixfix annotations	110
7.3	Explicit notation	113
7.4	The Pure syntax	114
7.4.1	Priority grammars	114
7.4.2	The Pure grammar	115
7.5	Lexical matters	119
7.6	Syntax and translations	120
7.7	Syntax translation functions	121
7.8	Inspecting the syntax	122
8	Other commands	124
8.1	Inspecting the context	124
8.2	History commands	127
8.3	System commands	127
9	Generic tools and packages	129
9.1	Configuration options	129
9.2	Basic proof tools	130
9.2.1	Miscellaneous methods and attributes	130
9.2.2	Low-level equational reasoning	132
9.2.3	Further tactic emulations	133
9.3	The Simplifier	136
9.3.1	Simplification methods	136
9.3.2	Declaring rules	138
9.3.3	Simplification procedures	138
9.3.4	Forward simplification	139

9.4	The Classical Reasoner	140
9.4.1	Basic methods	140
9.4.2	Automated methods	141
9.4.3	Combined automated methods	142
9.4.4	Declaring rules	144
9.4.5	Classical operations	145
9.5	Object-logic setup	145

III Object-Logics 147

10 Isabelle/HOL 148

10.1	Typedef axiomatization	148
10.2	Adhoc tuples	149
10.3	Records	149
10.3.1	Basic concepts	150
10.3.2	Record specifications	151
10.3.3	Record operations	151
10.3.4	Derived rules and proof tools	153
10.4	Datatypes	153
10.5	Functorial structure of types	154
10.6	Recursive functions	155
10.6.1	Proof methods related to recursive definitions	157
10.6.2	Functions with explicit partiality	159
10.6.3	Old-style recursive function definitions (TFL)	160
10.7	Inductive and coinductive definitions	162
10.7.1	Derived rules	163
10.7.2	Monotonicity theorems	164
10.8	Arithmetic proof support	164
10.9	Intuitionistic proof search	165
10.10	Coherent Logic	165
10.11	Checking and refuting propositions	166
10.12	Unstructured case analysis and induction	167
10.13	Executable code	169
10.14	Definition by specification	178

11 Isabelle/HOLCF 180

11.1	Mixfix syntax for continuous operations	180
11.2	Recursive domains	180

12 Isabelle/ZF	182
12.1 Type checking	182
12.2 (Co)Inductive sets and datatypes	182
12.2.1 Set definitions	182
12.2.2 Primitive recursive functions	184
12.2.3 Cases and induction: emulating tactic scripts	185
 IV Appendix	 186
A Isabelle/Isar quick reference	187
A.1 Proof commands	187
A.1.1 Primitives and basic syntax	187
A.1.2 Abbreviations and synonyms	188
A.1.3 Derived elements	188
A.1.4 Diagnostic commands	188
A.2 Proof methods	189
A.3 Attributes	190
A.4 Rule declarations and methods	190
A.5 Emulating tactic scripts	191
A.5.1 Commands	191
A.5.2 Methods	191
 B Predefined Isabelle symbols	 192
 C ML tactic expressions	 198
C.1 Resolution tactics	198
C.2 Simplifier tactics	199
C.3 Classical Reasoner tactics	199
C.4 Miscellaneous tactics	199
C.5 Tacticals	200

Part I

Basic Concepts

Introduction

1.1 Overview

The *Isabelle* system essentially provides a generic infrastructure for building deductive systems (programmed in Standard ML), with a special focus on interactive theorem proving in higher-order logics. Many years ago, even end-users would refer to certain ML functions (goal commands, tactics, tacticals etc.) to pursue their everyday theorem proving tasks.

In contrast *Isar* provides an interpreted language environment of its own, which has been specifically tailored for the needs of theory and proof development. Compared to raw ML, the Isabelle/Isar top-level provides a more robust and comfortable development platform, with proper support for theory development graphs, managed transactions with unlimited undo etc. The Isabelle/Isar version of the *Proof General* user interface [1, 2] provides a decent front-end for interactive theory and proof development in this advanced theorem proving environment, even though it is somewhat biased towards old-style proof scripts.

Apart from the technical advances over bare-bones ML programming, the main purpose of the Isar language is to provide a conceptually different view on machine-checked proofs [36, 37]. *Isar* stands for *Intelligible semi-automated reasoning*. Drawing from both the traditions of informal mathematical proof texts and high-level programming languages, Isar offers a versatile environment for structured formal proof documents. Thus properly written Isar proofs become accessible to a broader audience than unstructured tactic scripts (which typically only provide operational information for the machine). Writing human-readable proof texts certainly requires some additional efforts by the writer to achieve a good presentation, both of formal and informal parts of the text. On the other hand, human-readable formal texts gain some value in their own right, independently of the mechanic proof-checking process.

Despite its grand design of structured proof texts, Isar is able to assimilate the old tactical style as an “improper” sub-language. This provides an easy upgrade path for existing tactic scripts, as well as some means for interactive

experimentation and debugging of structured proofs. Isabelle/Isar supports a broad range of proof styles, both readable and unreadable ones.

The generic Isabelle/Isar framework (see chapter 2) works reasonably well for any Isabelle object-logic that conforms to the natural deduction view of the Isabelle/Pure framework. Specific language elements introduced by the major object-logics are described in chapter 10 (Isabelle/HOL), chapter 11 (Isabelle/HOLCF), and chapter 12 (Isabelle/ZF). The main language elements are already provided by the Isabelle/Pure framework. Nevertheless, examples given in the generic parts will usually refer to Isabelle/HOL as well.

Isar commands may be either *proper* document constructors, or *improper commands*. Some proof methods and attributes introduced later are classified as improper as well. Improper Isar language elements, which are marked by “*” in the subsequent chapters; they are often helpful when developing proof documents, but their use is discouraged for the final human-readable outcome. Typical examples are diagnostic commands that print terms or theorems according to the current context; other commands emulate old-style tactical theorem proving.

The Isabelle/Isar Framework

Isabelle/Isar [36, 37, 17, 40, 38] is intended as a generic framework for developing formal mathematical documents with full proof checking. Definitions and proofs are organized as theories. An assembly of theory sources may be presented as a printed document; see also chapter 4.

The main objective of Isar is the design of a human-readable structured proof language, which is called the “primary proof format” in Isar terminology. Such a primary proof language is somewhere in the middle between the extremes of primitive proof objects and actual natural language. In this respect, Isar is a bit more formalistic than Mizar [33, 31, 41], using logical symbols for certain reasoning schemes where Mizar would prefer English words; see [42] for further comparisons of these systems.

So Isar challenges the traditional way of recording informal proofs in mathematical prose, as well as the common tendency to see fully formal proofs directly as objects of some logical calculus (e.g. λ -terms in a version of type theory). In fact, Isar is better understood as an interpreter of a simple block-structured language for describing the data flow of local facts and goals, interspersed with occasional invocations of proof methods. Everything is reduced to logical inferences internally, but these steps are somewhat marginal compared to the overall bookkeeping of the interpretation process. Thanks to careful design of the syntax and semantics of Isar language elements, a formal record of Isar instructions may later appear as an intelligible text to the attentive reader.

The Isar proof language has emerged from careful analysis of some inherent virtues of the existing logical framework of Isabelle/Pure [26, 27], notably composition of higher-order natural deduction rules, which is a generalization of Gentzen’s original calculus [7]. The approach of generic inference systems in Pure is continued by Isar towards actual proof texts.

Concrete applications require another intermediate layer: an object-logic. Isabelle/HOL [20] (simply-typed set-theory) is being used most of the time; Isabelle/ZF [23] is less extensively developed, although it would probably fit better for classical mathematics.

In order to illustrate natural deduction in Isar, we shall refer to the background theory and library of Isabelle/HOL. This includes common notions of predicate logic, naive set-theory etc. using fairly standard mathematical notation. From the perspective of generic natural deduction there is nothing special about the logical connectives of HOL (\wedge , \vee , \forall , \exists , etc.), only the resulting reasoning principles are relevant to the user. There are similar rules available for set-theory operators (\cap , \cup , \bigcap , \bigcup , etc.), or any other theory developed in the library (lattice theory, topology etc.).

Subsequently we briefly review fragments of Isar proof texts corresponding directly to such general deduction schemes. The examples shall refer to set-theory, to minimize the danger of understanding connectives of predicate logic as something special.

The following deduction performs \cap -introduction, working forwards from assumptions towards the conclusion. We give both the Isar text, and depict the primitive rule involved, as determined by unification of the problem against rules that are declared in the library context.

assume $x \in A$ and $x \in B$ then have $x \in A \cap B$..	$\frac{x \in A \quad x \in B}{x \in A \cap B}$
--	--

Note that **assume** augments the proof context, **then** indicates that the current fact shall be used in the next step, and **have** states an intermediate goal. The two dots “..” refer to a complete proof of this claim, using the indicated facts and a canonical rule from the context. We could have been more explicit here by spelling out the final proof step via the **by** command:

assume $x \in A$ **and** $x \in B$
then have $x \in A \cap B$ **by** (*rule IntI*)

The format of the \cap -introduction rule represents the most basic inference, which proceeds from given premises to a conclusion, without any nested proof context involved.

The next example performs backwards introduction on $\bigcap \mathcal{A}$, the intersection of all sets within a given set. This requires a nested proof of set membership within a local context, where A is an arbitrary-but-fixed member of the collection:

have $x \in \bigcap \mathcal{A}$ proof fix A assume $A \in \mathcal{A}$ show $x \in A$ <i><proof></i> qed	$\frac{\begin{array}{c} [A][A \in \mathcal{A}] \\ \vdots \\ x \in A \end{array}}{x \in \bigcap \mathcal{A}}$
--	--

This Isar reasoning pattern again refers to the primitive rule depicted above. The system determines it in the “**proof**” step, which could have been spelt out more explicitly as “**proof** (*rule InterI*)”. Note that the rule involves both a local parameter A and an assumption $A \in \mathcal{A}$ in the nested reasoning. This kind of compound rule typically demands a genuine sub-proof in Isar, working backwards rather than forwards as seen before. In the proof body we encounter the **fix-assume-show** outline of nested sub-proofs that is typical for Isar. The final **show** is like **have** followed by an additional refinement of the enclosing claim, using the rule derived from the proof body.

The next example involves $\bigcup \mathcal{A}$, which can be characterized as the set of all x such that $\exists A. x \in A \wedge A \in \mathcal{A}$. The elimination rule for $x \in \bigcup \mathcal{A}$ does not mention \exists and \wedge at all, but admits to obtain directly a local A such that $x \in A$ and $A \in \mathcal{A}$ hold. This corresponds to the following Isar proof and inference rule, respectively:

<pre> assume $x \in \bigcup \mathcal{A}$ then have C proof fix A assume $x \in A$ and $A \in \mathcal{A}$ show C <i><proof></i> qed </pre>	$ \frac{x \in \bigcup \mathcal{A} \quad \begin{array}{c} [A][x \in A, A \in \mathcal{A}] \\ \vdots \\ C \end{array}}{C} $
--	---

Although the Isar proof follows the natural deduction rule closely, the text reads not as natural as anticipated. There is a double occurrence of an arbitrary conclusion C , which represents the final result, but is irrelevant for now. This issue arises for any elimination rule involving local parameters. Isar provides the derived language element **obtain**, which is able to perform the same elimination proof more conveniently:

```

assume  $x \in \bigcup \mathcal{A}$ 
then obtain  $A$  where  $x \in A$  and  $A \in \mathcal{A}$  ..

```

Here we avoid to mention the final conclusion C and return to plain forward reasoning. The rule involved in the “**..**” proof is the same as before.

2.1 The Pure framework

The Pure logic [26, 27] is an intuitionistic fragment of higher-order logic [6]. In type-theoretic parlance, there are three levels of λ -calculus with corresponding arrows $\Rightarrow/\wedge/\Longrightarrow$:

$\alpha \Rightarrow \beta$	syntactic function space (terms depending on terms)
$\bigwedge x. B(x)$	universal quantification (proofs depending on terms)
$A \Longrightarrow B$	implication (proofs depending on proofs)

Here only the types of syntactic terms, and the propositions of proof terms have been shown. The λ -structure of proofs can be recorded as an optional feature of the Pure inference kernel [4], but the formal system can never depend on them due to *proof irrelevance*.

On top of this most primitive layer of proofs, Pure implements a generic calculus for nested natural deduction rules, similar to [32]. Here object-logic inferences are internalized as formulae over \bigwedge and \Longrightarrow . Combining such rule statements may involve higher-order unification [25].

2.1.1 Primitive inferences

Term syntax provides explicit notation for abstraction $\lambda x :: \alpha. b(x)$ and application $b\ a$, while types are usually implicit thanks to type-inference; terms of type *prop* are called propositions. Logical statements are composed via $\bigwedge x :: \alpha. B(x)$ and $A \Longrightarrow B$. Primitive reasoning operates on judgments of the form $\Gamma \vdash \varphi$, with standard introduction and elimination rules for \bigwedge and \Longrightarrow that refer to fixed parameters x_1, \dots, x_m and hypotheses A_1, \dots, A_n from the context Γ ; the corresponding proof terms are left implicit. The subsequent inference rules define $\Gamma \vdash \varphi$ inductively, relative to a collection of axioms:

$$\begin{array}{c}
\frac{(A \text{ axiom})}{\vdash A} \quad \overline{A \vdash A} \\[10pt]
\frac{\Gamma \vdash B(x) \quad x \notin \Gamma}{\Gamma \vdash \bigwedge x. B(x)} \quad \frac{\Gamma \vdash \bigwedge x. B(x)}{\Gamma \vdash B(a)} \\[10pt]
\frac{\Gamma \vdash B}{\Gamma - A \vdash A \Longrightarrow B} \quad \frac{\Gamma_1 \vdash A \Longrightarrow B \quad \Gamma_2 \vdash A}{\Gamma_1 \cup \Gamma_2 \vdash B}
\end{array}$$

Furthermore, Pure provides a built-in equality $\equiv :: \alpha \Rightarrow \alpha \Rightarrow \text{prop}$ with axioms for reflexivity, substitution, extensionality, and $\alpha\beta\eta$ -conversion on λ -terms.

An object-logic introduces another layer on top of Pure, e.g. with types *i* for individuals and *o* for propositions, term constants $\text{Trueprop} :: o \Rightarrow \text{prop}$ as (implicit) derivability judgment and connectives like $\wedge :: o \Rightarrow o \Rightarrow o$ or $\vee :: (i \Rightarrow o) \Rightarrow o$, and axioms for object-level rules such as *conjI*: $A \Longrightarrow B \Longrightarrow A \wedge B$ or *allI*: $(\bigwedge x. B\ x) \Longrightarrow \forall x. B\ x$. Derived object rules are

represented as theorems of Pure. After the initial object-logic setup, further axiomatizations are usually avoided; plain definitions and derived principles are used exclusively.

2.1.2 Reasoning with rules

Primitive inferences mostly serve foundational purposes. The main reasoning mechanisms of Pure operate on nested natural deduction rules expressed as formulae, using \bigwedge to bind local parameters and \implies to express entailment. Multiple parameters and premises are represented by repeating these connectives in a right-associative manner.

Since \bigwedge and \implies commute thanks to the theorem $(A \implies (\bigwedge x. B\ x)) \equiv (\bigwedge x. A \implies B\ x)$, we may assume w.l.o.g. that rule statements always observe the normal form where quantifiers are pulled in front of implications at each level of nesting. This means that any Pure proposition may be presented as a *Hereditary Harrop Formula* [13] which is of the form $\bigwedge x_1 \dots x_m. H_1 \implies \dots H_n \implies A$ for $m, n \geq 0$, and A atomic, and H_1, \dots, H_n being recursively of the same format. Following the convention that outermost quantifiers are implicit, Horn clauses $A_1 \implies \dots A_n \implies A$ are a special case of this.

For example, \cap -introduction rule encountered before is represented as a Pure theorem as follows:

$$\text{IntI: } x \in A \implies x \in B \implies x \in A \cap B$$

This is a plain Horn clause, since no further nesting on the left is involved. The general \bigcap -introduction corresponds to a Hereditary Harrop Formula with one additional level of nesting:

$$\text{InterI: } (\bigwedge A. A \in \mathcal{A} \implies x \in A) \implies x \in \bigcap \mathcal{A}$$

Goals are also represented as rules: $A_1 \implies \dots A_n \implies C$ states that the sub-goals A_1, \dots, A_n entail the result C ; for $n = 0$ the goal is finished. To allow C being a rule statement itself, we introduce the protective marker $\#$:: $prop \Rightarrow prop$, which is defined as identity and hidden from the user. We initialize and finish goal states as follows:

$$\overline{C \implies \#C} \text{ (init)} \quad \frac{\#C}{C} \text{ (finish)}$$

Goal states are refined in intermediate proof steps until a finished form is achieved. Here the two main reasoning principles are *resolution*, for back-chaining a rule against a sub-goal (replacing it by zero or more sub-goals),

and *assumption*, for solving a sub-goal (finding a short-circuit with local assumptions). Below \bar{x} stands for x_1, \dots, x_n ($n \geq 0$).

$$\frac{\begin{array}{l} \text{rule: } \bar{A} \bar{a} \Longrightarrow B \bar{a} \\ \text{goal: } (\bigwedge \bar{x}. \bar{H} \bar{x} \Longrightarrow B' \bar{x}) \Longrightarrow C \\ \text{goal unifier: } (\lambda \bar{x}. B (\bar{a} \bar{x})) \theta = B' \theta \end{array}}{(\bigwedge \bar{x}. \bar{H} \bar{x} \Longrightarrow \bar{A} (\bar{a} \bar{x})) \theta \Longrightarrow C \theta} \text{ (resolution)}$$

$$\frac{\begin{array}{l} \text{goal: } (\bigwedge \bar{x}. \bar{H} \bar{x} \Longrightarrow A \bar{x}) \Longrightarrow C \\ \text{assm unifier: } A \theta = H_i \theta \text{ (for some } H_i) \end{array}}{C \theta} \text{ (assumption)}$$

The following trace illustrates goal-oriented reasoning in Isabelle/Pure:

$$\begin{array}{l} (A \wedge B \Longrightarrow B \wedge A) \Longrightarrow \#(A \wedge B \Longrightarrow B \wedge A) \text{ (init)} \\ (A \wedge B \Longrightarrow B) \Longrightarrow (A \wedge B \Longrightarrow A) \Longrightarrow \# \dots \text{ (resolution } B \Longrightarrow A \Longrightarrow B \wedge A) \\ (A \wedge B \Longrightarrow A \wedge B) \Longrightarrow (A \wedge B \Longrightarrow A) \Longrightarrow \# \dots \text{ (resolution } A \wedge B \Longrightarrow B) \\ \quad (A \wedge B \Longrightarrow A) \Longrightarrow \# \dots \text{ (assumption)} \\ \quad (A \wedge B \Longrightarrow B \wedge A) \Longrightarrow \# \dots \text{ (resolution } A \wedge B \Longrightarrow A) \\ \quad \quad \# \dots \text{ (assumption)} \\ \quad A \wedge B \Longrightarrow B \wedge A \text{ (finish)} \end{array}$$

Compositions of *assumption* after *resolution* occurs quite often, typically in elimination steps. Traditional Isabelle tactics accommodate this by a combined *elim_resolution* principle. In contrast, Isar uses a slightly more refined combination, where the assumptions to be closed are marked explicitly, using again the protective marker $\#$:

$$\frac{\begin{array}{l} \text{sub-proof: } \bar{G} \bar{a} \Longrightarrow B \bar{a} \\ \text{goal: } (\bigwedge \bar{x}. \bar{H} \bar{x} \Longrightarrow B' \bar{x}) \Longrightarrow C \\ \text{goal unifier: } (\lambda \bar{x}. B (\bar{a} \bar{x})) \theta = B' \theta \\ \text{assm unifiers: } (\lambda \bar{x}. G_j (\bar{a} \bar{x})) \theta = \#H_i \theta \\ \quad \text{(for each marked } G_j \text{ some } \#H_i) \end{array}}{(\bigwedge \bar{x}. \bar{H} \bar{x} \Longrightarrow \bar{G}' (\bar{a} \bar{x})) \theta \Longrightarrow C \theta} \text{ (refinement)}$$

Here the *sub-proof* rule stems from the main **fix-assume-show** outline of Isar (cf. §2.2.3): each assumption indicated in the text results in a marked premise G above. The marking enforces resolution against one of the sub-goal's premises. Consequently, **fix-assume-show** enables to fit the result of a sub-proof quite robustly into a pending sub-goal, while maintaining a good measure of flexibility.

2.2 The Isar proof language

Structured proofs are presented as high-level expressions for composing entities of Pure (propositions, facts, and goals). The Isar proof language allows to organize reasoning within the underlying rule calculus of Pure, but Isar is not another logical calculus!

Isar is an exercise in sound minimalism. Approximately half of the language is introduced as primitive, the rest defined as derived concepts. The following grammar describes the core language (category *proof*), which is embedded into theory specification elements such as **theorem**; see also §2.2.2 for the separate category *statement*.

```

theory-stmt  =  theorem statement proof | definition ... | ...
proof        =  prfx* proof method? stmt* qed method?
prfx         =  using facts
              |  unfolding facts
stmt         =  { stmt* }
              |  next
              |  note name = facts
              |  let term = term
              |  fix var+
              |  assume <<inference>> name: props
              |  then? goal
goal         =  have name: props proof
              |  show name: props proof

```

Simultaneous propositions or facts may be separated by the **and** keyword.

The syntax for terms and propositions is inherited from Pure (and the object-logic). A *pattern* is a *term* with schematic variables, to be bound by higher-order matching.

Facts may be referenced by name or proposition. For example, the result of “**have** *a*: *A* ⟨*proof*⟩” becomes available both as *a* and ⟨*A*⟩. Moreover, fact expressions may involve attributes that modify either the theorem or the background context. For example, the expression “*a* [*OF* *b*]” refers to the composition of two facts according to the *resolution* inference of §2.1.2, while “*a* [*intro*]” declares a fact as introduction rule in the context.

The special fact called “*this*” always refers to the last result, as produced by **note**, **assume**, **have**, or **show**. Since **note** occurs frequently together with **then** we provide some abbreviations:

```

from a  ≡  note a then
with a  ≡  from a and this

```

The *method* category is essentially a parameter and may be populated later. Methods use the facts indicated by **then** or **using**, and then operate on the goal state. Some basic methods are predefined: “*—*” leaves the goal unchanged, “*this*” applies the facts as rules to the goal, “*rule*” applies the facts to another rule and the result to the goal (both “*this*” and “*rule*” refer to *resolution* of §2.1.2). The secondary arguments to “*rule*” may be specified explicitly as in “(*rule a*)”, or picked from the context. In the latter case, the system first tries rules declared as *elim* or *dest*, followed by those declared as *intro*.

The default method for **proof** is “*rule*” (arguments picked from the context), for **qed** it is “*—*”. Further abbreviations for terminal proof steps are “**by** *method*₁ *method*₂” for “**proof** *method*₁ **qed** *method*₂”, and “*..*” for “**by** *rule*”, and “*.*” for “**by** *this*”. The **unfolding** element operates directly on the current facts and goal by applying equalities.

Block structure can be indicated explicitly by “{ ... }”, although the body of a sub-proof already involves implicit nesting. In any case, **next** jumps into the next section of a block, i.e. it acts like closing an implicit block scope and opening another one; there is no direct correspondence to subgoals here.

The remaining elements **fix** and **assume** build up a local context (see §2.2.1), while **show** refines a pending sub-goal by the rule resulting from a nested sub-proof (see §2.2.3). Further derived concepts will support calculational reasoning (see §2.2.4).

2.2.1 Context elements

In judgments $\Gamma \vdash \varphi$ of the primitive framework, Γ essentially acts like a proof context. Isar elaborates this idea towards a higher-level notion, with additional information for type-inference, term abbreviations, local facts, hypotheses etc.

The element **fix** $x :: \alpha$ declares a local parameter, i.e. an arbitrary-but-fixed entity of a given type; in results exported from the context, x may become anything. The **assume** $\langle\langle\textit{inference}\rangle\rangle$ element provides a general interface to hypotheses: “**assume** $\langle\langle\textit{inference}\rangle\rangle A$ ” produces $A \vdash A$ locally, while the included inference tells how to discharge A from results $A \vdash B$ later on. There is no user-syntax for $\langle\langle\textit{inference}\rangle\rangle$, i.e. it may only occur internally when derived commands are defined in ML.

At the user-level, the default inference for **assume** is *discharge* as given below. The additional variants **presume** and **def** are defined as follows:

presume $A \equiv$ **assume** $\langle\text{weak-discharge}\rangle A$
def $x \equiv a \equiv$ **fix** x **assume** $\langle\text{expansion}\rangle x \equiv a$

$$\frac{\Gamma \vdash B}{\Gamma - A \vdash \#A \Longrightarrow B} \text{ (discharge)}$$

$$\frac{\Gamma \vdash B}{\Gamma - A \vdash A \Longrightarrow B} \text{ (weak-discharge)}$$

$$\frac{\Gamma \vdash B \ x}{\Gamma - (x \equiv a) \vdash B \ a} \text{ (expansion)}$$

Note that *discharge* and *weak-discharge* differ in the marker for A , which is relevant when the result of a **fix-assume-show** outline is composed with a pending goal, cf. §2.2.3.

The most interesting derived context element in Isar is **obtain** [37, §5.3], which supports generalized elimination steps in a purely forward manner. The **obtain** command takes a specification of parameters \bar{x} and assumptions \bar{A} to be added to the context, together with a proof of a case rule stating that this extension is conservative (i.e. may be removed from closed results later on):

```

<facts> obtain  $\bar{x}$  where  $\bar{A} \ \bar{x}$  <proof>  $\equiv$ 
  have case:  $\bigwedge thesis. (\bigwedge \bar{x}. \bar{A} \ \bar{x} \Longrightarrow thesis) \Longrightarrow thesis$ 
proof –
  fix thesis
  assume [intro]:  $\bigwedge \bar{x}. \bar{A} \ \bar{x} \Longrightarrow thesis$ 
  show thesis using <facts> <proof>
qed
fix  $\bar{x}$  assume  $\langle\text{elimination case}\rangle \bar{A} \ \bar{x}$ 

  case:  $\Gamma \vdash \bigwedge thesis. (\bigwedge \bar{x}. \bar{A} \ \bar{x} \Longrightarrow thesis) \Longrightarrow thesis$ 
  result:  $\Gamma \cup \bar{A} \ \bar{y} \vdash B$ 
  -----
   $\Gamma \vdash B$  (elimination)

```

Here the name “*thesis*” is a specific convention for an arbitrary-but-fixed proposition; in the primitive natural deduction rules shown before we have occasionally used C . The whole statement of “**obtain** x **where** $A \ x$ ” may be read as a claim that $A \ x$ may be assumed for some arbitrary-but-fixed x . Also note that “**obtain** A **and** B ” without parameters is similar to “**have** A **and** B ”, but the latter involves multiple sub-goals.

The subsequent Isar proof texts explain all context elements introduced above using the formal proof language itself. After finishing a local proof within a block, we indicate the exported result via **note**.

<pre> { fix x have B x <proof> } note <$\bigwedge x. B x$> </pre>	<pre> { assume A have B <proof> } note <$A \implies B$> </pre>
<pre> { def x \equiv a have B x <proof> } note <$B a$> </pre>	<pre> { obtain x where A x <proof> have B <proof> } note </pre>

This illustrates the meaning of Isar context elements without goals getting in between.

2.2.2 Structured statements

The category *statement* of top-level theorem specifications is defined as follows:

<i>statement</i>	\equiv	<i>name: props and ...</i> <i>context* conclusion</i>
<i>context</i>	\equiv	fixes <i>vars</i> and ... assumes <i>name: props and ...</i>
<i>conclusion</i>	\equiv	shows <i>name: props and ...</i> obtains <i>vars and ... where name: props and ...</i> ...

A simple *statement* consists of named propositions. The full form admits local context elements followed by the actual conclusions, such as “**fixes** *x* **assumes** *A x* **shows** *B x*”. The final result emerges as a Pure rule after discharging the context: $\bigwedge x. A x \implies B x$.

The **obtains** variant is another abbreviation defined below; unlike **obtain** (cf. §2.2.1) there may be several “cases” separated by “|”, each consisting of several parameters (*vars*) and several premises (*props*). This specifies multi-branch elimination rules.

```

obtains  $\bar{x}$  where  $\bar{A} \bar{x}$  | ...  $\equiv$ 
  fixes thesis
  assumes [intro]:  $\bigwedge \bar{x}. \bar{A} \bar{x} \implies \textit{thesis}$  and ...
  shows thesis

```

Presenting structured statements in such an “open” format usually simplifies the subsequent proof, because the outer structure of the problem is already laid out directly. E.g. consider the following canonical patterns for **shows** and **obtains**, respectively:

theorem

fixes x **and** y

assumes $A\ x$ **and** $B\ y$

shows $C\ x\ y$

proof –

from $\langle A\ x \rangle$ **and** $\langle B\ y \rangle$

show $C\ x\ y$ $\langle proof \rangle$

qed

theorem

obtains x **and** y

where $A\ x$ **and** $B\ y$

proof –

have $A\ a$ **and** $B\ b$ $\langle proof \rangle$

then show *thesis* ..

qed

Here local facts $\langle A\ x \rangle$ and $\langle B\ y \rangle$ are referenced immediately; there is no need to decompose the logical rule structure again. In the second proof the final “**then show** *thesis* ..” involves the local rule case $\bigwedge x\ y. A\ x \implies B\ y \implies thesis$ for the particular instance of terms a and b produced in the body.

2.2.3 Structured proof refinement

By breaking up the grammar for the Isar proof language, we may understand a proof text as a linear sequence of individual proof commands. These are interpreted as transitions of the Isar virtual machine (Isar/VM), which operates on a block-structured configuration in single steps. This allows users to write proof texts in an incremental manner, and inspect intermediate configurations for debugging.

The basic idea is analogous to evaluating algebraic expressions on a stack machine: $(a + b) \cdot c$ then corresponds to a sequence of single transitions for each symbol $(, a, +, b,), \cdot, c$. In Isar the algebraic values are facts or goals, and the operations are inferences.

The Isar/VM state maintains a stack of nodes, each node contains the local proof context, the linguistic mode, and a pending goal (optional). The mode determines the type of transition that may be performed next, it essentially alternates between forward and backward reasoning, with an intermediate stage for chained facts (see figure 2.1).

For example, in *state* mode Isar acts like a mathematical scratch-pad, accepting declarations like **fix**, **assume**, and claims like **have**, **show**. A goal statement changes the mode to *prove*, which means that we may now refine the problem via **unfolding** or **proof**. Then we are again in *state* mode of a proof body, which may issue **show** statements to solve pending sub-goals. A concluding **qed** will return to the original *state* mode one level upwards.

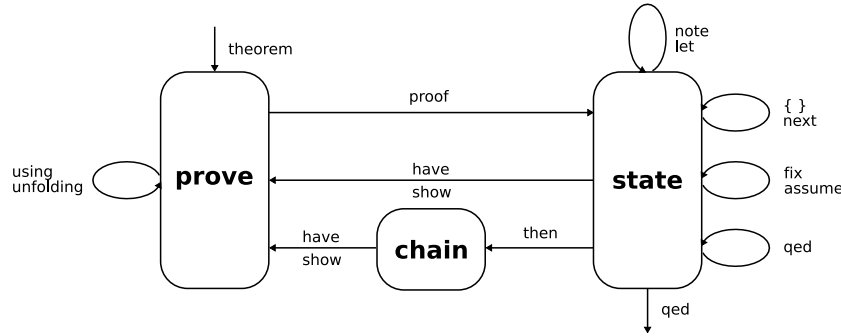


Figure 2.1: Isar/VM modes

The subsequent Isar/VM trace indicates block structure, linguistic mode, goal state, and inferences:

```

have  $A \longrightarrow B$  begin prove  $(A \longrightarrow B) \Longrightarrow \#(A \longrightarrow B)$  (init)
proof state  $(A \Longrightarrow B) \Longrightarrow \#(A \longrightarrow B)$  (resolution impI)
  assume  $A$  state
  show  $B$  begin prove
     $\langle \text{proof} \rangle$  end state  $\#(A \longrightarrow B)$  (refinement  $\#A \Longrightarrow B$ )
  qed end state  $A \longrightarrow B$  (finish)

```

Here the *refinement* inference from §2.1.2 mediates composition of Isar sub-proofs nicely. Observe that this principle incorporates some degree of freedom in proof composition. In particular, the proof body allows parameters and assumptions to be re-ordered, or commuted according to Hereditary Harrop Form. Moreover, context elements that are not used in a sub-proof may be omitted altogether. For example:

<pre> have $\bigwedge x y. A x \Longrightarrow B y \Longrightarrow C x y$ proof – fix x and y assume $A x$ and $B y$ show $C x y$ $\langle \text{proof} \rangle$ qed </pre>	<pre> have $\bigwedge x y. A x \Longrightarrow B y \Longrightarrow C x y$ proof – fix x assume $A x$ fix y assume $B y$ show $C x y$ $\langle \text{proof} \rangle$ qed </pre>
<pre> have $\bigwedge x y. A x \Longrightarrow B y \Longrightarrow C x y$ proof – fix y assume $B y$ fix x assume $A x$ show $C x y$ sorry qed </pre>	<pre> have $\bigwedge x y. A x \Longrightarrow B y \Longrightarrow C x y$ proof – fix y assume $B y$ fix x show $C x y$ sorry qed </pre>

Such “peephole optimizations” of Isar texts are practically important to improve readability, by rearranging contexts elements according to the natural

flow of reasoning in the body, while still observing the overall scoping rules.

This illustrates the basic idea of structured proof processing in Isar. The main mechanisms are based on natural deduction rule composition within the Pure framework. In particular, there are no direct operations on goal states within the proof body. Moreover, there is no hidden automated reasoning involved, just plain unification.

2.2.4 Calculational reasoning

The existing Isar infrastructure is sufficiently flexible to support calculational reasoning (chains of transitivity steps) as derived concept. The generic proof elements introduced below depend on rules declared as *trans* in the context. It is left to the object-logic to provide a suitable rule collection for mixed relations of $=$, $<$, \leq , \subset , \subseteq etc. Due to the flexibility of rule composition (§2.1.2), substitution of equals by equals is covered as well, even substitution of inequalities involving monotonicity conditions; see also [37, §6] and [3].

The generic calculational mechanism is based on the observation that rules such as *trans*: $x = y \implies y = z \implies x = z$ proceed from the premises towards the conclusion in a deterministic fashion. Thus we may reason in forward mode, feeding intermediate results into rules selected from the context. The course of reasoning is organized by maintaining a secondary fact called “*calculation*”, apart from the primary “*this*” already provided by the Isar primitives. In the definitions below, *OF* refers to *resolution* (§2.1.2) with multiple rule arguments, and *trans* represents to a suitable rule from the context:

```

also0  ≡ note calculation = this
alson+1 ≡ note calculation = trans [OF calculation this]
finally ≡ also from calculation

```

The start of a calculation is determined implicitly in the text: here **also** sets *calculation* to the current result; any subsequent occurrence will update *calculation* by combination with the next result and a transitivity rule. The calculational sequence is concluded via **finally**, where the final result is exposed for use in a concluding claim.

Here is a canonical proof pattern, using **have** to establish the intermediate results:

```

have a = b sorry
also have ... = c sorry
also have ... = d sorry
finally have a = d .

```

The term “...” above is a special abbreviation provided by the Isabelle/Isar syntax layer: it statically refers to the right-hand side argument of the previous statement given in the text. Thus it happens to coincide with relevant sub-expressions in the calculational chain, but the exact correspondence is dependent on the transitivity rules being involved.

Symmetry rules such as $x = y \implies y = x$ are like transivities with only one premise. Isar maintains a separate rule collection declared via the *sym* attribute, to be used in fact expressions “*a* [*symmetric*]”, or single-step proofs “**assume** $x = y$ **then have** $y = x$..”.

2.3 Example: First-Order Logic

```
theory First_Order_Logic
imports Pure
begin
```

In order to commence a new object-logic within Isabelle/Pure we introduce abstract syntactic categories *i* for individuals and *o* for object-propositions. The latter is embedded into the language of Pure propositions by means of a separate judgment.

```
typedecl i
typedecl o
```

```
judgment
```

```
Trueprop :: o  $\Rightarrow$  prop    (- 5)
```

Note that the object-logic judgement is implicit in the syntax: writing *A* produces *Trueprop A* internally. From the Pure perspective this means “*A* is derivable in the object-logic”.

2.3.1 Equational reasoning

Equality is axiomatized as a binary predicate on individuals, with reflexivity as introduction, and substitution as elimination principle. Note that the latter is particularly convenient in a framework like Isabelle, because syntactic congruences are implicitly produced by unification of $B\ x$ against expressions containing occurrences of x .

```
axiomatization
```

```
equal :: i  $\Rightarrow$  i  $\Rightarrow$  o    (infix = 50)
```

```
where
```

```
refl [intro]:  $x = x$  and
```


subst [elim]: $x = y \implies B\ x \implies B\ y$

Substitution is very powerful, but also hard to control in full generality. We derive some common symmetry / transitivity schemes of as particular consequences.

```

theorem sym [sym]:
  assumes  $x = y$ 
  shows  $y = x$ 
proof –
  have  $x = x$  ..
  with  $\langle x = y \rangle$  show  $y = x$  ..
qed

```

```

theorem forw_subst [trans]:
  assumes  $y = x$  and  $B\ x$ 
  shows  $B\ y$ 
proof –
  from  $\langle y = x \rangle$  have  $x = y$  ..
  from this and  $\langle B\ x \rangle$  show  $B\ y$  ..
qed

```

```

theorem back_subst [trans]:
  assumes  $B\ x$  and  $x = y$ 
  shows  $B\ y$ 
proof –
  from  $\langle x = y \rangle$  and  $\langle B\ x \rangle$ 
  show  $B\ y$  ..
qed

```

```

theorem trans [trans]:
  assumes  $x = y$  and  $y = z$ 
  shows  $x = z$ 
proof –
  from  $\langle y = z \rangle$  and  $\langle x = y \rangle$ 
  show  $x = z$  ..
qed

```

2.3.2 Basic group theory

As an example for equational reasoning we consider some bits of group theory. The subsequent locale definition postulates group operations and axioms; we also derive some consequences of this specification.

locale *group* =

```

fixes prod ::  $i \Rightarrow i \Rightarrow i$  (infix  $\circ$  70)
and inv ::  $i \Rightarrow i$  ( $(-^1)$  [1000] 999)
and unit ::  $i$  (1)
assumes assoc:  $(x \circ y) \circ z = x \circ (y \circ z)$ 
and left_unit:  $1 \circ x = x$ 
and left_inv:  $x^{-1} \circ x = 1$ 
begin

theorem right_inv:  $x \circ x^{-1} = 1$ 
proof -
  have  $x \circ x^{-1} = 1 \circ (x \circ x^{-1})$  by (rule left_unit [symmetric])
  also have  $\dots = (1 \circ x) \circ x^{-1}$  by (rule assoc [symmetric])
  also have  $1 = (x^{-1})^{-1} \circ x^{-1}$  by (rule left_inv [symmetric])
  also have  $\dots \circ x = (x^{-1})^{-1} \circ (x^{-1} \circ x)$  by (rule assoc)
  also have  $x^{-1} \circ x = 1$  by (rule left_inv)
  also have  $((x^{-1})^{-1} \circ \dots) \circ x^{-1} = (x^{-1})^{-1} \circ (1 \circ x^{-1})$  by (rule assoc)
  also have  $1 \circ x^{-1} = x^{-1}$  by (rule left_unit)
  also have  $(x^{-1})^{-1} \circ \dots = 1$  by (rule left_inv)
  finally show  $x \circ x^{-1} = 1$  .
qed

theorem right_unit:  $x \circ 1 = x$ 
proof -
  have  $1 = x^{-1} \circ x$  by (rule left_inv [symmetric])
  also have  $x \circ \dots = (x \circ x^{-1}) \circ x$  by (rule assoc [symmetric])
  also have  $x \circ x^{-1} = 1$  by (rule right_inv)
  also have  $\dots \circ x = x$  by (rule left_unit)
  finally show  $x \circ 1 = x$  .
qed

```

Reasoning from basic axioms is often tedious. Our proofs work by producing various instances of the given rules (potentially the symmetric form) using the pattern “**have** eq **by** (*rule* r)” and composing the chain of results via **also/finally**. These steps may involve any of the transitivity rules declared in §2.3.1, namely *trans* in combining the first two results in *right_inv* and in the final steps of both proofs, *forw_subst* in the first combination of *right_unit*, and *back_subst* in all other calculational steps.

Occasional substitutions in calculations are adequate, but should not be over-emphasized. The other extreme is to compose a chain by plain transitivity only, with replacements occurring always in topmost position. For example:

```

have  $x \circ 1 = x \circ (x^{-1} \circ x)$  unfolding left_inv ..
also have  $\dots = (x \circ x^{-1}) \circ x$  unfolding assoc ..

```

also have $\dots = 1 \circ x$ **unfolding** *right_inv* ..
 also have $\dots = x$ **unfolding** *left_unit* ..
 finally have $x \circ 1 = x$.

Here we have re-used the built-in mechanism for unfolding definitions in order to normalize each equational problem. A more realistic object-logic would include proper setup for the Simplifier (§9.3), the main automated tool for equational reasoning in Isabelle. Then “**unfolding** *left_inv* ..” would become “**by** (*simp only: left_inv*)” etc.

end

2.3.3 Propositional logic

We axiomatize basic connectives of propositional logic: implication, disjunction, and conjunction. The associated rules are modeled after Gentzen’s system of Natural Deduction [7].

axiomatization

imp :: $o \Rightarrow o \Rightarrow o$ (**infixr** \longrightarrow 25) **where**
impI [*intro*]: $(A \Longrightarrow B) \Longrightarrow A \longrightarrow B$ **and**
impD [*dest*]: $(A \longrightarrow B) \Longrightarrow A \Longrightarrow B$

axiomatization

disj :: $o \Rightarrow o \Rightarrow o$ (**infixr** \vee 30) **where**
disjI₁ [*intro*]: $A \Longrightarrow A \vee B$ **and**
disjI₂ [*intro*]: $B \Longrightarrow A \vee B$ **and**
disjE [*elim*]: $A \vee B \Longrightarrow (A \Longrightarrow C) \Longrightarrow (B \Longrightarrow C) \Longrightarrow C$

axiomatization

conj :: $o \Rightarrow o \Rightarrow o$ (**infixr** \wedge 35) **where**
conjI [*intro*]: $A \Longrightarrow B \Longrightarrow A \wedge B$ **and**
conjD₁: $A \wedge B \Longrightarrow A$ **and**
conjD₂: $A \wedge B \Longrightarrow B$

The conjunctive destructions have the disadvantage that decomposing $A \wedge B$ involves an immediate decision which component should be projected. The more convenient simultaneous elimination $A \wedge B \Longrightarrow (A \Longrightarrow B \Longrightarrow C) \Longrightarrow C$ can be derived as follows:

theorem *conjE* [*elim*]:

assumes $A \wedge B$
 obtains A and B

proof

from $\langle A \wedge B \rangle$ **show** A **by** (*rule conjD₁*)
from $\langle A \wedge B \rangle$ **show** B **by** (*rule conjD₂*)

qed

Here is an example of swapping conjuncts with a single intermediate elimination step:

```

assume  $A \wedge B$ 
then obtain  $B$  and  $A$  ..
then have  $B \wedge A$  ..

```

Note that the analogous elimination rule for disjunction “**assumes** $A \vee B$ **obtains** $A \mid B$ ” coincides with the original axiomatization of *disjE*.

We continue propositional logic by introducing absurdity with its characteristic elimination. Plain truth may then be defined as a proposition that is trivially true.

axiomatization

```

 $false :: o \ (\perp)$  where
 $falseE \ [elim]: \perp \Longrightarrow A$ 

```

definition

```

 $true :: o \ (\top)$  where
 $\top \equiv \perp \longrightarrow \perp$ 

```

```

theorem  $trueI \ [intro]: \top$ 
unfolding  $true\_def$  ..

```

Now negation represents an implication towards absurdity:

definition

```

 $not :: o \Rightarrow o \ (\neg \_ [40] \ 40)$  where
 $\neg A \equiv A \longrightarrow \perp$ 

```

theorem $notI \ [intro]:$

```

assumes  $A \Longrightarrow \perp$ 

```

```

shows  $\neg A$ 

```

unfolding not_def

proof

```

assume  $A$ 

```

```

then show  $\perp$  by ( $rule \ \langle A \Longrightarrow \perp \rangle$ )

```

qed

theorem $notE \ [elim]:$

```

assumes  $\neg A$  and  $A$ 

```

```

shows  $B$ 

```

proof –

```

from  $\langle \neg A \rangle$  have  $A \longrightarrow \perp$  unfolding  $not\_def$  .

```

```

from  $\langle A \longrightarrow \bot \rangle$  and  $\langle A \rangle$  have  $\bot$  ..
then show  $B$  ..
qed

```

2.3.4 Classical logic

Subsequently we state the principle of classical contradiction as a local assumption. Thus we refrain from forcing the object-logic into the classical perspective. Within that context, we may derive well-known consequences of the classical principle.

```

locale classical =
  assumes classical:  $(\neg C \Longrightarrow C) \Longrightarrow C$ 
begin

```

```

theorem double_negation:

```

```

  assumes  $\neg \neg C$ 
  shows  $C$ 
proof (rule classical)
  assume  $\neg C$ 
  with  $\langle \neg \neg C \rangle$  show  $C$  ..
qed

```

```

theorem tertium_non_datur:  $C \vee \neg C$ 

```

```

proof (rule double_negation)
  show  $\neg \neg (C \vee \neg C)$ 
  proof
    assume  $\neg (C \vee \neg C)$ 
    have  $\neg C$ 
    proof
      assume  $C$  then have  $C \vee \neg C$  ..
      with  $\langle \neg (C \vee \neg C) \rangle$  show  $\bot$  ..
    qed
    then have  $C \vee \neg C$  ..
    with  $\langle \neg (C \vee \neg C) \rangle$  show  $\bot$  ..
  qed
qed

```

These examples illustrate both classical reasoning and non-trivial propositional proofs in general. All three rules characterize classical logic independently, but the original rule is already the most convenient to use, because it leaves the conclusion unchanged. Note that $(\neg C \Longrightarrow C) \Longrightarrow C$ fits again into our format for eliminations, despite the additional twist that the context refers to the main conclusion. So we may write *classical* as the Isar state-

ment “**obtains** \neg *thesis*”. This also explains nicely how classical reasoning really works: whatever the main *thesis* might be, we may always assume its negation!

end

2.3.5 Quantifiers

Representing quantifiers is easy, thanks to the higher-order nature of the underlying framework. According to the well-known technique introduced by Church [6], quantifiers are operators on predicates, which are syntactically represented as λ -terms of type $i \Rightarrow o$. Binder notation turns *All* $(\lambda x. B\ x)$ into $\forall x. B\ x$ etc.

axiomatization

All :: $(i \Rightarrow o) \Rightarrow o$ (**binder** \forall 10) **where**
allI [*intro*]: $(\bigwedge x. B\ x) \Longrightarrow \forall x. B\ x$ **and**
allD [*dest*]: $(\forall x. B\ x) \Longrightarrow B\ a$

axiomatization

Ex :: $(i \Rightarrow o) \Rightarrow o$ (**binder** \exists 10) **where**
exI [*intro*]: $B\ a \Longrightarrow (\exists x. B\ x)$ **and**
exE [*elim*]: $(\exists x. B\ x) \Longrightarrow (\bigwedge x. B\ x \Longrightarrow C) \Longrightarrow C$

The statement of *exE* corresponds to “**assumes** $\exists x. B\ x$ **obtains** x **where** $B\ x$ ” in Isar. In the subsequent example we illustrate quantifier reasoning involving all four rules:

theorem

assumes $\exists x. \forall y. R\ x\ y$
shows $\forall y. \exists x. R\ x\ y$

proof — \forall introduction

obtain x **where** $\forall y. R\ x\ y$ **using** $\langle \exists x. \forall y. R\ x\ y \rangle$.. — \exists elimination

fix y **have** $R\ x\ y$ **using** $\langle \forall y. R\ x\ y \rangle$.. — \forall destruction

then show $\exists x. R\ x\ y$.. — \exists introduction

qed

2.3.6 Canonical reasoning patterns

The main rules of first-order predicate logic from §2.3.3 and §2.3.5 can now be summarized as follows, using the native Isar statement format of §2.2.2.

impI: assumes $A \implies B$ shows $A \longrightarrow B$
impD: assumes $A \longrightarrow B$ and A shows B
*disjI*₁: assumes A shows $A \vee B$
*disjI*₂: assumes B shows $A \vee B$
disjE: assumes $A \vee B$ obtains $A \mid B$
conjI: assumes A and B shows $A \wedge B$
conjE: assumes $A \wedge B$ obtains A and B
falseE: assumes \perp shows A
trueI: shows \top
notI: assumes $A \implies \perp$ shows $\neg A$
notE: assumes $\neg A$ and A shows B
allI: assumes $\bigwedge x. B\ x$ shows $\forall x. B\ x$
allE: assumes $\forall x. B\ x$ shows $B\ a$
exI: assumes $B\ a$ shows $\exists x. B\ x$
exE: assumes $\exists x. B\ x$ obtains a where $B\ a$

This essentially provides a declarative reading of Pure rules as Isar reasoning patterns: the rule statements tells how a canonical proof outline shall look like. Since the above rules have already been declared as *intro*, *elim*, *dest* — each according to its particular shape — we can immediately write Isar proof texts as follows:

<pre> have A \longrightarrow B proof assume A show B $\langle proof \rangle$ qed </pre>	<pre> have A \longrightarrow B and A $\langle proof \rangle$ then have B .. </pre>
<pre> have A $\langle proof \rangle$ then have A \vee B .. </pre>	<pre> have A \vee B $\langle proof \rangle$ then have C proof assume A then show C $\langle proof \rangle$ next assume B then show C $\langle proof \rangle$ qed </pre>
<pre> have B $\langle proof \rangle$ then have A \vee B .. </pre>	<pre> have A \wedge B $\langle proof \rangle$ then obtain A and B .. </pre>

have \perp $\langle proof \rangle$
then have A ..

have \top ..

have $\neg A$
proof
 assume A
 then show \perp $\langle proof \rangle$
qed

have $\neg A$ **and** A $\langle proof \rangle$
then have B ..

have $\forall x. B\ x$
proof
 fix x
 show $B\ x$ $\langle proof \rangle$
qed

have $\forall x. B\ x$ $\langle proof \rangle$
then have $B\ a$..

have $\exists x. B\ x$
proof
 show $B\ a$ $\langle proof \rangle$
qed

have $\exists x. B\ x$ $\langle proof \rangle$
then obtain a **where** $B\ a$..

Of course, these proofs are merely examples. As sketched in §2.2.3, there is a fair amount of flexibility in expressing Pure deductions in Isar. Here the user is asked to express himself adequately, aiming at proof texts of literary quality.

end

Part II

General Language Elements

Outer syntax

The rather generic framework of Isabelle/Isar syntax emerges from three main syntactic categories: *commands* of the top-level Isar engine (covering theory and proof elements), *methods* for general goal refinements (analogous to traditional “tactics”), and *attributes* for operations on facts (within a certain context). Subsequently we give a reference of basic syntactic entities underlying Isabelle/Isar syntax in a bottom-up manner. Concrete theory and proof language elements will be introduced later on.

In order to get started with writing well-formed Isabelle/Isar documents, the most important aspect to be noted is the difference of *inner* versus *outer* syntax. Inner syntax is that of Isabelle types and terms of the logic, while outer syntax is that of Isabelle/Isar theory sources (specifications and proofs). As a general rule, inner syntax entities may occur only as *atomic entities* within outer syntax. For example, the string “ $x + y$ ” and identifier z are legal term specifications within a theory, while $x + y$ without quotes is not.

Printed theory documents usually omit quotes to gain readability (this is a matter of L^AT_EX macro setup, say via `\isabellestyle`, see also [39]). Experienced users of Isabelle/Isar may easily reconstruct the lost technical information, while mere readers need not care about quotes at all.

Isabelle/Isar input may contain any number of input termination characters “;” (semicolon) to separate commands explicitly. This is particularly useful in interactive shell sessions to make clear where the current command is intended to end. Otherwise, the interpreter loop will continue to issue a secondary prompt “#” until an end-of-command is clearly recognized from the input syntax, e.g. encounter of the next command keyword.

More advanced interfaces such as Proof General [1] do not require explicit semicolons, the amount of input text is determined automatically by inspecting the present content of the Emacs text buffer. In the printed presentation of Isabelle/Isar documents semicolons are omitted altogether for readability.

! • Proof General requires certain syntax classification tables in order to achieve properly synchronized interaction with the Isabelle/Isar process. These tables need to be consistent with the Isabelle version and particular logic image to be used

in a running session (common object-logics may well change the outer syntax). The standard setup should work correctly with any of the “official” logic images derived from Isabelle/HOL (including HOLCF etc.). Users of alternative logics may need to tell Proof General explicitly, e.g. by giving an option `-k ZF` (in conjunction with `-l ZF`, to specify the default logic image). Note that option `-L` does both of this at the same time.

3.1 Lexical matters

The outer lexical syntax consists of three main categories of syntax tokens:

1. *major keywords* — the command names that are available in the present logic session;
2. *minor keywords* — additional literal tokens required by the syntax of commands;
3. *named tokens* — various categories of identifiers etc.

Major keywords and minor keywords are guaranteed to be disjoint. This helps user-interfaces to determine the overall structure of a theory text, without knowing the full details of command syntax. Internally, there is some additional information about the kind of major keywords, which approximates the command type (theory command, proof command etc.).

Keywords override named tokens. For example, the presence of a command called `term` inhibits the identifier `term`, but the string `"term"` can be used instead. By convention, the outer syntax always allows quoted strings in addition to identifiers, wherever a named entity is expected.

When tokenizing a given input sequence, the lexer repeatedly takes the longest prefix of the input that forms a valid token. Spaces, tabs, newlines and formfeeds between tokens serve as explicit separators.

The categories for named tokens are defined once and for all as follows.

$$\begin{aligned}
 \textit{ident} &= \textit{letter quasiletter}^* \\
 \textit{longident} &= \textit{ident}(\textit{. ident})^+ \\
 \textit{symident} &= \textit{sym}^+ \mid \backslash \langle \textit{ident} \rangle \\
 \textit{nat} &= \textit{digit}^+ \\
 \textit{float} &= \textit{nat}.\textit{nat} \mid -\textit{nat}.\textit{nat} \\
 \textit{var} &= ?\textit{ident} \mid ?\textit{ident}.\textit{nat}
 \end{aligned}$$

```

typefree  = 'ident
typevar   = ?typefree | ?typefree.nat
string    = " ... "
altstring = ' ... '
verbatim  = { * ... * }

letter    = latin | \<latin> | \<latin latin> | greek |
           \<^isub> | \<^isup>
quasiletter = letter | digit | _ | '
latin     = a | ... | z | A | ... | Z
digit     = 0 | ... | 9
sym       = ! | # | $ | % | & | * | + | - | / |
           < | = | > | ? | @ | ^ | _ | | | ~
greek     = \<alpha> | \<beta> | \<gamma> | \<delta> |
           \<epsilon> | \<zeta> | \<eta> | \<theta> |
           \<iota> | \<kappa> | \<mu> | \<nu> |
           \<xi> | \<pi> | \<rho> | \<sigma> | \<tau> |
           \<upsilon> | \<phi> | \<chi> | \<psi> |
           \<omega> | \<Gamma> | \<Delta> | \<Theta> |
           \<Lambda> | \<Xi> | \<Pi> | \<Sigma> |
           \<Upsilon> | \<Phi> | \<Psi> | \<Omega>

```

A *var* or *typevar* describes an unknown, which is internally a pair of base name and index (ML type `indexname`). These components are either separated by a dot as in `?x.1` or `?x7.3` or run together as in `?x1`. The latter form is possible if the base name does not end with digits. If the index is 0, it may be dropped altogether: `?x` and `?x0` and `?x.0` all refer to the same unknown, with basename *x* and index 0.

The syntax of *string* admits any characters, including newlines; “” (double-quote) and “\” (backslash) need to be escaped by a backslash; arbitrary character codes may be specified as “\ddd”, with three decimal digits. Alternative strings according to *altstring* are analogous, using single back-quotes instead.

The body of *verbatim* may consist of any text not containing “*}”; this allows convenient inclusion of quotes without further escapes. There is no way to escape “*}”. If the quoted text is L^AT_EX source, one may usually add some blank or comment to avoid the critical character sequence.

Source comments take the form `(* ... *)` and may be nested, although the user-interface might prevent this. Note that this form indicates source comments only, which are stripped after lexical analysis of the input. The Isar syntax also provides proper *document comments* that are considered as part of the text (see §3.2.3).

Common mathematical symbols such as \forall are represented in Isabelle as `\<forall>`. There are infinitely many Isabelle symbols like this, although proper presentation is left to front-end tools such as \LaTeX or Proof General with the X-Symbol package. A list of predefined Isabelle symbols that work well with these tools is given in appendix B. Note that `\<lambda>` does not belong to the *letter* category, since it is already used differently in the Pure term language.

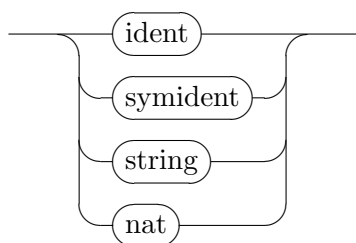
3.2 Common syntax entities

We now introduce several basic syntactic entities, such as names, terms, and theorem specifications, which are factored out of the actual Isar language elements to be described later.

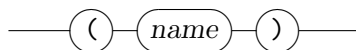
3.2.1 Names

Entity *name* usually refers to any name of types, constants, theorems etc. that are to be *declared* or *defined* (so qualified identifiers are excluded here). Quoted strings provide an escape for non-identifier names or those ruled out by outer syntax keywords (e.g. quoted `"let"`). Already existing objects are usually referenced by *nameref*.

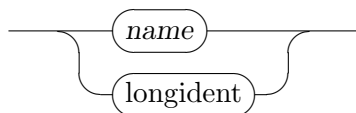
name



parname



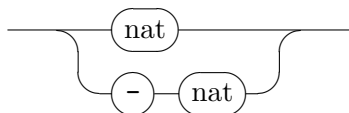
nameref



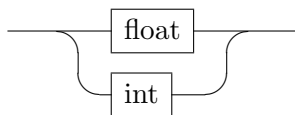
3.2.2 Numbers

The outer lexical syntax (§3.1) admits natural numbers and floating point numbers. These are combined as *int* and *real* as follows.

int



real

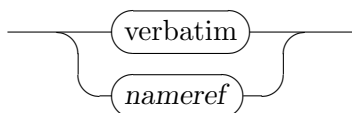


Note that there is an overlap with the category *name*, which also includes *nat*.

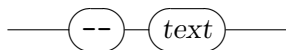
3.2.3 Comments

Large chunks of plain *text* are usually given verbatim, i.e. enclosed in `{* ... *}`. For convenience, any of the smaller text units conforming to *nameref* are admitted as well. A marginal *comment* is of the form `-- text`. Any number of these may occur within Isabelle/Isar commands.

text

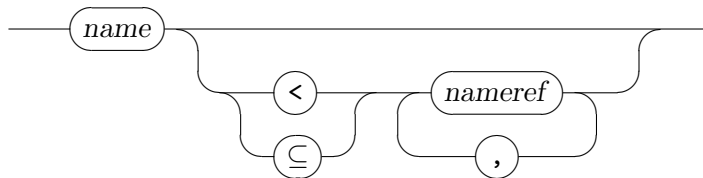
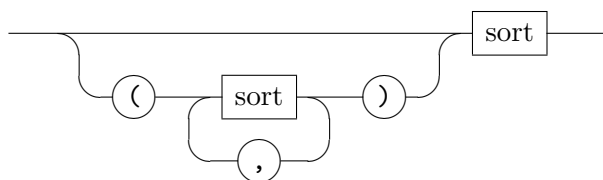


comment



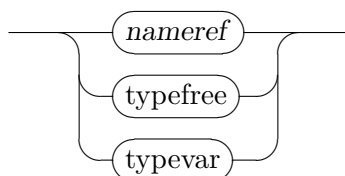
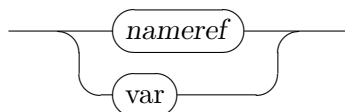
3.2.4 Type classes, sorts and arities

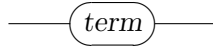
Classes are specified by plain names. Sorts have a very simple inner syntax, which is either a single class name c or a list $\{c_1, \dots, c_n\}$ referring to the intersection of these classes. The syntax of type arities is given directly at the outer level.

classdecl*sort**arity*

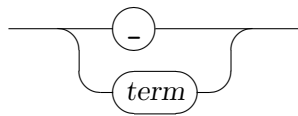
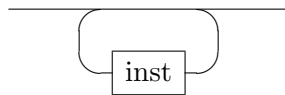
3.2.5 Types and terms

The actual inner Isabelle syntax, that of types and terms of the logic, is far too sophisticated in order to be modelled explicitly at the outer theory level. Basically, any such entity has to be quoted to turn it into a single token (the parsing and type-checking is performed internally later). For convenience, a slightly more liberal convention is adopted: quotes may be omitted for any type or term that is already atomic at the outer level. For example, one may just write x instead of quoted $"x"$. Note that symbolic identifiers (e.g. $++$ or \forall) are available as well, provided these have not been superseded by commands or other keywords already (such as $=$ or $+$).

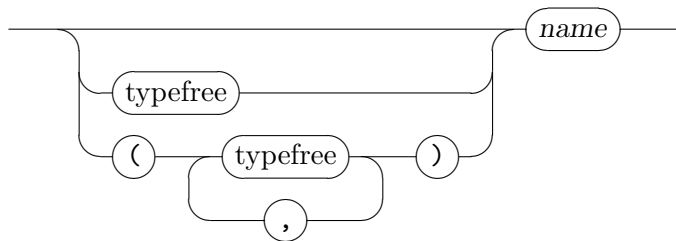
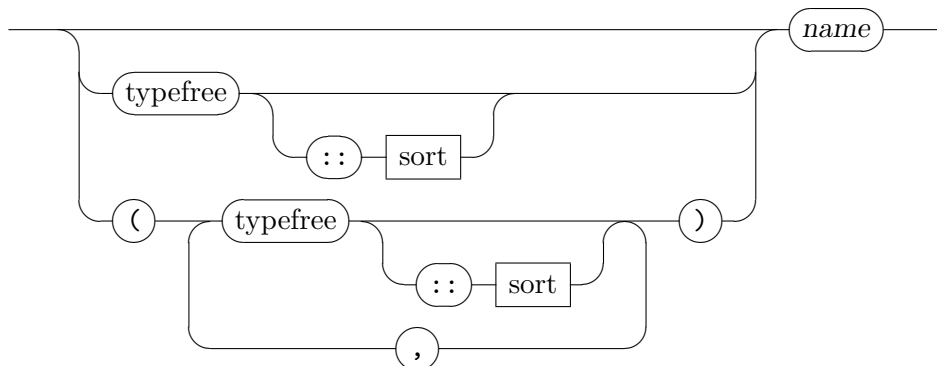
type*term*

prop

Positional instantiations are indicated by giving a sequence of terms, or the placeholder “_” (underscore), which means to skip a position.

inst*insts*

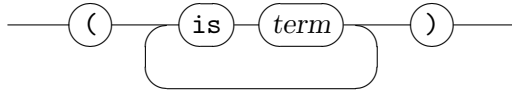
Type declarations and definitions usually refer to *typespec* on the left-hand side. This models basic type constructor application at the outer syntax level. Note that only plain postfix notation is available here, but no infixes.

typespec*typespecsorts*

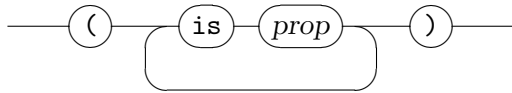
3.2.6 Term patterns and declarations

Wherever explicit propositions (or term fragments) occur in a proof text, casual binding of schematic term variables may be given specified via patterns of the form “(is $p_1 \dots p_n$)”. This works both for *term* and *prop*.

termpat

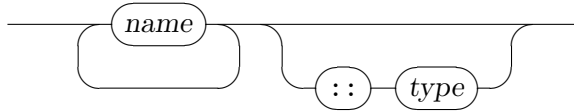


proppat

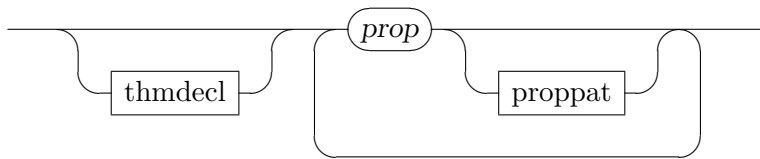


Declarations of local variables $x :: \tau$ and logical propositions $a : \varphi$ represent different views on the same principle of introducing a local scope. In practice, one may usually omit the typing of *vars* (due to type-inference), and the naming of propositions (due to implicit references of current facts). In any case, Isar proof elements usually admit to introduce multiple such items simultaneously.

vars



props

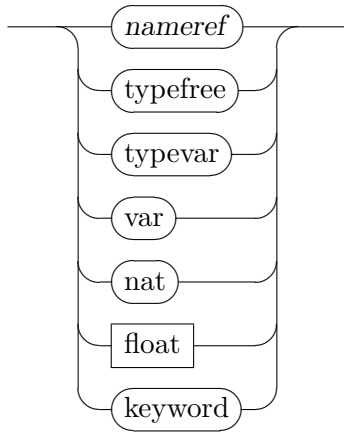


The treatment of multiple declarations corresponds to the complementary focus of *vars* versus *props*. In “ $x_1 \dots x_n :: \tau$ ” the typing refers to all variables, while in $a : \varphi_1 \dots \varphi_n$ the naming refers to all propositions collectively. Isar language elements that refer to *vars* or *props* typically admit separate typings or namings via another level of iteration, with explicit **and** separators; e.g. see **fix** and **assume** in §6.2.1.

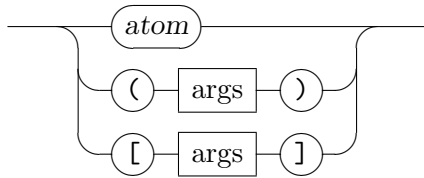
3.2.7 Attributes and theorems

Attributes have their own “semi-inner” syntax, in the sense that input conforming to *args* below is parsed by the attribute a second time. The attribute argument specifications may be any sequence of atomic entities (identifiers, strings etc.), or properly bracketed argument lists. Below *atom* refers to any atomic entity, including any keyword conforming to *symident*.

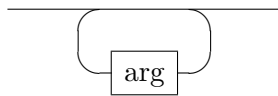
atom



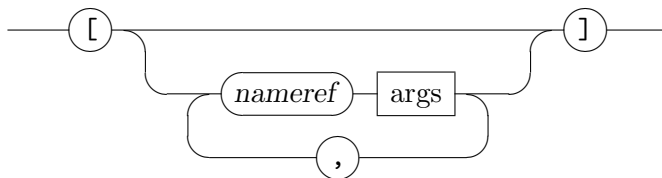
arg



args



attributes



Theorem specifications come in several flavors: *axmdecl* and *thmdecl* usually refer to axioms, assumptions or results of goal statements, while *thmdef*

collects lists of existing theorems. Existing theorems are given by *thmref* and *thmrefs*, the former requires an actual singleton result.

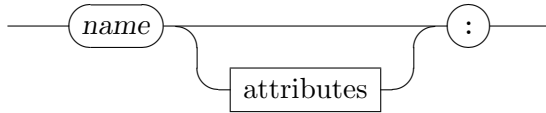
There are three forms of theorem references:

1. named facts a ,
2. selections from named facts $a(i)$ or $a(j - k)$,
3. literal fact propositions using *altstring* syntax ‘ φ ’ (see also method *fact*).

Any kind of theorem specification may include lists of attributes both on the left and right hand sides; attributes are applied to any immediately preceding fact. If names are omitted, the theorems are not stored within the theorem database of the theory or proof context, but any given attributes are applied nonetheless.

An extra pair of brackets around attributes (like “ $[[\textit{simplproc } a]]$ ”) abbreviates a theorem reference involving an internal dummy fact, which will be ignored later on. So only the effect of the attribute on the background context will persist. This form of in-place declarations is particularly useful with commands like **declare** and **using**.

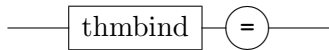
axmdecl



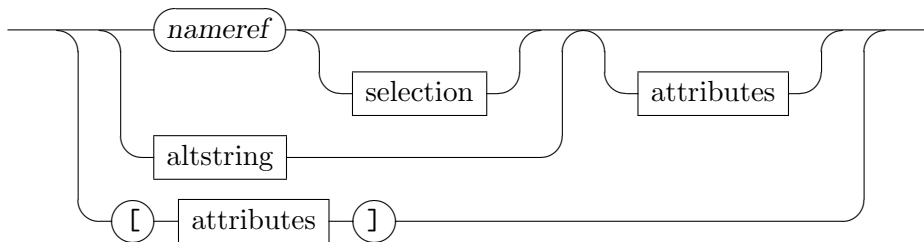
thmdecl



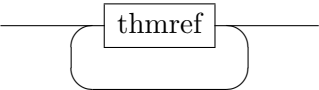
thmdef



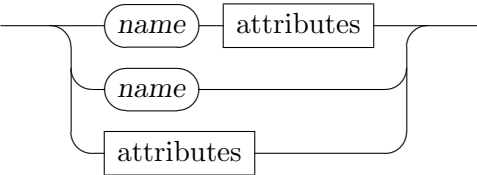
thmref



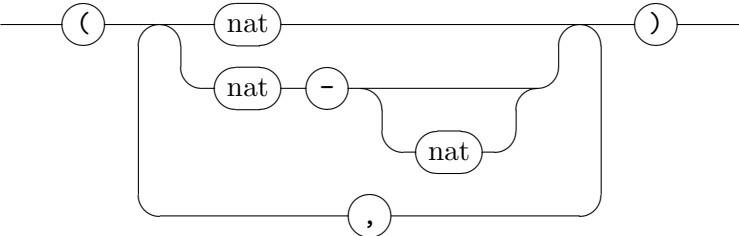
thmrefs



thmbind



selection



Document preparation

Isabelle/Isar provides a simple document preparation system based on regular PDF- \LaTeX technology, with full support for hyper-links and bookmarks. Thus the results are well suited for WWW browsing and as printed copies.

Isabelle generates \LaTeX output while running a *logic session* (see also [39]). Getting started with a working configuration for common situations is quite easy by using the Isabelle `mkdir` and `make` tools. First invoke

```
isabelle mkdir Foo
```

to initialize a separate directory for session `Foo` (it is safe to experiment, since `isabelle mkdir` never overwrites existing files). Ensure that `Foo/ROOT.ML` holds ML commands to load all theories required for this session; furthermore `Foo/document/root.tex` should include any special \LaTeX macro packages required for your document (the default is usually sufficient as a start).

The session is controlled by a separate `IsaMakefile` (with crude source dependencies by default). This file is located one level up from the `Foo` directory location. Now invoke

```
isabelle make Foo
```

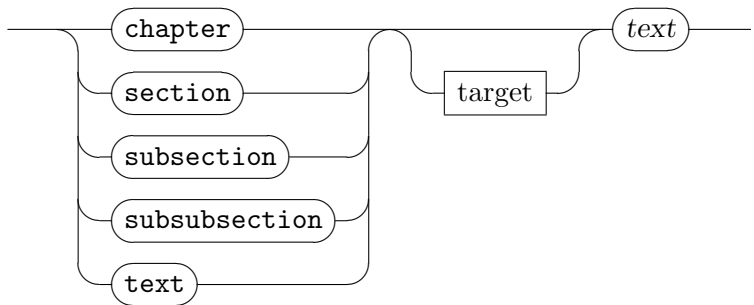
to run the `Foo` session, with browser information and document preparation enabled. Unless any errors are reported by Isabelle or \LaTeX , the output will appear inside the directory defined by the `ISABELLE_BROWSER_INFO` setting (as reported by the batch job in verbose mode).

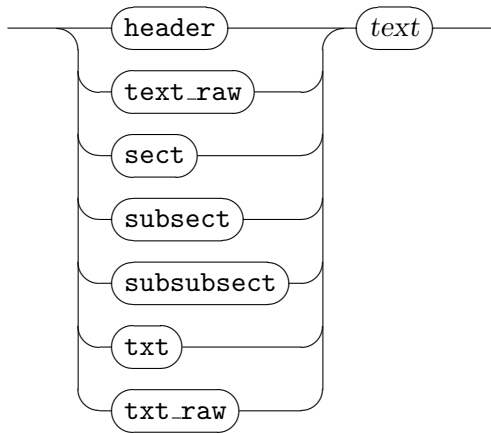
You may also consider to tune the `usedir` options in `IsaMakefile`, for example to switch the output format between `pdf` and `dvi`, or activate the `-D` option to retain a second copy of the generated \LaTeX sources (for manual inspection or separate runs of `latex`).

See *The Isabelle System Manual* [39] for further details on Isabelle logic sessions and theory presentation. The Isabelle/HOL tutorial [19] also covers theory presentation to some extent.

header	:	<i>toplevel</i>	\rightarrow	<i>toplevel</i>
chapter	:	<i>local_theory</i>	\rightarrow	<i>local_theory</i>
section	:	<i>local_theory</i>	\rightarrow	<i>local_theory</i>
subsection	:	<i>local_theory</i>	\rightarrow	<i>local_theory</i>
subsubsection	:	<i>local_theory</i>	\rightarrow	<i>local_theory</i>
text	:	<i>local_theory</i>	\rightarrow	<i>local_theory</i>
text_raw	:	<i>local_theory</i>	\rightarrow	<i>local_theory</i>
sect	:	<i>proof</i>	\rightarrow	<i>proof</i>
subject	:	<i>proof</i>	\rightarrow	<i>proof</i>
subsubject	:	<i>proof</i>	\rightarrow	<i>proof</i>
txt	:	<i>proof</i>	\rightarrow	<i>proof</i>
txt_raw	:	<i>proof</i>	\rightarrow	<i>proof</i>

Note that formal comments (§3.2.3) are similar to markup commands, but have a different status within Isabelle/Isar syntax.





header provides plain text markup just preceding the formal beginning of a theory. The corresponding \LaTeX macro is `\isamarkupheader`, which acts like **section** by default.

chapter, **section**, **subsection**, and **subsubsection** mark chapter and section headings within the main theory body or local theory targets. The corresponding \LaTeX macros are `\isamarkupchapter`, `\isamarkupsection`, `\isamarkupsubsection` etc.

sect, **subject**, and **subsubsect** mark section headings within proofs. The corresponding \LaTeX macros are `\isamarkupsect`, `\isamarkupsubsect` etc.

text and **txt** specify paragraphs of plain text. This corresponds to a \LaTeX environment `\begin{isamarkuptext} ... \end{isamarkuptext}` etc.

text_raw and **txt_raw** insert \LaTeX source into the output, without additional markup. Thus the full range of document manipulations becomes available, at the risk of messing up document output.

Except for **text_raw** and **txt_raw**, the text passed to any of the above markup commands may refer to formal entities via *document antiquotations*, see also §4.2. These are interpreted in the present theory or proof context, or the named *target*.

The proof markup commands closely resemble those for theory specifications, but have a different formal status and produce different \LaTeX macros. The default definitions coincide for analogous commands such as **section** and **sect**.

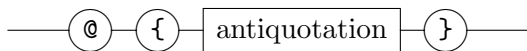
4.2 Document Antiquotations

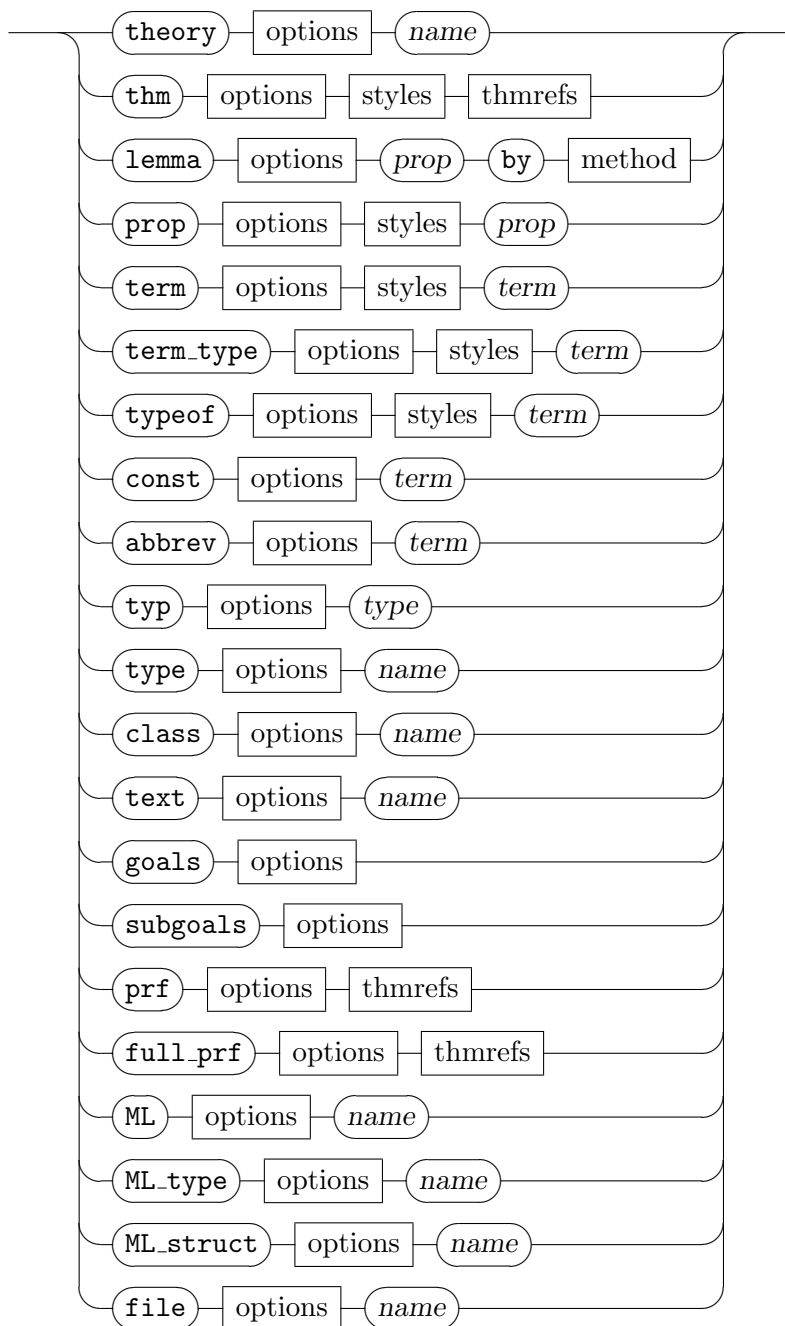
$theory$: antiquotation
 thm : antiquotation
 $lemma$: antiquotation
 $prop$: antiquotation
 $term$: antiquotation
 $term_type$: antiquotation
 $typeof$: antiquotation
 $const$: antiquotation
 $abbrev$: antiquotation
 typ : antiquotation
 $type$: antiquotation
 $class$: antiquotation
 $text$: antiquotation
 $goals$: antiquotation
 $subgoals$: antiquotation
 prf : antiquotation
 $full_prf$: antiquotation
 ML : antiquotation
 ML_type : antiquotation
 ML_struct : antiquotation
 $file$: antiquotation

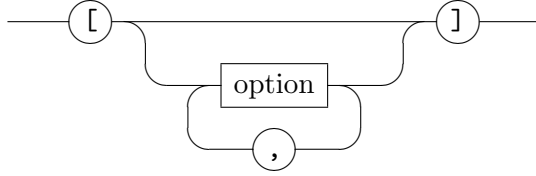
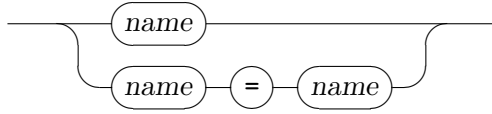
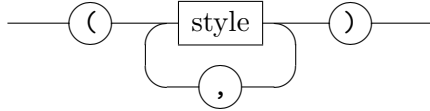
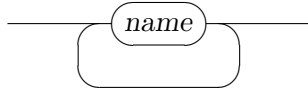
The overall content of an Isabelle/Isar theory may alternate between formal and informal text. The main body consists of formal specification and proof commands, interspersed with markup commands (§4.1) or document comments (§3.2.3). The argument of markup commands quotes informal text to be printed in the resulting document, but may again refer to formal entities via *document antiquotations*.

For example, embedding of “ $@\{term\ [show_types]\ f\ x = a + x\}$ ” within a text block makes $(f::'a \Rightarrow 'a)\ (x::'a) = (a::'a) + x$ appear in the final L^AT_EX document.

Antiquotations usually spare the author tedious typing of logical entities in full detail. Even more importantly, some degree of consistency-checking between the main body of formal text and its informal explanation is achieved, since terms and types appearing in antiquotations are checked within the current theory or proof context.



antiquotation

options*option**styles**style*

Note that the syntax of antiquotations may *not* include source comments (`* ... *`) nor verbatim text `{* ... *}`.

`@{theory A}` prints the name A , which is guaranteed to refer to a valid ancestor theory in the current context.

`@{thm $a_1 \dots a_n$ }` prints theorems $a_1 \dots a_n$. Full fact expressions are allowed here, including attributes (§3.2.7).

`@{prop φ }` prints a well-typed proposition φ .

`@{lemma φ by m }` proves a well-typed proposition φ by method m and prints the original φ .

`@{term t }` prints a well-typed term t .

`@{term_type t }` prints a well-typed term t annotated with its type.

`@{typeof t }` prints the type of a well-typed term t .

`@{const c }` prints a logical or syntactic constant c .

$\text{@}\{abbrev\ c\ x_1 \dots x_n\}$ prints a constant abbreviation $c\ x_1 \dots x_n \equiv rhs$ as defined in the current context.

$\text{@}\{typ\ \tau\}$ prints a well-formed type τ .

$\text{@}\{type\ \kappa\}$ prints a (logical or syntactic) type constructor κ .

$\text{@}\{class\ c\}$ prints a class c .

$\text{@}\{text\ s\}$ prints uninterpreted source text s . This is particularly useful to print portions of text according to the Isabelle document style, without demanding well-formedness, e.g. small pieces of terms that should not be parsed or type-checked yet.

$\text{@}\{goals\}$ prints the current *dynamic* goal state. This is mainly for support of tactic-emulation scripts within Isar. Presentation of goal states does not conform to the idea of human-readable proof documents!

When explaining proofs in detail it is usually better to spell out the reasoning via proper Isar proof commands, instead of peeking at the internal machine configuration.

$\text{@}\{subgoals\}$ is similar to $\text{@}\{goals\}$, but does not print the main goal.

$\text{@}\{prf\ a_1 \dots a_n\}$ prints the (compact) proof terms corresponding to the theorems $a_1 \dots a_n$. Note that this requires proof terms to be switched on for the current logic session.

$\text{@}\{full_prf\ a_1 \dots a_n\}$ is like $\text{@}\{prf\ a_1 \dots a_n\}$, but prints the full proof terms, i.e. also displays information omitted in the compact proof term, which is denoted by “_” placeholders there.

$\text{@}\{ML\ s\}$, $\text{@}\{ML_type\ s\}$, and $\text{@}\{ML_struct\ s\}$ check text s as ML value, type, and structure, respectively. The source is printed verbatim.

$\text{@}\{file\ path\}$ checks that $path$ refers to a file (or directory) and prints it verbatim.

Styled antiquotations

The antiquotations *thm*, *prop* and *term* admit an extra *style* specification to modify the printed result. A style is specified by a name with a possibly empty number of arguments; multiple styles can be sequenced with commas. The following standard styles are available:

lhs extracts the first argument of any application form with at least two arguments — typically meta-level or object-level equality, or any other binary relation.

rhs is like *lhs*, but extracts the second argument.

concl extracts the conclusion C from a rule in Horn-clause normal form $A_1 \implies \dots A_n \implies C$.

prem n extract premise number n from from a rule in Horn-clause normal form $A_1 \implies \dots A_n \implies C$

General options

The following options are available to tune the printed output of antiquotations. Note that many of these coincide with global ML flags of the same names.

show_types = *bool* and *show_sorts* = *bool* control printing of explicit type and sort constraints.

show_structs = *bool* controls printing of implicit structures.

show_abbrevs = *bool* controls folding of abbreviations.

long_names = *bool* forces names of types and constants etc. to be printed in their fully qualified internal form.

short_names = *bool* forces names of types and constants etc. to be printed unqualified. Note that internalizing the output again in the current context may well yield a different result.

unique_names = *bool* determines whether the printed version of qualified names should be made sufficiently long to avoid overlap with names declared further back. Set to *false* for more concise output.

eta_contract = *bool* prints terms in η -contracted form.

display = *bool* indicates if the text is to be output as multi-line “display material”, rather than a small piece of text without line breaks (which is the default).

In this mode the embedded entities are printed in the same style as the main theory text.

break = bool controls line breaks in non-display material.

quotes = bool indicates if the output should be enclosed in double quotes.

mode = name adds *name* to the print mode to be used for presentation.
Note that the standard setup for L^AT_EX output is already present by default, including the modes *latex* and *xsymbols*.

margin = nat and *indent = nat* change the margin or indentation for pretty printing of display material.

goals_limit = nat determines the maximum number of goals to be printed (for goal-based antiquotation).

source = bool prints the original source text of the antiquotation arguments, rather than its internal representation. Note that formal checking of *thm*, *term*, etc. is still enabled; use the *text* antiquotation for unchecked output.

Regular *term* and *typ* antiquotations with *source = false* involve a full round-trip from the original source to an internalized logical entity back to a source form, according to the syntax of the current context. Thus the printed output is not under direct control of the author, it may even fluctuate a bit as the underlying theory is changed later on.

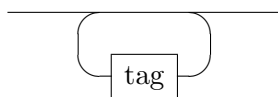
In contrast, *source = true* admits direct printing of the given source text, with the desirable well-formedness check in the background, but without modification of the printed text.

For boolean flags, “*name = true*” may be abbreviated as “*name*”. All of the above flags are disabled by default, unless changed from ML, say in the `ROOT.ML` of the logic session.

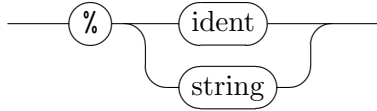
4.3 Markup via command tags

Each Isabelle/Isar command may be decorated by additional presentation tags, to indicate some modification in the way it is printed in the document.

tags



tag



Some tags are pre-declared for certain classes of commands, serving as default markup if no tags are given in the text:

<i>theory</i>	theory begin/end
<i>proof</i>	all proof commands
<i>ML</i>	all commands involving ML code

The Isabelle document preparation system [39] allows tagged command regions to be presented specifically, e.g. to fold proof texts, or drop parts of the text completely.

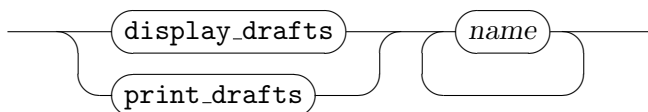
For example “**by** %*invisible auto*” causes that piece of proof to be treated as *invisible* instead of *proof* (the default), which may be shown or hidden depending on the document setup. In contrast, “**by** %*visible auto*” forces this text to be shown invariably.

Explicit tag specifications within a proof apply to all subsequent commands of the same level of nesting. For example, “**proof** %*visible* ... **qed**” forces the whole sub-proof to be typeset as *visible* (unless some of its parts are tagged differently).

Command tags merely produce certain markup environments for typesetting. The meaning of these is determined by \LaTeX macros, as defined in `~/lib/texinputs/isabelle.sty` or by the document author. The Isabelle document preparation tools also provide some high-level options to specify the meaning of arbitrary tags to “keep”, “drop”, or “fold” the corresponding parts of the text. Logic sessions may also specify “document versions”, where given tags are interpreted in some particular way. Again see [39] for further details.

4.4 Draft presentation

`display_drafts*` : *any* \rightarrow
`print_drafts*` : *any* \rightarrow



display_drafts *paths* and **print_drafts** *paths* perform simple output of a given list of raw source files. Only those symbols that do not require additional L^AT_EX packages are displayed properly, everything else is left verbatim.

Theory specifications

The Isabelle/Isar theory format integrates specifications and proofs, supporting interactive development with unlimited undo operation. There is an integrated document preparation system (see chapter 4), for typesetting formal developments together with informal text. The resulting hyper-linked PDF documents can be used both for WWW presentation and printed copies.

The Isar proof language (see chapter 6) is embedded into the theory language as a proper sub-language. Proof mode is entered by stating some **theorem** or **lemma** at the theory level, and left again with the final conclusion (e.g. via **qed**). Some theory specification mechanisms also require a proof, such as **typedef** in HOL, which demands non-emptiness of the representing sets.

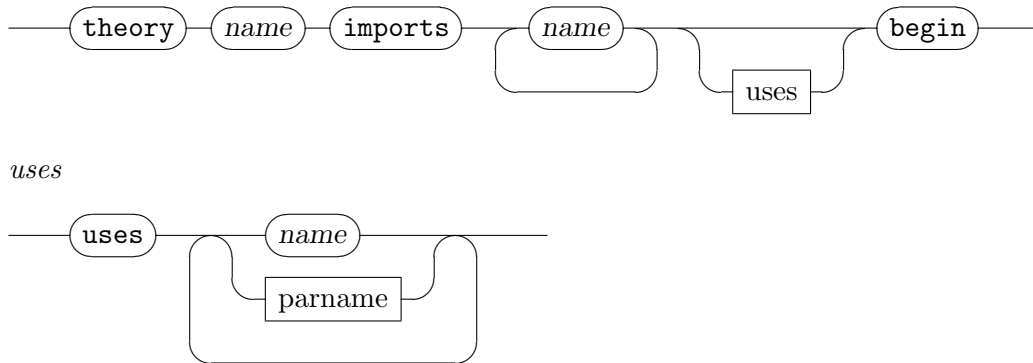
5.1 Defining theories

theory : *oplevel* \rightarrow *theory*
end : *theory* \rightarrow *oplevel*

Isabelle/Isar theories are defined via theory files, which may contain both specifications and proofs; occasionally definitional mechanisms also require some explicit proof. The theory body may be sub-structured by means of *local theory targets*, such as **locale** and **class**.

The first proper command of a theory is **theory**, which indicates imports of previous theories and optional dependencies on other source files (usually in ML). Just preceding the initial **theory** command there may be an optional **header** declaration, which is only relevant to document preparation: see also the other section markup commands in §4.1.

A theory is concluded by a final **end** command, one that does not belong to a local theory target. No further commands may follow such a global **end**, although some user-interfaces might pretend that trailing input is admissible.



theory A **imports** $B_1 \dots B_n$ **begin** starts a new theory A based on the merge of existing theories $B_1 \dots B_n$.

Due to the possibility to import more than one ancestor, the resulting theory structure of an Isabelle session forms a directed acyclic graph (DAG). Isabelle’s theory loader ensures that the sources contributing to the development graph are always up-to-date: changed files are automatically reloaded whenever a theory header specification is processed.

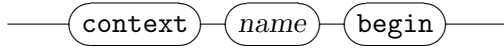
The optional **uses** specification declares additional dependencies on extra files (usually ML sources). Files will be loaded immediately (as ML), unless the name is parenthesized. The latter case records a dependency that needs to be resolved later in the text, usually via explicit **use** for ML files; other file formats require specific load commands defined by the corresponding tools or packages.

end concludes the current theory definition. Note that local theory targets involve a local **end**, which is clear from the nesting.

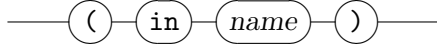
5.2 Local theory targets

A local theory target is a context managed separately within the enclosing theory. Contexts may introduce parameters (fixed variables) and assumptions (hypotheses). Definitions and theorems depending on the context may be added incrementally later on. Named contexts refer to locales (cf. §5.5) or type classes (cf. §5.6); the name “—” signifies the global theory context.

context : $theory \rightarrow local_theory$
end : $local_theory \rightarrow theory$



target



context *c* **begin** recommences an existing locale or class context *c*. Note that locale and class definitions allow to include the **begin** keyword as well, in order to continue the local theory immediately after the initial specification.

end concludes the current local theory and continues the enclosing global theory. Note that a global **end** has a different meaning: it concludes the theory itself (§5.1).

(**in** *c*) given after any local theory command specifies an immediate target, e.g. “**definition** (**in** *c*) ...” or “**theorem** (**in** *c*) ...”. This works both in a local or global theory context; the current target context will be suspended for this command only. Note that “(**in** *–*)” will always produce a global result independently of the current target context.

The exact meaning of results produced within a local theory context depends on the underlying target infrastructure (locale, type class etc.). The general idea is as follows, considering a context named *c* with parameter *x* and assumption *A*[*x*].

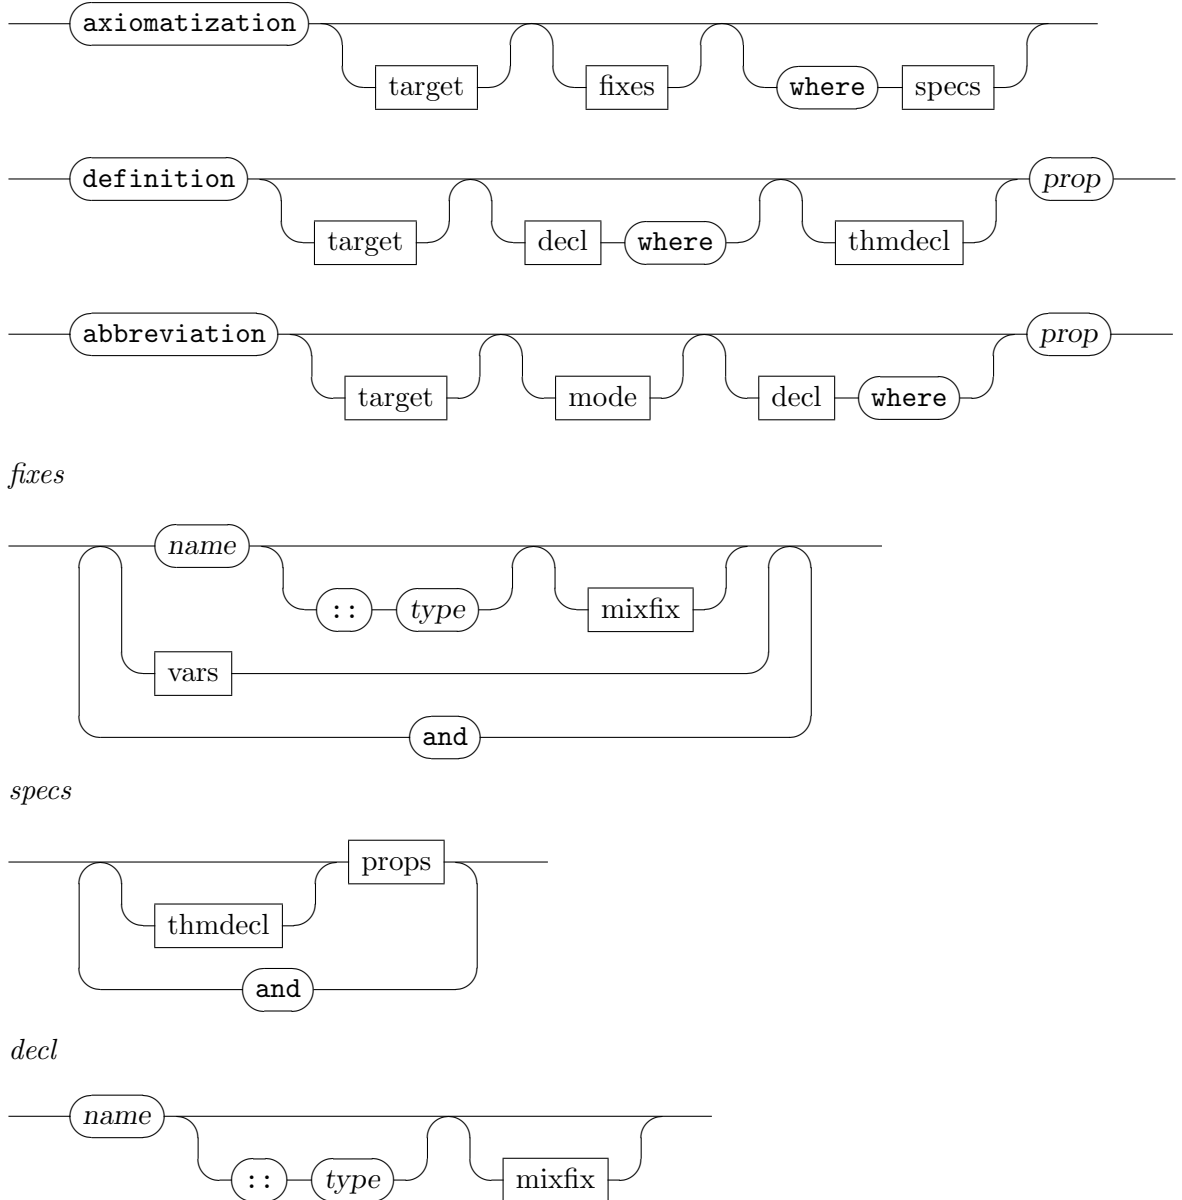
Definitions are exported by introducing a global version with additional arguments; a syntactic abbreviation links the long form with the abstract version of the target context. For example, $a \equiv t[x]$ becomes $c.a \text{ ?}x \equiv t[\text{?}x]$ at the theory level (for arbitrary *?x*), together with a local abbreviation $c \equiv c.a \text{ } x$ in the target context (for the fixed parameter *x*).

Theorems are exported by discharging the assumptions and generalizing the parameters of the context. For example, $a: B[x]$ becomes $c.a: A[\text{?}x] \implies B[\text{?}x]$, again for arbitrary *?x*.

5.3 Basic specification elements

axiomatization	: <i>theory</i> \rightarrow <i>theory</i>	(<i>axiomatic!</i>)
definition	: <i>local_theory</i> \rightarrow <i>local_theory</i>	
defn	: <i>attribute</i>	
abbreviation	: <i>local_theory</i> \rightarrow <i>local_theory</i>	
print_abbrevs*	: <i>context</i> \rightarrow	

These specification mechanisms provide a slightly more abstract view than the underlying primitives of **consts**, **defs** (see §5.9.3), and **axioms** (see §5.10). In particular, type-inference is commonly available, and result names need not be given.



axiomatization $c_1 \dots c_m$ **where** $\varphi_1 \dots \varphi_n$ introduces several constants simultaneously and states axiomatic properties for these. The constants are marked as being specified once and for all, which prevents additional specifications being issued later on.

Note that axiomatic specifications are only appropriate when declaring a new logical system; axiomatic specifications are restricted to global theory contexts. Normal applications should only use definitional mechanisms!

definition *c where eq* produces an internal definition $c \equiv t$ according to the specification given as *eq*, which is then turned into a proven fact. The given proposition may deviate from internal meta-level equality according to the rewrite rules declared as *defn* by the object-logic. This usually covers object-level equality $x = y$ and equivalence $A \leftrightarrow B$. End-users normally need not change the *defn* setup.

Definitions may be presented with explicit arguments on the LHS, as well as additional conditions, e.g. $f\ x\ y = t$ instead of $f \equiv \lambda x\ y. t$ and $y \neq 0 \implies g\ x\ y = u$ instead of an unrestricted $g \equiv \lambda x\ y. u$.

abbreviation *c where eq* introduces a syntactic constant which is associated with a certain term according to the meta-level equality *eq*.

Abbreviations participate in the usual type-inference process, but are expanded before the logic ever sees them. Pretty printing of terms involves higher-order rewriting with rules stemming from reverted abbreviations. This needs some care to avoid overlapping or looping syntactic replacements!

The optional *mode* specification restricts output to a particular print mode; using “*input*” here achieves the effect of one-way abbreviations. The mode may also include an “**output**” qualifier that affects the concrete syntax declared for abbreviations, cf. **syntax** in §7.6.

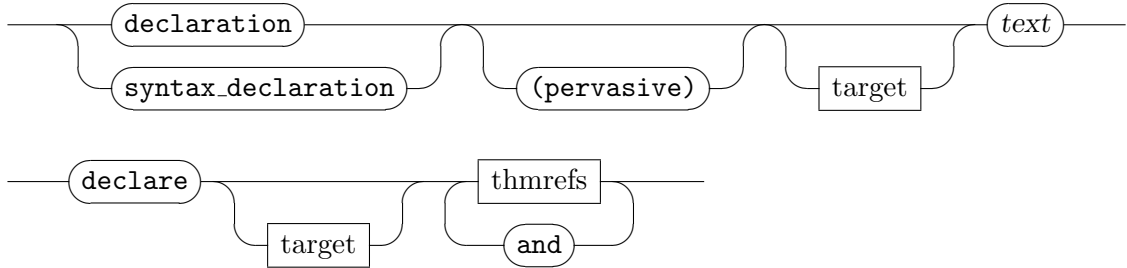
print_abbrevs prints all constant abbreviations of the current context.

5.4 Generic declarations

Arbitrary operations on the background context may be wrapped-up as generic declaration elements. Since the underlying concept of local theories may be subject to later re-interpretation, there is an additional dependency on a morphism that tells the difference of the original declaration context wrt. the application context encountered later on. A fact declaration is an

important special case: it consists of a theorem which is applied to the context by means of an attribute.

declaration : $local_theory \rightarrow local_theory$
syntax_declaration : $local_theory \rightarrow local_theory$
declare : $local_theory \rightarrow local_theory$



declaration d adds the declaration function d of ML type **declaration**, to the current local theory under construction. In later application contexts, the function is transformed according to the morphisms being involved in the interpretation hierarchy.

If the *(pervasive)* option is given, the corresponding declaration is applied to all possible contexts involved, including the global background theory.

syntax_declaration is similar to **declaration**, but is meant to affect only “syntactic” tools by convention (such as notation and type-checking information).

declare $thms$ declares theorems to the current local theory context. No theorem binding is involved here, unlike **theorems** or **lemmas** (cf. §5.10), so **declare** only has the effect of applying attributes as included in the theorem specification.

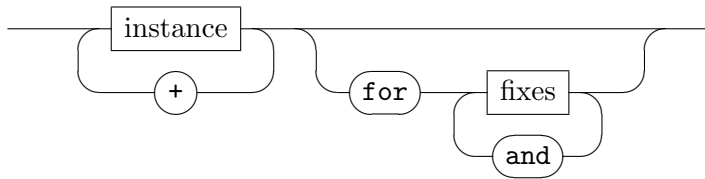
5.5 Locales

Locales are parametric named local contexts, consisting of a list of declaration elements that are modeled after the Isar proof context commands (cf. §6.2.1).

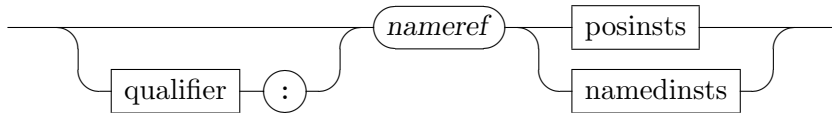
5.5.1 Locale expressions

A *locale expression* denotes a structured context composed of instances of existing locales. The context consists of a list of instances of declaration elements from the locales. Two locale instances are equal if they are of the same locale and the parameters are instantiated with equivalent terms. Declaration elements from equal instances are never repeated, thus avoiding duplicate declarations.

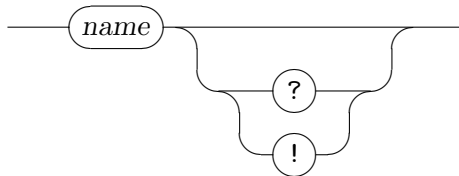
localeexpr



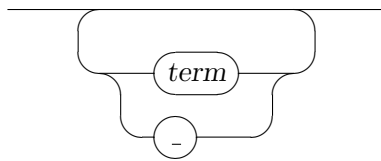
instance



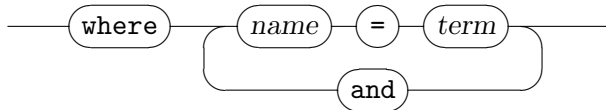
qualifier



posinsts



namedinsts



A locale instance consists of a reference to a locale and either positional or named parameter instantiations. Identical instantiations (that is, those

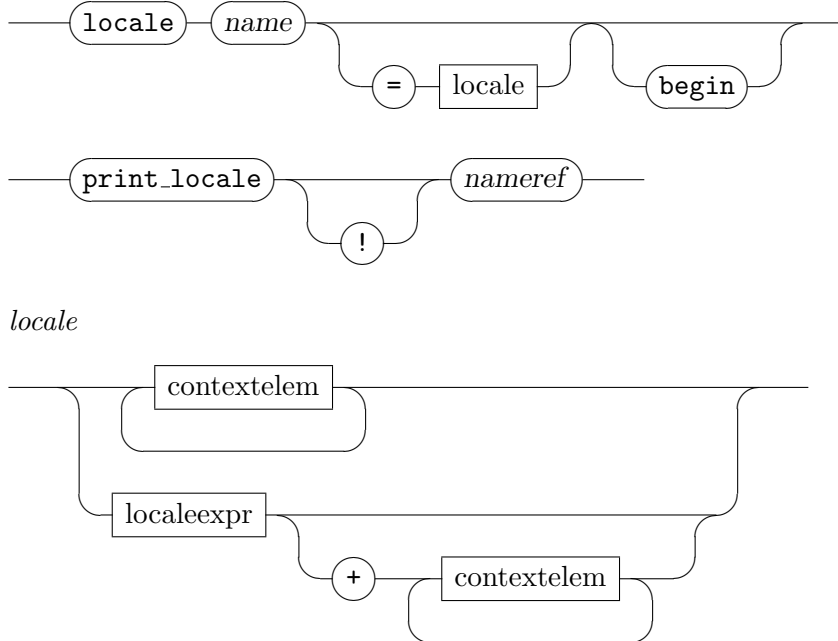
that instantiate a parameter by itself) may be omitted. The notation ‘_’ enables to omit the instantiation for a parameter inside a positional instantiation.

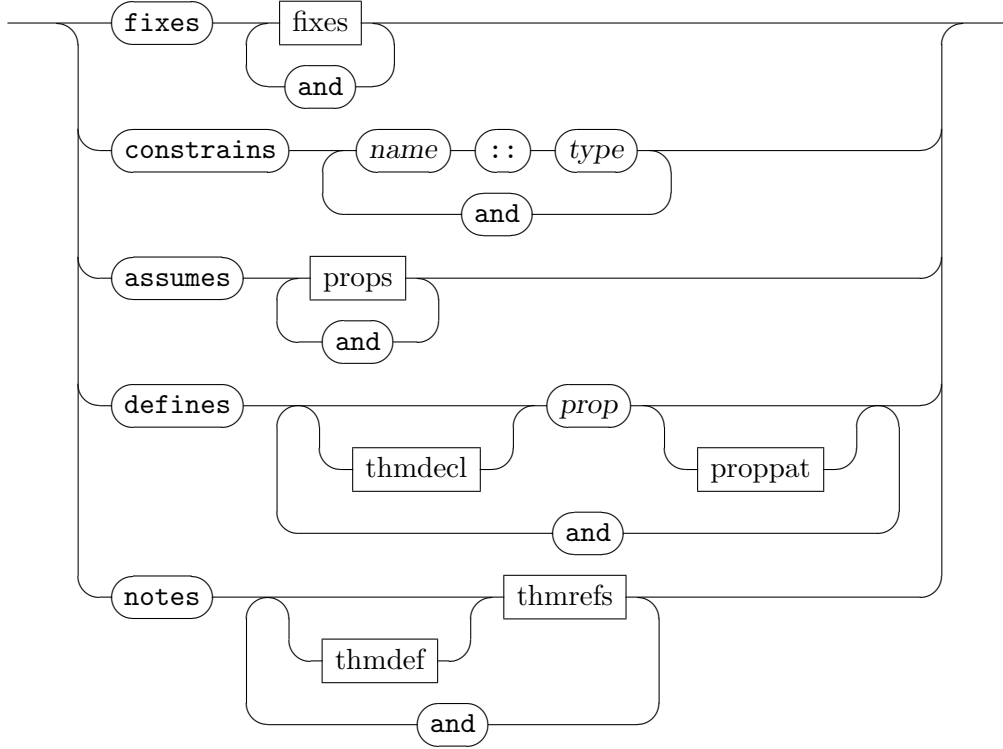
Terms in instantiations are from the context the locale expressions is declared in. Local names may be added to this context with the optional for clause. In addition, syntax declarations from one instance are effective when parsing subsequent instances of the same expression.

Instances have an optional qualifier which applies to names in declarations. Names include local definitions and theorem names. If present, the qualifier itself is either optional (“?”), which means that it may be omitted on input of the qualified name, or mandatory (“!”). If neither “?” nor “!” are present, the command’s default is used. For **interpretation** and **interpret** the default is “mandatory”, for **locale** and **sublocale** the default is “optional”.

5.5.2 Locale declarations

locale : *theory* \rightarrow *local_theory*
print_locale* : *context* \rightarrow
print_locales* : *context* \rightarrow
intro_locales : *method*
unfold_locales : *method*



contextelem

locale $loc = import + body$ defines a new locale loc as a context consisting of a certain view of existing locales ($import$) plus some additional elements ($body$). Both $import$ and $body$ are optional; the degenerate form **locale** loc defines an empty locale, which may still be useful to collect declarations of facts later on. Type-inference on locale expressions automatically takes care of the most general typing that the combined context elements may acquire.

The $import$ consists of a structured locale expression; see §6.2.1 above. Its for clause defines the local parameters of the $import$. In addition, locale parameters whose instance is omitted automatically extend the (possibly empty) for clause: they are inserted at its beginning. This means that these parameters may be referred to from within the expression and also in the subsequent context elements and provides a notational convenience for the inheritance of parameters in locale declarations.

The $body$ consists of context elements.

fixes $x :: \tau$ (mx) declares a local parameter of type τ and mixfix annotation mx (both are optional). The special syntax declaration

“(structure)” means that x may be referenced implicitly in this context.

constrains $x :: \tau$ introduces a type constraint τ on the local parameter x . This element is deprecated. The type constraint should be introduced in the for clause or the relevant **fixes** element.

assumes $a: \varphi_1 \dots \varphi_n$ introduces local premises, similar to **assume** within a proof (cf. §6.2.1).

defines $a: x \equiv t$ defines a previously declared parameter. This is similar to **def** within a proof (cf. §6.2.1), but **defines** takes an equational proposition instead of variable-term pair. The left-hand side of the equation may have additional arguments, e.g. “**defines** $f x_1 \dots x_n \equiv t$ ”.

notes $a = b_1 \dots b_n$ reconsiders facts within a local context. Most notably, this may include arbitrary declarations in any attribute specifications included here, e.g. a local *simp* rule.

The initial *import* specification of a locale expression maintains a dynamic relation to the locales being referenced (benefiting from any later fact declarations in the obvious manner).

Note that “(is $p_1 \dots p_n$)” patterns given in the syntax of **assumes** and **defines** above are illegal in locale definitions. In the long goal format of §6.2.4, term bindings may be included as expected, though.

Locale specifications are “closed up” by turning the given text into a predicate definition *loc_axioms* and deriving the original assumptions as local lemmas (modulo local definitions). The predicate statement covers only the newly specified assumptions, omitting the content of included locale expressions. The full cumulative view is only provided on export, involving another predicate *loc* that refers to the complete specification text.

In any case, the predicate arguments are those locale parameters that actually occur in the respective piece of text. Also note that these predicates operate at the meta-level in theory, but the locale packages attempts to internalize statements according to the object-logic setup (e.g. replacing \bigwedge by \forall , and \implies by \longrightarrow in HOL; see also §9.5). Separate introduction rules *loc_axioms.intro* and *loc.intro* are provided as well.

print_locale *locale* prints the contents of the named locale. The command omits **notes** elements by default. Use **print_locale!** to have them included.

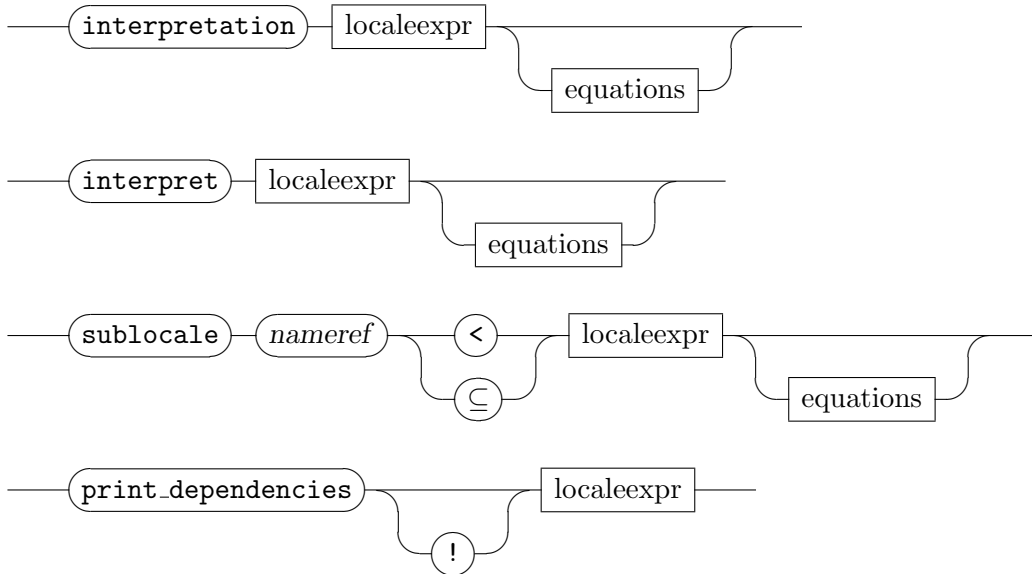
print_locales prints the names of all locales of the current theory.

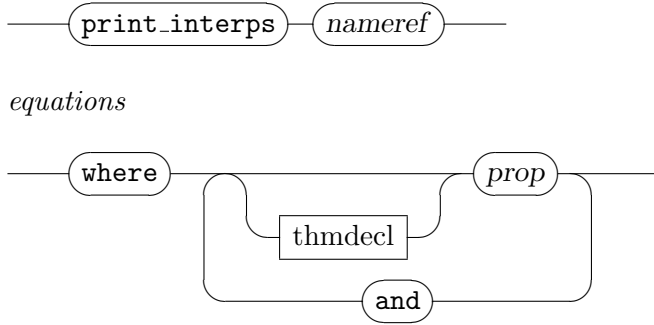
intro_locales and *unfold_locales* repeatedly expand all introduction rules of locale predicates of the theory. While *intro_locales* only applies the *loc.intro* introduction rules and therefore does not descend to assumptions, *unfold_locales* is more aggressive and applies *loc.axioms.intro* as well. Both methods are aware of locale specifications entailed by the context, both from target statements, and from interpretations (see below). New goals that are entailed by the current context are discharged automatically.

5.5.3 Locale interpretations

Locale expressions may be instantiated, and the instantiated facts added to the current context. This requires a proof of the instantiated specification and is called *locale interpretation*. Interpretation is possible in locales **sublocale**, theories (command **interpretation**) and also within a proof body (command **interpret**).

interpretation : $theory \rightarrow proof(prove)$
interpret : $proof(state) \mid proof(chain) \rightarrow proof(prove)$
sublocale : $theory \rightarrow proof(prove)$
print_dependencies* : $context \rightarrow$
print_interps* : $context \rightarrow$





interpretation *expr* **where** *eqns* interprets *expr* in the theory. The command generates proof obligations for the instantiated specifications (assumes and defines elements). Once these are discharged by the user, instantiated facts are added to the theory in a post-processing phase.

Additional equations, which are unfolded during post-processing, may be given after the keyword **where**. This is useful for interpreting concepts introduced through definitions. The equations must be proved.

The command is aware of interpretations already active in the theory, but does not simplify the goal automatically. In order to simplify the proof obligations use methods *intro_locales* or *unfold_locales*. Post-processing is not applied to facts of interpretations that are already active. This avoids duplication of interpreted facts, in particular. Note that, in the case of a locale with import, parts of the interpretation may already be active. The command will only process facts for new parts.

Adding facts to locales has the effect of adding interpreted facts to the theory for all interpretations as well. That is, interpretations dynamically participate in any facts added to locales. Note that if a theory inherits additional facts for a locale through one parent and an interpretation of that locale through another parent, the additional facts will not be interpreted.

interpret *expr* **where** *eqns* interprets *expr* in the proof context and is otherwise similar to interpretation in theories. Note that rewrite rules given to **interpret** after the **where** keyword should be explicitly universally quantified.

sublocale *name* \subseteq *expr* **where** *eqns* interprets *expr* in the locale *name*.

A proof that the specification of *name* implies the specification of *expr* is required. As in the localized version of the theorem command, the proof is in the context of *name*. After the proof obligation has been

discharged, the facts of *expr* become part of locale *name* as *derived* context elements and are available when the context *name* is subsequently entered. Note that, like `import`, this is dynamic: facts added to a locale part of *expr* after interpretation become also available in *name*.

Only specification fragments of *expr* that are not already part of *name* (be it imported, derived or a derived fragment of the import) are considered in this process. This enables circular interpretations provided that no infinite chains are generated in the locale hierarchy.

If interpretations of *name* exist in the current theory, the command adds interpretations for *expr* as well, with the same qualifier, although only for fragments of *expr* that are not interpreted in the theory already.

Equations given after **where** amend the morphism through which *expr* is interpreted. This enables to map definitions from the interpreted locales to entities of *name*. This feature is experimental.

print_dependencies *expr* is useful for understanding the effect of an interpretation of *expr*. It lists all locale instances for which interpretations would be added to the current context. Variant **print_dependencies!** prints all locale instances that would be considered for interpretation, and would be interpreted in an empty context (that is, without interpretations).

print_interps *locale* lists all interpretations of *locale* in the current theory or proof context, including those due to a combination of a **interpretation** or **interpret** and one or several **sublocale** declarations.

! Since attributes are applied to interpreted theorems, interpretation may modify

- the context of common proof tools, e.g. the Simplifier or Classical Reasoner.

As the behavior of such tools is *not* stable under interpretation morphisms, manual declarations might have to be added to the target context of the interpretation to revert such declarations.

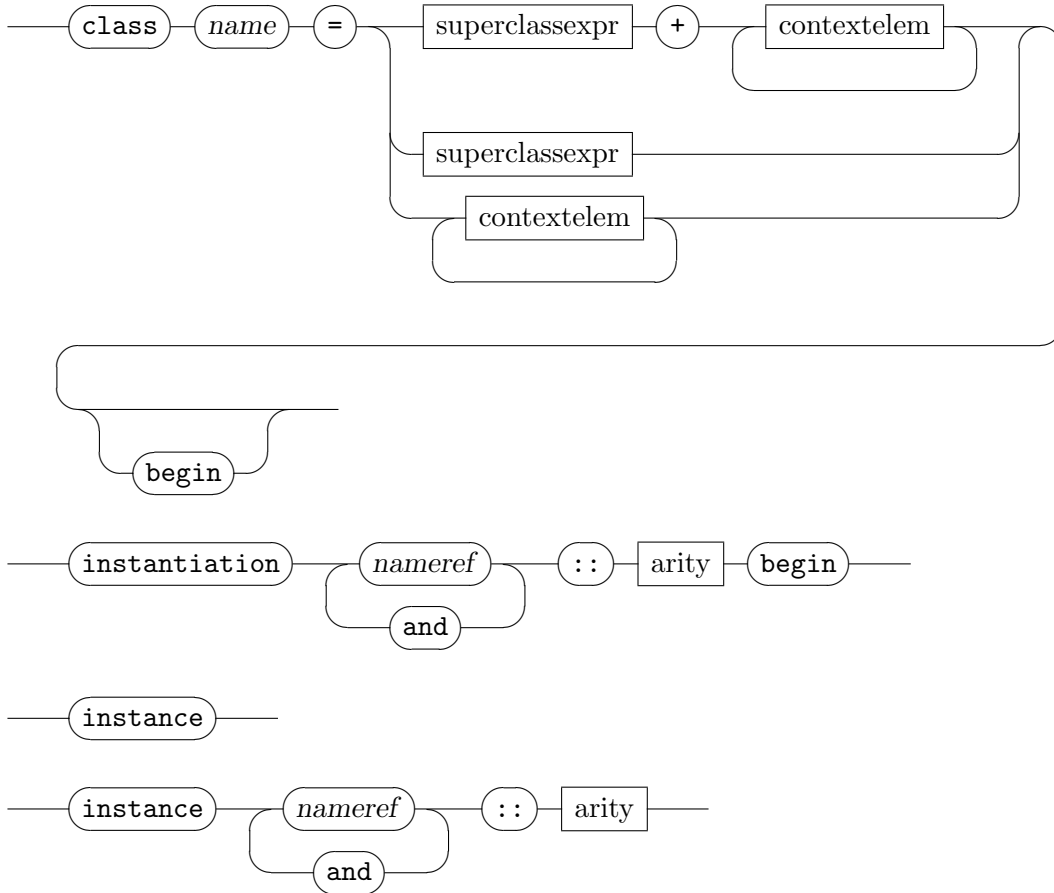
! An interpretation in a theory or proof context may subsume previous interpretations.

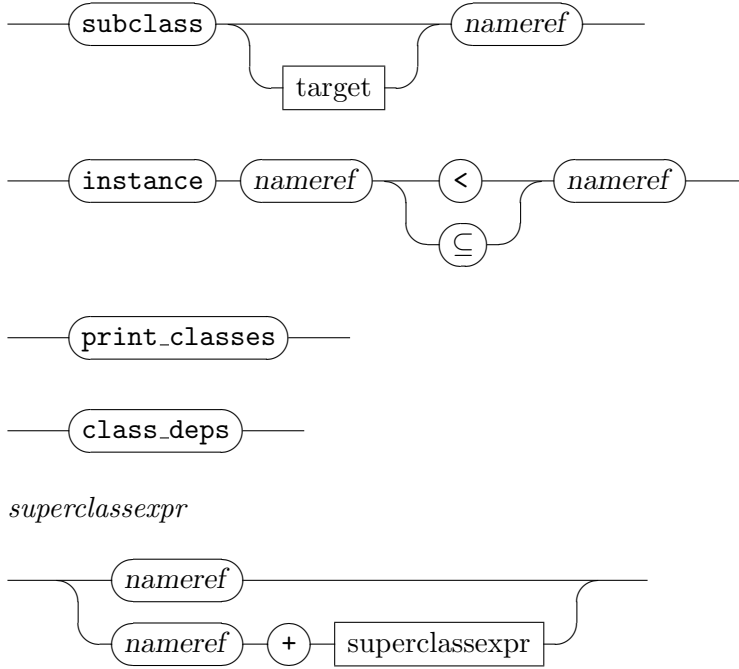
- This happens if the same specification fragment is interpreted twice and the instantiation of the second interpretation is more general than the interpretation of the first. The locale package does not attempt to remove subsumed interpretations.

5.6 Classes

A class is a particular locale with *exactly one* type variable α . Beyond the underlying locale, a corresponding type class is established which is interpreted logically as axiomatic type class [35] whose logical content are the assumptions of the locale. Thus, classes provide the full generality of locales combined with the commodity of type classes (notably type-inference). See [9] for a short tutorial.

class : $theory \rightarrow local_theory$
instantiation : $theory \rightarrow local_theory$
instance : $local_theory \rightarrow local_theory$
instance : $theory \rightarrow proof(prove)$
subclass : $local_theory \rightarrow local_theory$
print_classes* : $context \rightarrow$
class_deps* : $context \rightarrow$
intro_classes : *method*





class $c = \text{superclasses} + \text{body}$ defines a new class c , inheriting from *superclasses*. This introduces a locale c with import of all locales *superclasses*.

Any **fixes** in *body* are lifted to the global theory level (*class operations* f_1, \dots, f_n of class c), mapping the local type parameter α to a schematic type variable $?\alpha :: c$.

Likewise, **assumes** in *body* are also lifted, mapping each local parameter $f :: \tau[\alpha]$ to its corresponding global constant $f :: \tau[?\alpha :: c]$. The corresponding introduction rule is provided as *c-class-axioms.intro*. This rule should be rarely needed directly — the *intro_classes* method takes care of the details of class membership proofs.

instantiation $t :: (s_1, \dots, s_n)s$ **begin** opens a theory target (cf. §5.2) which allows to specify class operations f_1, \dots, f_n corresponding to sort s at the particular type instance $(\alpha_1 :: s_1, \dots, \alpha_n :: s_n) t$. A plain **instance** command in the target body poses a goal stating these type arities. The target is concluded by an **end** command.

Note that a list of simultaneous type constructors may be given; this corresponds nicely to mutually recursive type definitions, e.g. in Isabelle/HOL.

instance in an instantiation target body sets up a goal stating the type arities claimed at the opening **instantiation**. The proof would usually proceed by *intro_classes*, and then establish the characteristic theorems of the type classes involved. After finishing the proof, the background theory will be augmented by the proven type arities.

On the theory level, **instance** $t :: (s_1, \dots, s_n)s$ provides a convenient way to instantiate a type class with no need to specify operations: one can continue with the instantiation proof immediately.

subclass c in a class context for class d sets up a goal stating that class c is logically contained in class d . After finishing the proof, class d is proven to be subclass c and the locale c is interpreted into d simultaneously.

A weakend form of this is available through a further variant of **instance**: **instance** $c_1 \subseteq c_2$ opens a proof that class c_2 implies c_1 without reference to the underlying locales; this is useful if the properties to prove the logical connection are not sufficient on the locale level but on the theory level.

print_classes prints all classes in the current theory.

class_deps visualizes all classes and their subclass relations as a Hasse diagram.

intro_classes repeatedly expands all class introduction rules of this theory.

Note that this method usually needs not be named explicitly, as it is already included in the default proof step (e.g. of **proof**). In particular, instantiation of trivial (syntactic) classes may be performed by a single “..” proof step.

5.6.1 The class target

A named context may refer to a locale (cf. §5.2). If this locale is also a class c , apart from the common locale target behaviour the following happens.

- Local constant declarations $g[\alpha]$ referring to the local type parameter α and local parameters $f[\alpha]$ are accompanied by theory-level constants $g[? \alpha :: c]$ referring to theory-level class operations $f[? \alpha :: c]$.
- Local theorem bindings are lifted as are assumptions.

- Local syntax refers to local operations $g[\alpha]$ and global operations $g[? \alpha :: c]$ uniformly. Type inference resolves ambiguities. In rare cases, manual type annotations are needed.

5.6.2 Co-regularity of type classes and arities

The class relation together with the collection of type-constructor arities must obey the principle of *co-regularity* as defined below.

For the subsequent formulation of co-regularity we assume that the class relation is closed by transitivity and reflexivity. Moreover the collection of arities $t :: (\bar{s})c$ is completed such that $t :: (\bar{s})c$ and $c \subseteq c'$ implies $t :: (\bar{s})c'$ for all such declarations.

Treating sorts as finite sets of classes (meaning the intersection), the class relation $c_1 \subseteq c_2$ is extended to sorts as follows:

$$s_1 \subseteq s_2 \equiv \forall c_2 \in s_2. \exists c_1 \in s_1. c_1 \subseteq c_2$$

This relation on sorts is further extended to tuples of sorts (of the same length) in the component-wise way.

Co-regularity of the class relation together with the arities relation means:

$$t :: (\bar{s}_1)c_1 \implies t :: (\bar{s}_2)c_2 \implies c_1 \subseteq c_2 \implies \bar{s}_1 \subseteq \bar{s}_2$$

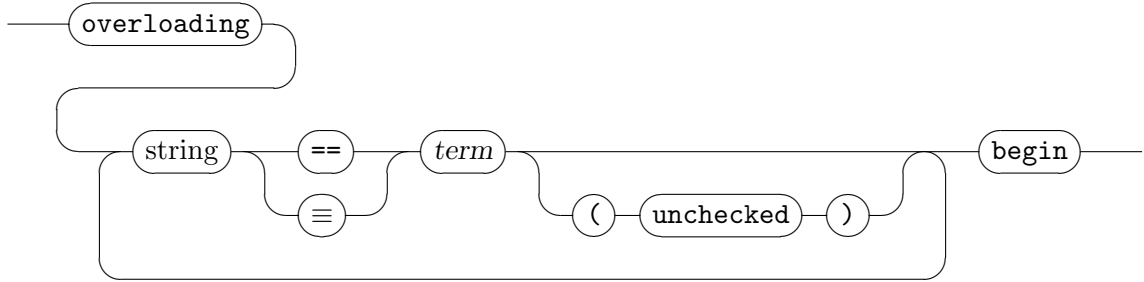
for all such arities. In other words, whenever the result classes of some type-constructor arities are related, then the argument sorts need to be related in the same way.

Co-regularity is a very fundamental property of the order-sorted algebra of types. For example, it entails principle types and most general unifiers, e.g. see [21].

5.7 Unrestricted overloading

Isabelle/Pure's definitional schemes support certain forms of overloading (see §5.9.3). Overloading means that a constant being declared as $c :: \alpha \text{ decl}$ may be defined separately on type instances $c :: (\beta_1, \dots, \beta_n) t \text{ decl}$ for each type constructor t . At most occasions overloading will be used in a Haskell-like fashion together with type classes by means of **instantiation** (see §5.6). Sometimes low-level overloading is desirable. The **overloading** target provides a convenient view for end-users.

$$\mathbf{overloading} : theory \rightarrow local_theory$$

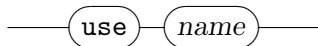


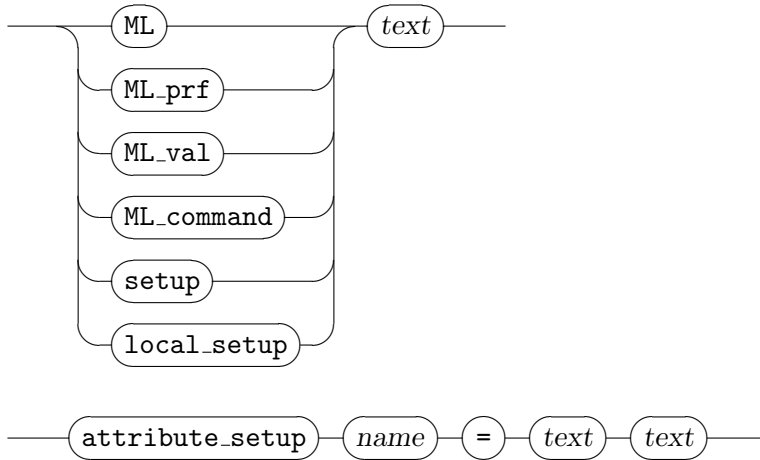
overloading $x_1 \equiv c_1 :: \tau_1$ **and** $\dots x_n \equiv c_n :: \tau_n$ **begin** opens a theory target (cf. §5.2) which allows to specify constants with overloaded definitions. These are identified by an explicitly given mapping from variable names x_i to constants c_i at particular type instances. The definitions themselves are established using common specification tools, using the names x_i as reference to the corresponding constants. The target is concluded by **end**.

A (*unchecked*) option disables global dependency checks for the corresponding definition, which is occasionally useful for exotic overloading (see §5.9.3 for a precise description). It is at the discretion of the user to avoid malformed theory specifications!

5.8 Incorporating ML code

use : $local_theory \rightarrow local_theory$
ML : $local_theory \rightarrow local_theory$
ML_prf : $proof \rightarrow proof$
ML_val : $any \rightarrow$
ML_command : $any \rightarrow$
setup : $theory \rightarrow theory$
local_setup : $local_theory \rightarrow local_theory$
attribute_setup : $theory \rightarrow theory$





use *file* reads and executes ML commands from *file*. The current theory context is passed down to the ML toplevel and may be modified, using `Context.>>` or derived ML commands. The file name is checked with the **uses** dependency declaration given in the theory header (see also §5.1).

Top-level ML bindings are stored within the (global or local) theory context.

ML *text* is similar to **use**, but executes ML commands directly from the given *text*. Top-level ML bindings are stored within the (global or local) theory context.

ML_prf is analogous to **ML** but works within a proof context.

Top-level ML bindings are stored within the proof context in a purely sequential fashion, disregarding the nested proof structure. ML bindings introduced by **ML_prf** are discarded at the end of the proof.

ML_val and **ML_command** are diagnostic versions of **ML**, which means that the context may not be updated. **ML_val** echos the bindings produced at the ML toplevel, but **ML_command** is silent.

setup *text* changes the current theory context by applying *text*, which refers to an ML expression of type `theory -> theory`. This enables to initialize any object-logic specific tools and packages written in ML, for example.

local_setup is similar to **setup** for a local theory context, and an ML expression of type `local_theory -> local_theory`. This allows to

invoke local theory specification packages without going through concrete outer syntax, for example.

attribute_setup *name* = *text description* defines an attribute in the current theory. The given *text* has to be an ML expression of type **attribute context_parser**, cf. basic parsers defined in structure **Args** and **Attrib**.

In principle, attributes can operate both on a given theorem and the implicit context, although in practice only one is modified and the other serves as parameter. Here are examples for these two cases:

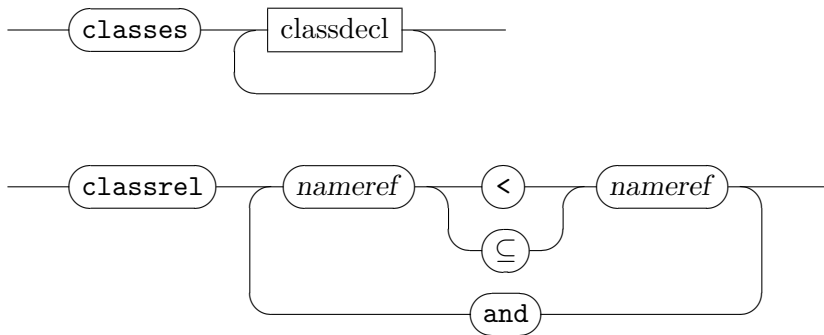
```
attribute_setup my_rule = ⟨⟨
  Attrib.thms >> (fn ths =>
    Thm.rule_attribute (fn context: Context.generic => fn th: thm =>
      let val th' = th OF ths
      in th' end))) ⟩ my_rule
```

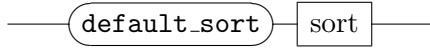
```
attribute_setup my_declaration = ⟨⟨
  Attrib.thms >> (fn ths =>
    Thm.declaration_attribute (fn th: thm => fn context: Context.generic =>
      let val context' = context
      in context' end))) ⟩ my_declaration
```

5.9 Primitive specification elements

5.9.1 Type classes and sorts

classes : *theory* → *theory*
classrel : *theory* → *theory* (axiomatic!)
default_sort : *local_theory* → *local_theory*





classes $c \subseteq c_1, \dots, c_n$ declares class c to be a subclass of existing classes c_1, \dots, c_n . Isabelle implicitly maintains the transitive closure of the class hierarchy. Cyclic class structures are not permitted.

classrel $c_1 \subseteq c_2$ states subclass relations between existing classes c_1 and c_2 . This is done axiomatically! The **subclass** and **instance** commands (see §5.6) provide a way to introduce proven class relations.

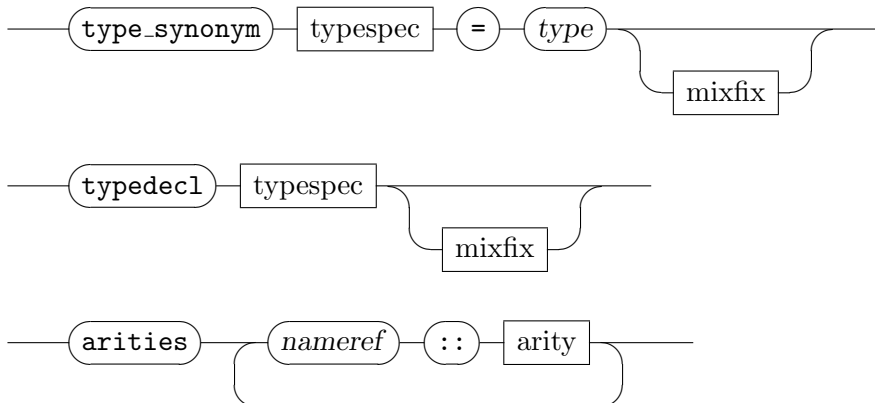
default_sort s makes sort s the new default sort for any type variable that is given explicitly in the text, but lacks a sort constraint (wrt. the current context). Type variables generated by type inference are not affected.

Usually the default sort is only changed when defining a new object-logic. For example, the default sort in Isabelle/HOL is *type*, the class of all HOL types.

When merging theories, the default sorts of the parents are logically intersected, i.e. the representations as lists of classes are joined.

5.9.2 Types and type abbreviations

type_synonym : $local_theory \rightarrow local_theory$
typedecl : $local_theory \rightarrow local_theory$
arities : $theory \rightarrow theory$ (axiomatic!)



type_synonym $(\alpha_1, \dots, \alpha_n) \ t = \tau$ introduces a *type synonym* $(\alpha_1, \dots, \alpha_n) \ t$ for the existing type τ . Unlike actual type definitions, as are available in Isabelle/HOL for example, type synonyms are merely syntactic abbreviations without any logical significance. Internally, type synonyms are fully expanded.

typeddecl $(\alpha_1, \dots, \alpha_n) \ t$ declares a new type constructor t . If the object-logic defines a base sort s , then the constructor is declared to operate on that, via the axiomatic specification **arities** $t :: (s, \dots, s)s$.

arities $t :: (s_1, \dots, s_n)s$ augments Isabelle’s order-sorted signature of types by new type constructor arities. This is done axiomatically! The **instantiation** target (see §5.6) provides a way to introduce proven type arities.

5.9.3 Constants and definitions

Definitions essentially express abbreviations within the logic. The simplest form of a definition is $c :: \sigma \equiv t$, where c is a newly declared constant. Isabelle also allows derived forms where the arguments of c appear on the left, abbreviating a prefix of λ -abstractions, e.g. $c \equiv \lambda x \ y. \ t$ may be written more conveniently as $c \ x \ y \equiv t$. Moreover, definitions may be weakened by adding arbitrary pre-conditions: $A \implies c \ x \ y \equiv t$.

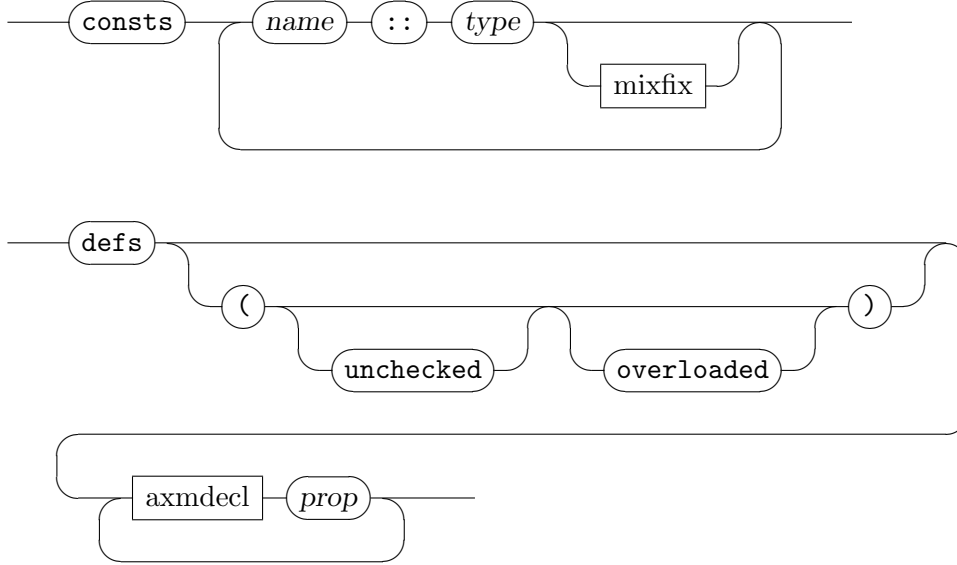
The built-in well-formedness conditions for definitional specifications are:

- Arguments (on the left-hand side) must be distinct variables.
- All variables on the right-hand side must also appear on the left-hand side.
- All type variables on the right-hand side must also appear on the left-hand side; this prohibits $0 :: nat \equiv length \ (\square :: \alpha \ list)$ for example.
- The definition must not be recursive. Most object-logics provide definitional principles that can be used to express recursion safely.

The right-hand side of overloaded definitions may mention overloaded constants recursively at type instances corresponding to the immediate argument types β_1, \dots, β_n . Incomplete specification patterns impose global constraints on all occurrences, e.g. $d :: \alpha \times \alpha$ on the left-hand side means

that all corresponding occurrences on some right-hand side need to be an instance of this, general $d :: \alpha \times \beta$ will be disallowed.

consts : $theory \rightarrow theory$
defs : $theory \rightarrow theory$



consts $c :: \sigma$ declares constant c to have any instance of type scheme σ . The optional mixfix annotations may attach concrete syntax to the constants declared.

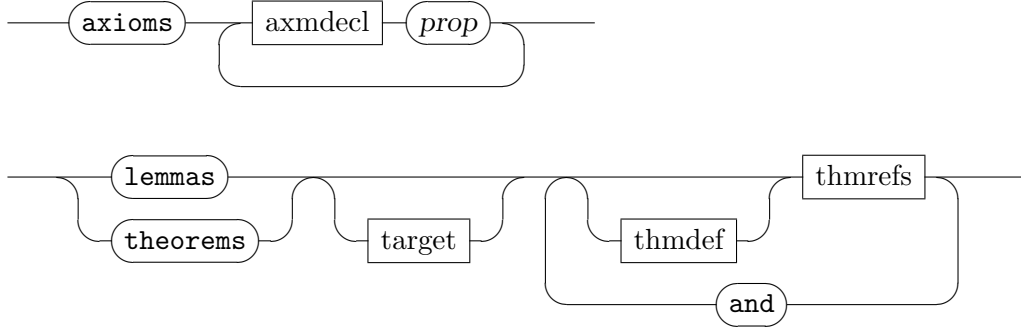
defs $name: eqn$ introduces eqn as a definitional axiom for some existing constant.

The (*unchecked*) option disables global dependency checks for this definition, which is occasionally useful for exotic overloading. It is at the discretion of the user to avoid malformed theory specifications!

The (*overloaded*) option declares definitions to be potentially overloaded. Unless this option is given, a warning message would be issued for any definitional equation with a more special type than that of the corresponding constant declaration.

5.10 Axioms and theorems

axioms : $theory \rightarrow theory$ (axiomatic!)
lemmas : $local_theory \rightarrow local_theory$
theorems : $local_theory \rightarrow local_theory$



axioms a : φ introduces arbitrary statements as axioms of the meta-logic. In fact, axioms are “axiomatic theorems”, and may be referred later just as any other theorem.

Axioms are usually only introduced when declaring new logical systems. Everyday work is typically done the hard way, with proper definitions and proven theorems.

lemmas $a = b_1 \dots b_n$ retrieves and stores existing facts in the theory context, or the specified target context (see also §5.2). Typical applications would also involve attributes, to declare Simplifier rules, for example.

theorems is essentially the same as **lemmas**, but marks the result as a different kind of facts.

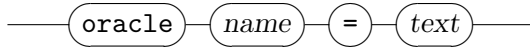
5.11 Oracles

Oracles allow Isabelle to take advantage of external reasoners such as arithmetic decision procedures, model checkers, fast tautology checkers or computer algebra systems. Invoked as an oracle, an external reasoner can create arbitrary Isabelle theorems.

It is the responsibility of the user to ensure that the external reasoner is as trustworthy as the application requires. Another typical source of errors is the linkup between Isabelle and the external tool, not just its concrete implementation, but also the required translation between two different logical environments.

Isabelle merely guarantees well-formedness of the propositions being asserted, and records within the internal derivation object how presumed theorems depend on unproven suppositions.

oracle : *theory* \rightarrow *theory* (*axiomatic!*)

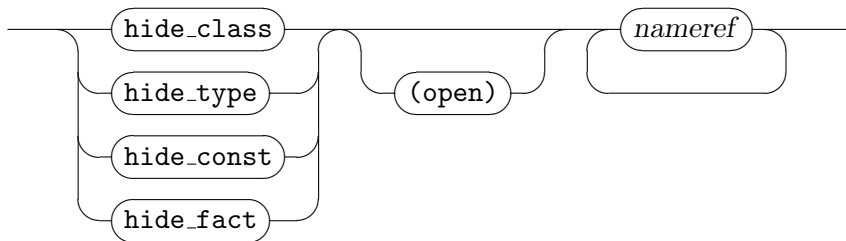


oracle *name* = *text* turns the given ML expression *text* of type '*a* \rightarrow *cterm*' into an ML function of type '*a* \rightarrow *thm*', which is bound to the global identifier **name**. This acts like an infinitary specification of axioms! Invoking the oracle only works within the scope of the resulting theory.

See `~/src/HOL/ex/Iff_Oracle.thy` for a worked example of defining a new primitive rule as oracle, and turning it into a proof method.

5.12 Name spaces

hide_class : *theory* \rightarrow *theory*
hide_type : *theory* \rightarrow *theory*
hide_const : *theory* \rightarrow *theory*
hide_fact : *theory* \rightarrow *theory*



Isabelle organizes any kind of name declarations (of types, constants, theorems etc.) by separate hierarchically structured name spaces. Normally the user does not have to control the behavior of name spaces by hand, yet the following commands provide some way to do so.

hide_class *names* fully removes class declarations from a given name space; with the (*open*) option, only the base name is hidden.

Note that hiding name space accesses has no impact on logical declarations — they remain valid internally. Entities that are no longer accessible to the user are printed with the special qualifier “??” prefixed to the full internal name.

hide_type, **hide_const**, and **hide_fact** are similar to **hide_class**, but hide types, constants, and facts, respectively.

Proofs

Proof commands perform transitions of Isar/VM machine configurations, which are block-structured, consisting of a stack of nodes with three main components: logical proof context, current facts, and open goals. Isar/VM transitions are typed according to the following three different modes of operation:

proof(prove) means that a new goal has just been stated that is now to be *proven*; the next command may refine it by some proof method, and enter a sub-proof to establish the actual result.

proof(state) is like a nested theory mode: the context may be augmented by *stating* additional assumptions, intermediate results etc.

proof(chain) is intermediate between *proof(state)* and *proof(prove)*: existing facts (i.e. the contents of the special “*this*” register) have been just picked up in order to be used when refining the goal claimed next.

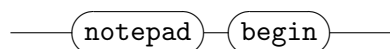
The proof mode indicator may be understood as an instruction to the writer, telling what kind of operation may be performed next. The corresponding typings of proof commands restricts the shape of well-formed proof texts to particular command sequences. So dynamic arrangements of commands eventually turn out as static texts of a certain structure.

Appendix A gives a simplified grammar of the (extensible) language emerging that way from the different types of proof commands. The main ideas of the overall Isar framework are explained in chapter 2.

6.1 Proof structure

6.1.1 Formal notepad

notepad : *local_theory* \rightarrow *proof(state)*



— (end) —

notepad begin opens a proof state without any goal statement. This allows to experiment with Isar, without producing any persistent result.

The notepad can be closed by **end** or discontinued by **oops**.

6.1.2 Blocks

$$\begin{array}{ll} \mathbf{next} & : \text{proof}(state) \rightarrow \text{proof}(state) \\ \{ & : \text{proof}(state) \rightarrow \text{proof}(state) \\ \} & : \text{proof}(state) \rightarrow \text{proof}(state) \end{array}$$

While Isar is inherently block-structured, opening and closing blocks is mostly handled rather casually, with little explicit user-intervention. Any local goal statement automatically opens *two* internal blocks, which are closed again when concluding the sub-proof (by **qed** etc.). Sections of different context within a sub-proof may be switched via **next**, which is just a single block-close followed by block-open again. The effect of **next** is to reset the local proof context; there is no goal focus involved here!

For slightly more advanced applications, there are explicit block parentheses as well. These typically achieve a stronger forward style of reasoning.

next switches to a fresh block within a sub-proof, resetting the local context to the initial one.

{ and **}** explicitly open and close blocks. Any current facts pass through “**{**” unchanged, while “**}**” causes any result to be *exported* into the enclosing context. Thus fixed variables are generalized, assumptions discharged, and local definitions unfolded (cf. §6.2.1). There is no difference of **assume** and **presume** in this mode of forward reasoning — in contrast to plain backward reasoning with the result exported at **show** time.

6.1.3 Omitting proofs

$$\mathbf{oops} : \text{proof} \rightarrow \text{local_theory} \mid \text{theory}$$

The **oops** command discontinues the current proof attempt, while considering the partial proof text as properly processed. This is conceptually quite different from “faking” actual proofs via **sorry** (see §6.3.2): **oops** does

not observe the proof structure at all, but goes back right to the theory level. Furthermore, **oops** does not produce any result theorem — there is no intended claim to be able to complete the proof anyhow.

A typical application of **oops** is to explain Isar proofs *within* the system itself, in conjunction with the document preparation tools of Isabelle described in chapter 4. Thus partial or even wrong proof attempts can be discussed in a logically sound manner. Note that the Isabelle \LaTeX macros can be easily adapted to print something like “...” instead of the keyword “oops”.

The **oops** command is undo-able, unlike **kill** (see §8.2). The effect is to get back to the theory just before the opening of the proof.

6.2 Statements

6.2.1 Context elements

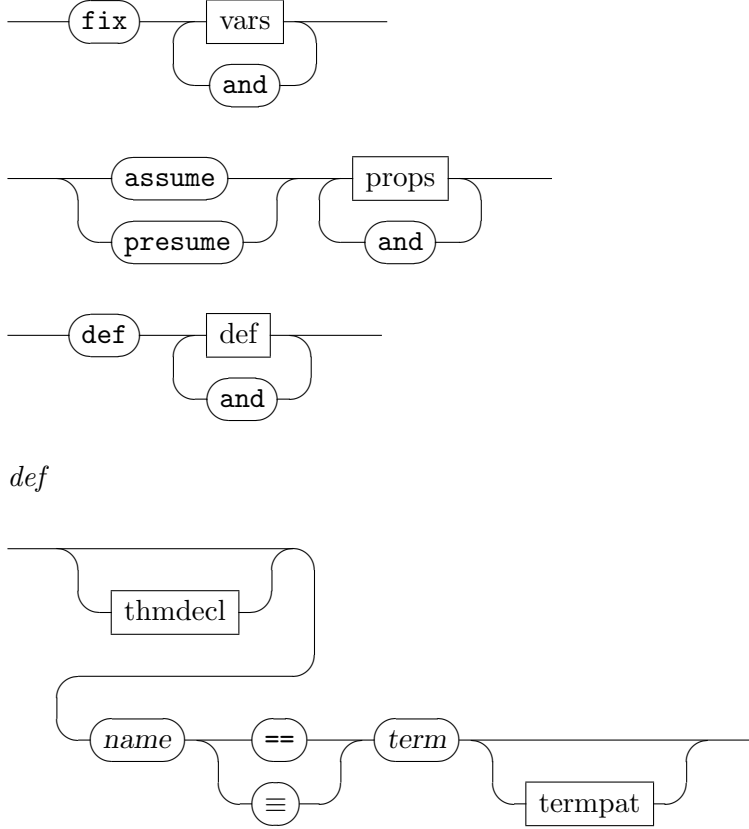
$$\begin{array}{ll} \mathbf{fix} & : \text{proof}(state) \rightarrow \text{proof}(state) \\ \mathbf{assume} & : \text{proof}(state) \rightarrow \text{proof}(state) \\ \mathbf{presume} & : \text{proof}(state) \rightarrow \text{proof}(state) \\ \mathbf{def} & : \text{proof}(state) \rightarrow \text{proof}(state) \end{array}$$

The logical proof context consists of fixed variables and assumptions. The former closely correspond to Skolem constants, or meta-level universal quantification as provided by the Isabelle/Pure logical framework. Introducing some *arbitrary, but fixed* variable via “**fix** x ” results in a local value that may be used in the subsequent proof as any other variable or constant. Furthermore, any result $\vdash \varphi[x]$ exported from the context will be universally closed wrt. x at the outermost level: $\vdash \bigwedge x. \varphi[x]$ (this is expressed in normal form using Isabelle’s meta-variables).

Similarly, introducing some assumption χ has two effects. On the one hand, a local theorem is created that may be used as a fact in subsequent proof steps. On the other hand, any result $\chi \vdash \varphi$ exported from the context becomes conditional wrt. the assumption: $\vdash \chi \implies \varphi$. Thus, solving an enclosing goal using such a result would basically introduce a new subgoal stemming from the assumption. How this situation is handled depends on the version of assumption command used: while **assume** insists on solving the subgoal by unification with some premise of the goal, **presume** leaves the subgoal unchanged in order to be proved later by the user.

Local definitions, introduced by “**def** $x \equiv t$ ”, are achieved by combining “**fix** x ” with another version of assumption that causes any hypothetical

equation $x \equiv t$ to be eliminated by the reflexivity rule. Thus, exporting some result $x \equiv t \vdash \varphi[x]$ yields $\vdash \varphi[t]$.



fix x introduces a local variable x that is *arbitrary, but fixed*.

assume $a: \varphi$ and **presume** $a: \varphi$ introduce a local fact $\varphi \vdash \varphi$ by assumption.

Subsequent results applied to an enclosing goal (e.g. by **show**) are handled as follows: **assume** expects to be able to unify with existing premises in the goal, while **presume** leaves φ as new subgoals.

Several lists of assumptions may be given (separated by **and**; the resulting list of current facts consists of all of these concatenated.

def $x \equiv t$ introduces a local (non-polymorphic) definition. In results exported from the context, x is replaced by t . Basically, “**def** $x \equiv t$ ” abbreviates “**fix** x **assume** $x \equiv t$ ”, with the resulting hypothetical equation solved by reflexivity.

The default name for the definitional equation is x_def . Several simultaneous definitions may be given at the same time.

The special name *prems* refers to all assumptions of the current context as a list of theorems. This feature should be used with great care! It is better avoided in final proof texts.

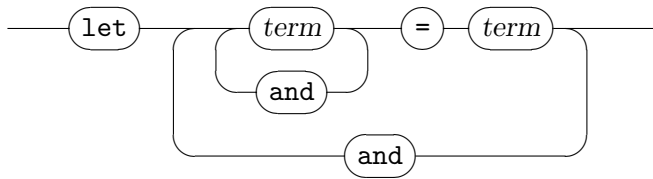
6.2.2 Term abbreviations

let : $proof(state) \rightarrow proof(state)$
is : *syntax*

Abbreviations may be either bound by explicit **let** $p \equiv t$ statements, or by annotating assumptions or goal statements with a list of patterns “(**is** $p_1 \dots p_n$)”. In both cases, higher-order matching is invoked to bind extra-logical term variables, which may be either named schematic variables of the form $?x$, or nameless dummies “_” (underscore). Note that in the **let** form the patterns occur on the left-hand side, while the **is** patterns are in postfix position.

Polymorphism of term bindings is handled in Hindley-Milner style, similar to ML. Type variables referring to local assumptions or open goal statements are *fixed*, while those of finished results or bound by **let** may occur in *arbitrary* instances later. Even though actual polymorphism should be rarely used in practice, this mechanism is essential to achieve proper incremental type-inference, as the user proceeds to build up the Isar proof text from left to right.

Term abbreviations are quite different from local definitions as introduced via **def** (see §6.2.1). The latter are visible within the logic as actual equations, while abbreviations disappear during the input process just after type checking. Also note that **def** does not support polymorphism.



The syntax of **is** patterns follows *termpat* or *proppat* (see §3.2.6).

let $p_1 = t_1$ **and** \dots $p_n = t_n$ binds any text variables in patterns p_1, \dots, p_n by simultaneous higher-order matching against terms t_1, \dots, t_n .

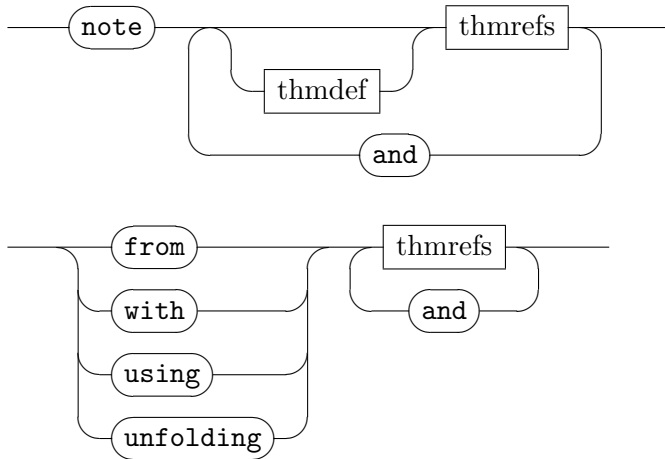
(**is** $p_1 \dots p_n$) resembles **let**, but matches p_1, \dots, p_n against the preceding statement. Also note that **is** is not a separate command, but part of others (such as **assume**, **have** etc.).

Some *implicit* term abbreviations for goals and facts are available as well. For any open goal, *thesis* refers to its object-level statement, abstracted over any meta-level parameters (if present). Likewise, *this* is bound for fact statements resulting from assumptions or finished goals. In case *this* refers to an object-logic statement that is an application $f\ t$, then t is bound to the special text variable “...” (three dots). The canonical application of this convenience are calculational proofs (see §6.5).

6.2.3 Facts and forward chaining

note : $proof(state) \rightarrow proof(state)$
then : $proof(state) \rightarrow proof(chain)$
from : $proof(state) \rightarrow proof(chain)$
with : $proof(state) \rightarrow proof(chain)$
using : $proof(prove) \rightarrow proof(prove)$
unfolding : $proof(prove) \rightarrow proof(prove)$

New facts are established either by assumption or proof of local statements. Any fact will usually be involved in further proofs, either as explicit arguments of proof methods, or when forward chaining towards the next goal via **then** (and variants); **from** and **with** are composite forms involving **note**. The **using** elements augments the collection of used facts *after* a goal has been stated. Note that the special theorem name *this* refers to the most recently established facts, but only *before* issuing a follow-up claim.



note $a = b_1 \dots b_n$ recalls existing facts b_1, \dots, b_n , binding the result as a . Note that attributes may be involved as well, both on the left and right hand sides.

then indicates forward chaining by the current facts in order to establish the goal to be claimed next. The initial proof method invoked to refine that will be offered the facts to do “anything appropriate” (see also §6.3.2). For example, method *rule* (see §6.3.3) would typically do an elimination rather than an introduction. Automatic methods usually insert the facts into the goal state before operation. This provides a simple scheme to control relevance of facts in automated proof search.

from b abbreviates “**note** b **then**”; thus **then** is equivalent to “**from** *this*”.

with $b_1 \dots b_n$ abbreviates “**from** $b_1 \dots b_n$ **and** *this*”; thus the forward chaining is from earlier facts together with the current ones.

using $b_1 \dots b_n$ augments the facts being currently indicated for use by a subsequent refinement step (such as **apply** or **proof**).

unfolding $b_1 \dots b_n$ is structurally similar to **using**, but unfolds definitional equations $b_1, \dots b_n$ throughout the goal state and facts.

Forward chaining with an empty list of theorems is the same as not chaining at all. Thus “**from** *nothing*” has no effect apart from entering *prove(chain)* mode, since *nothing* is bound to the empty list of theorems.

Basic proof methods (such as *rule*) expect multiple facts to be given in their proper order, corresponding to a prefix of the premises of the rule involved. Note that positions may be easily skipped using something like **from** $_$ **and** a **and** b , for example. This involves the trivial rule $PROP \psi \implies PROP \psi$, which is bound in Isabelle/Pure as “ $_$ ” (underscore).

Automated methods (such as *simp* or *auto*) just insert any given facts before their usual operation. Depending on the kind of procedure involved, the order of facts is less significant here.

6.2.4 Goals

```

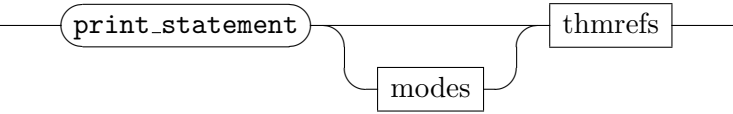
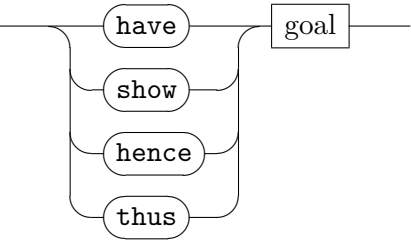
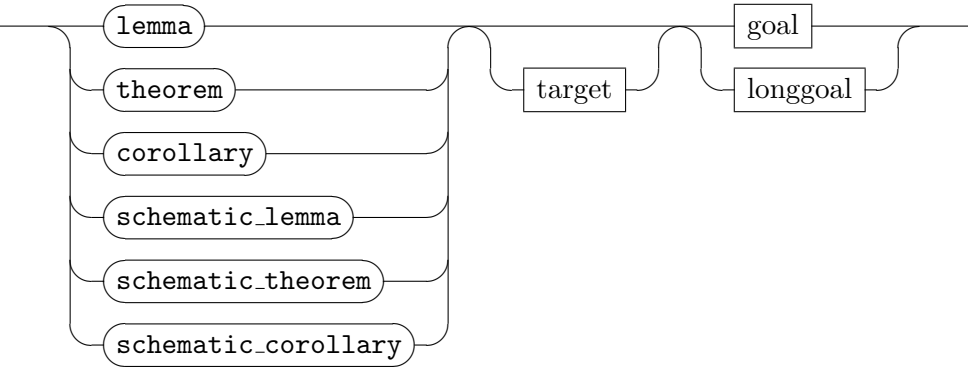
lemma      : local_theory → proof(prove)
theorem    : local_theory → proof(prove)
corollary  : local_theory → proof(prove)
schematic_lemma : local_theory → proof(prove)
schematic_theorem : local_theory → proof(prove)
schematic_corollary : local_theory → proof(prove)
have      : proof(state) | proof(chain) → proof(prove)
show      : proof(state) | proof(chain) → proof(prove)
hence     : proof(state) → proof(prove)
thus      : proof(state) → proof(prove)
print_statement* : context →

```

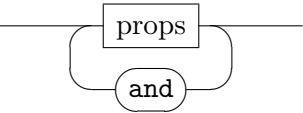
From a theory context, proof mode is entered by an initial goal command such as **lemma**, **theorem**, or **corollary**. Within a proof, new claims may be introduced locally as well; four variants are available here to indicate whether forward chaining of facts should be performed initially (via **then**), and whether the final result is meant to solve some pending goal.

Goals may consist of multiple statements, resulting in a list of facts eventually. A pending multi-goal is internally represented as a meta-level conjunction (&&&), which is usually split into the corresponding number of sub-goals prior to an initial method application, via **proof** (§6.3.2) or **apply** (§6.3.4). The *induct* method covered in §6.6 acts on multiple claims simultaneously.

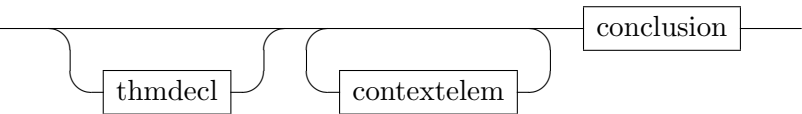
Claims at the theory level may be either in short or long form. A short goal merely consists of several simultaneous propositions (often just one). A long goal includes an explicit context specification for the subsequent conclusion, involving local parameters and assumptions. Here the role of each part of the statement is explicitly marked by separate keywords (see also §5.5); the local assumptions being introduced here are available as *assms* in the proof. Moreover, there are two kinds of conclusions: **shows** states several simultaneous propositions (essentially a big conjunction), while **obtains** claims several simultaneous contexts of (essentially a big disjunction of eliminated parameters and assumptions, cf. §6.4).



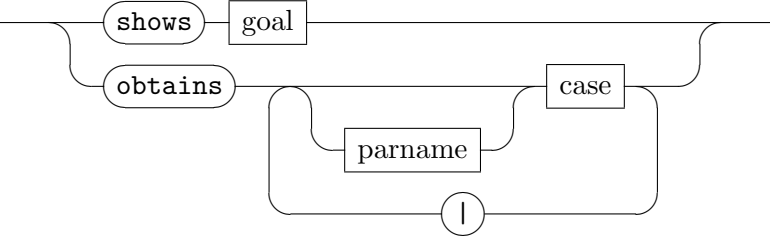
goal



longgoal



conclusion



case



lemma *a*: φ enters proof mode with φ as main goal, eventually resulting in some fact $\vdash \varphi$ to be put back into the target context. An additional *context* specification may build up an initial proof context for the subsequent claim; this includes local definitions and syntax as well, see the definition of *contextelem* in §5.5.

theorem *a*: φ and **corollary** *a*: φ are essentially the same as **lemma** *a*: φ , but the facts are internally marked as being of a different kind. This discrimination acts like a formal comment.

schematic_lemma, **schematic_theorem**, **schematic_corollary** are similar to **lemma**, **theorem**, **corollary**, respectively but allow the statement to contain unbound schematic variables.

Under normal circumstances, an Isar proof text needs to specify claims explicitly. Schematic goals are more like goals in Prolog, where certain results are synthesized in the course of reasoning. With schematic statements, the inherent compositionality of Isar proofs is lost, which also impacts performance, because proof checking is forced into sequential mode.

have *a*: φ claims a local goal, eventually resulting in a fact within the current logical context. This operation is completely independent of any pending sub-goals of an enclosing goal statements, so **have** may be freely used for experimental exploration of potential results within a proof body.

show *a*: φ is like **have** *a*: φ plus a second stage to refine some pending sub-goal for each one of the finished result, after having been exported into the corresponding context (at the head of the sub-proof of this **show** command).

To accommodate interactive debugging, resulting rules are printed before being applied internally. Even more, interactive execution of **show** predicts potential failure and displays the resulting error as a warning beforehand. Watch out for the following message:

Problem! Local statement will fail to solve any pending goal

hence abbreviates “**then have**”, i.e. claims a local goal to be proven by forward chaining the current facts. Note that **hence** is also equivalent to “**from this have**”.

thus abbreviates “**then show**”. Note that **thus** is also equivalent to “**from this show**”.

print_statement *a* prints facts from the current theory or proof context in long statement form, according to the syntax for **lemma** given above.

Any goal statement causes some term abbreviations (such as *?thesis*) to be bound automatically, see also §6.2.2.

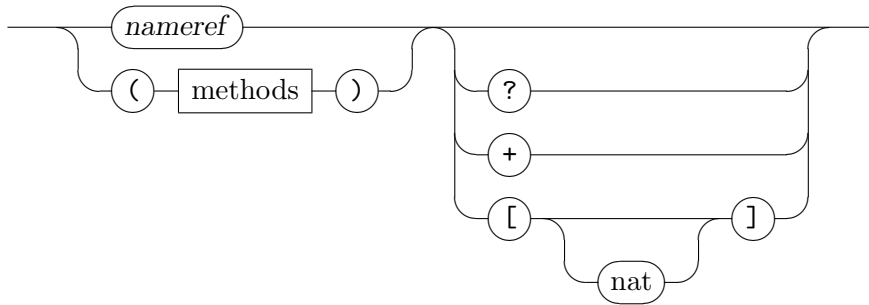
The optional case names of **obtains** have a twofold meaning: (1) during the of this claim they refer to the the local context introductions, (2) the resulting rule is annotated accordingly to support symbolic case splits when used with the *cases* method (cf. §6.6).

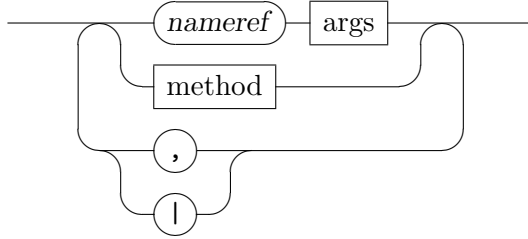
6.3 Refinement steps

6.3.1 Proof method expressions

Proof methods are either basic ones, or expressions composed of methods via “,” (sequential composition), “|” (alternative choices), “?” (try), “+” (repeat at least once), “[*n*]” (restriction to first *n* sub-goals, with default *n* = 1). In practice, proof methods are usually just a comma separated list of *nameref args* specifications. Note that parentheses may be dropped for single method specifications (with no arguments).

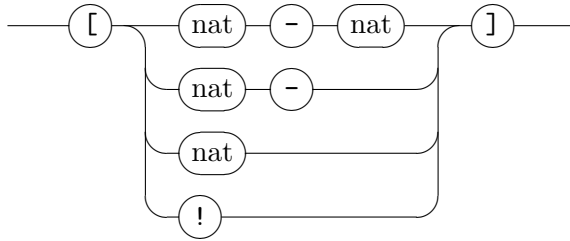
method



methods

Proper Isar proof methods do *not* admit arbitrary goal addressing, but refer either to the first sub-goal or all sub-goals uniformly. The goal restriction operator “[*n*]” evaluates a method expression within a sandbox consisting of the first *n* sub-goals (which need to exist). For example, the method “*simp_all*[3]” simplifies the first three sub-goals, while “(*rule foo*, *simp_all*)[]” simplifies all new goals that emerge from applying rule *foo* to the originally first one.

Improper methods, notably tactic emulations, offer a separate low-level goal addressing scheme as explicit argument to the individual tactic being involved. Here “[!]” refers to all goals, and “[*n*–]” to all goals starting from *n*.

goalspec

6.3.2 Initial and terminal proof steps

```

proof  : proof(prove) → proof(state)
qed    : proof(state) → proof(state) | local_theory | theory
by     : proof(prove) → proof(state) | local_theory | theory
..     : proof(prove) → proof(state) | local_theory | theory
.      : proof(prove) → proof(state) | local_theory | theory
sorry  : proof(prove) → proof(state) | local_theory | theory

```

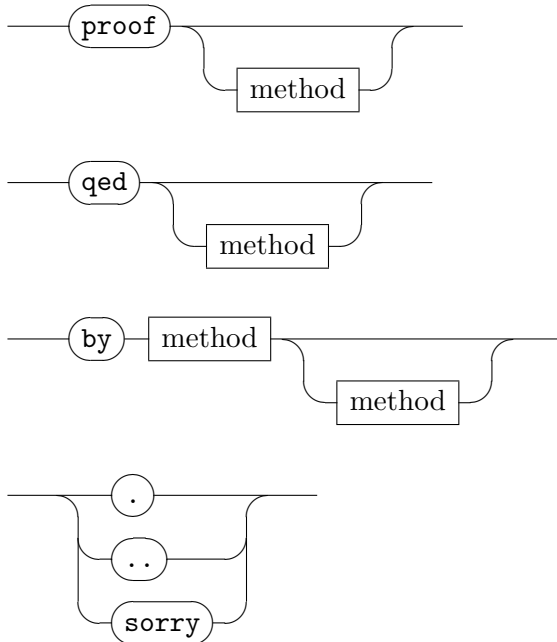
Arbitrary goal refinement via tactics is considered harmful. Structured proof composition in Isar admits proof methods to be invoked in two places only.

1. An *initial* refinement step **proof** m_1 reduces a newly stated goal to a number of sub-goals that are to be solved later. Facts are passed to m_1 for forward chaining, if so indicated by *proof(chain)* mode.
2. A *terminal* conclusion step **qed** m_2 is intended to solve remaining goals. No facts are passed to m_2 .

The only other (proper) way to affect pending goals in a proof body is by **show**, which involves an explicit statement of what is to be solved eventually. Thus we avoid the fundamental problem of unstructured tactic scripts that consist of numerous consecutive goal transformations, with invisible effects.

As a general rule of thumb for good proof style, initial proof methods should either solve the goal completely, or constitute some well-understood reduction to new sub-goals. Arbitrary automatic proof tools that are prone to leave a large number of badly structured sub-goals are no help in continuing the proof document in an intelligible manner.

Unless given explicitly by the user, the default initial method is “*rule*”, which applies a single standard elimination or introduction rule according to the topmost symbol involved. There is no separate default terminal method. Any remaining goals are always solved by assumption in the very last step.



proof m_1 refines the goal by proof method m_1 ; facts for forward chaining are passed if so indicated by *proof(chain)* mode.

qed m_2 refines any remaining goals by proof method m_2 and concludes the sub-proof by assumption. If the goal had been *show* (or *thus*), some pending sub-goal is solved as well by the rule resulting from the result *exported* into the enclosing goal context. Thus *qed* may fail for two reasons: either m_2 fails, or the resulting rule does not fit to any pending goal¹ of the enclosing context. Debugging such a situation might involve temporarily changing **show** into **have**, or weakening the local context by replacing occurrences of **assume** by **presume**.

by m_1 m_2 is a *terminal proof*; it abbreviates **proof** m_1 *qed* m_2 , but with backtracking across both methods. Debugging an unsuccessful **by** m_1 m_2 command can be done by expanding its definition; in many cases **proof** m_1 (or even *apply* m_1) is already sufficient to see the problem.

.. is a *default proof*; it abbreviates **by rule**.

. is a *trivial proof*; it abbreviates **by this**.

sorry is a *fake proof* pretending to solve the pending claim without further ado. This only works in interactive development, or if the **quick_and_dirty** flag is enabled (in ML). Facts emerging from fake proofs are not the real thing. Internally, each theorem container is tainted by an oracle invocation, which is indicated as “[!]” in the printed result.

The most important application of **sorry** is to support experimentation and top-down proof development.

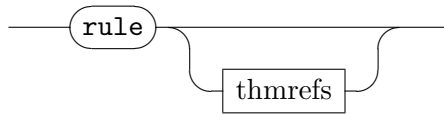
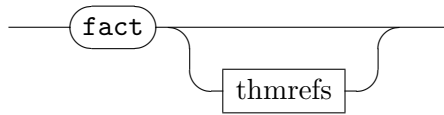
6.3.3 Fundamental methods and attributes

The following proof methods and attributes refer to basic logical operations of Isar. Further methods and attributes are provided by several generic and

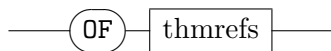
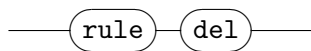
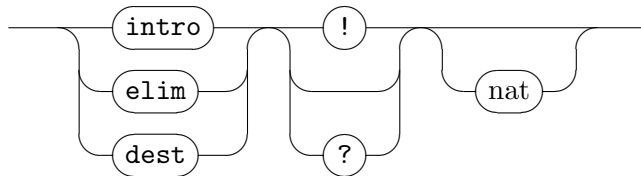
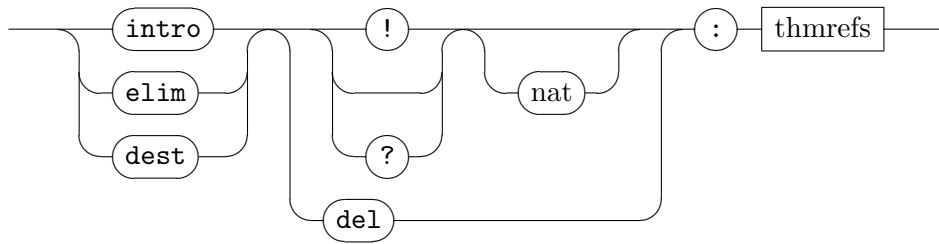
¹This includes any additional “strong” assumptions as introduced by **assume**.

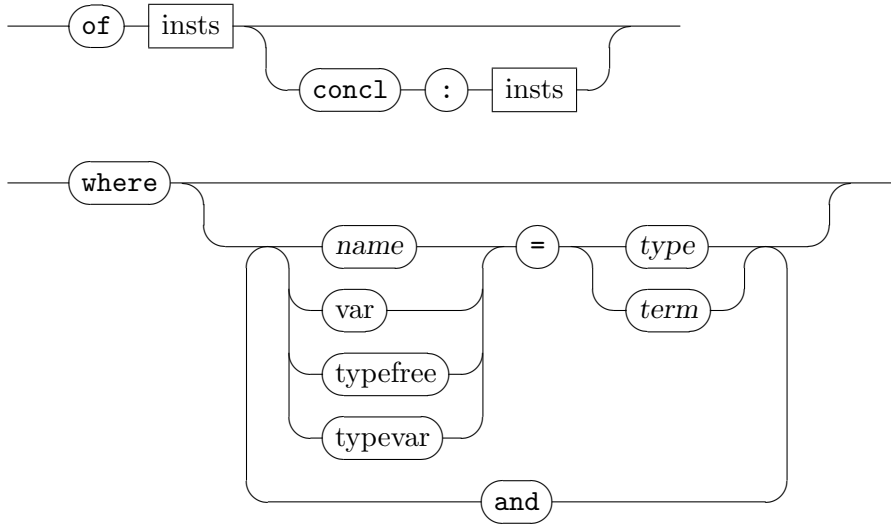
object-logic specific tools and packages (see chapter 9 and chapter 10).

$-$: *method*
fact : *method*
assumption : *method*
this : *method*
rule : *method*
intro : *attribute*
elim : *attribute*
dest : *attribute*
rule : *attribute*
OF : *attribute*
of : *attribute*
where : *attribute*



rulemod





“**–**” (minus) does nothing but insert the forward chaining facts as premises into the goal. Note that command **proof** without any method actually performs a single reduction step using the *rule* method; thus a plain *do-nothing* proof step would be “**proof –**” rather than **proof** alone.

fact $a_1 \dots a_n$ composes some fact from a_1, \dots, a_n (or implicitly from the current proof context) modulo unification of schematic type and term variables. The rule structure is not taken into account, i.e. meta-level implication is considered atomic. This is the same principle underlying literal facts (cf. §3.2.7): “**have** φ **by fact**” is equivalent to “**note** ‘ φ ’” provided that $\vdash \varphi$ is an instance of some known $\vdash \varphi$ in the proof context.

assumption solves some goal by a single assumption step. All given facts are guaranteed to participate in the refinement; this means there may be only 0 or 1 in the first place. Recall that **qed** (§6.3.2) already concludes any remaining sub-goals by assumption, so structured proofs usually need not quote the *assumption* method at all.

this applies all of the current facts directly as rules. Recall that “**.**” (dot) abbreviates “**by this**”.

rule $a_1 \dots a_n$ applies some rule given as argument in backward manner; facts are used to reduce the rule before applying it to the goal. Thus *rule* without facts is plain introduction, while with facts it becomes elimination.

When no arguments are given, the *rule* method tries to pick appropriate rules automatically, as declared in the current context using the *intro*, *elim*, *dest* attributes (see below). This is the default behavior of **proof** and “..” (double-dot) steps (see §6.3.2).

intro, *elim*, and *dest* declare introduction, elimination, and destruct rules, to be used with method *rule*, and similar tools. Note that the latter will ignore rules declared with “?”, while “!” are used most aggressively.

The classical reasoner (see §9.4) introduces its own variants of these attributes; use qualified names to access the present versions of Isabelle/Pure, i.e. *Pure.intro*.

rule del undeclares introduction, elimination, or destruct rules.

OF $a_1 \dots a_n$ applies some theorem to all of the given rules a_1, \dots, a_n (in parallel). This corresponds to the **op MRS** operation in ML, but note the reversed order. Positions may be effectively skipped by including “_” (underscore) as argument.

of $t_1 \dots t_n$ performs positional instantiation of term variables. The terms t_1, \dots, t_n are substituted for any schematic variables occurring in a theorem from left to right; “_” (underscore) indicates to skip a position. Arguments following a “*concl*:” specification refer to positions of the conclusion of a rule.

where $x_1 = t_1$ **and** $\dots x_n = t_n$ performs named instantiation of schematic type and term variables occurring in a theorem. Schematic variables have to be specified on the left-hand side (e.g. $?x1.3$). The question mark may be omitted if the variable name is a plain identifier without index. As type instantiations are inferred from term instantiations, explicit type instantiations are seldom necessary.

6.3.4 Emulating tactic scripts

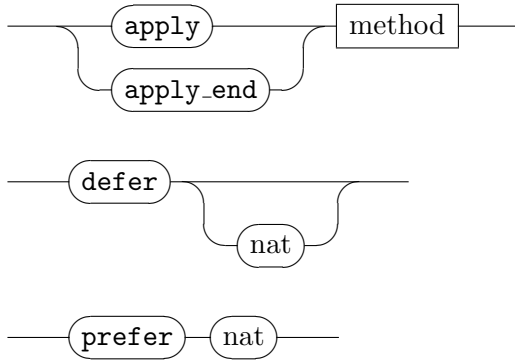
The Isar provides separate commands to accommodate tactic-style proof scripts within the same system. While being outside the orthodox Isar proof language, these might come in handy for interactive exploration and debugging, or even actual tactical proof within new-style theories (to benefit from document preparation, for example). See also §9.2.3 for actual tactics, that

have been encapsulated as proof methods. Proper proof methods may be used in scripts, too.

```

apply*   : proof(prove) → proof(prove)
apply_end* : proof(state) → proof(state)
done*    : proof(prove) → proof(state) | local_theory | theory
defer*   : proof → proof
prefer*  : proof → proof
back*    : proof → proof

```



apply m applies proof method m in initial position, but unlike **proof** it retains “*proof(prove)*” mode. Thus consecutive method applications may be given just as in tactic scripts.

Facts are passed to m as indicated by the goal’s forward-chain mode, and are *consumed* afterwards. Thus any further **apply** command would always work in a purely backward manner.

apply_end m applies proof method m as if in terminal position. Basically, this simulates a multi-step tactic script for **qed**, but may be given anywhere within the proof body.

No facts are passed to m here. Furthermore, the static context is that of the enclosing goal (as for actual **qed**). Thus the proof method may not refer to any assumptions introduced in the current body, for example.

done completes a proof script, provided that the current goal state is solved completely. Note that actual structured proof commands (e.g. “.” or **sorry**) may be used to conclude proof scripts as well.

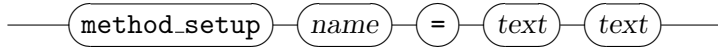
defer n and **prefer** n shuffle the list of pending goals: **defer** puts off sub-goal n to the end of the list ($n = 1$ by default), while **prefer** brings sub-goal n to the front.

back does back-tracking over the result sequence of the latest proof command. Basically, any proof command may return multiple results.

Any proper Isar proof method may be used with tactic script commands such as **apply**. A few additional emulations of actual tactics are provided as well; these would be never used in actual structured proofs, of course.

6.3.5 Defining proof methods

method_setup : *theory* \rightarrow *theory*



method_setup *name* = *text description* defines a proof method in the current theory. The given *text* has to be an ML expression of type `(Proof.context \rightarrow Proof.method) context_parser`, cf. basic parsers defined in structure **Args** and **Attrib**. There are also combinators like **METHOD** and **SIMPLE_METHOD** to turn certain tactic forms into official proof methods; the primed versions refer to tactics with explicit goal addressing.

Here are some example method definitions:

```

method_setup my_method1 = ⟨⟨
  Scan.succeed (K (SIMPLE_METHOD' (fn i: int => no_tac)))
⟩⟩ my first method (without any arguments)

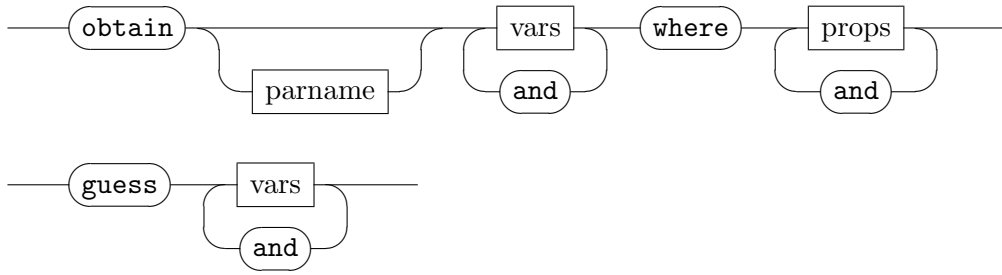
method_setup my_method2 = ⟨⟨
  Scan.succeed (fn ctxt: Proof.context =>
    SIMPLE_METHOD' (fn i: int => no_tac))
⟩⟩ my second method (with context)

method_setup my_method3 = ⟨⟨
  Attrib.thms >> (fn thms: thm list => fn ctxt: Proof.context =>
    SIMPLE_METHOD' (fn i: int => no_tac))
⟩⟩ my third method (with theorem arguments and context)
  
```

6.4 Generalized elimination

obtain : *proof*(*state*) | *proof*(*chain*) \rightarrow *proof*(*prove*)
guess* : *proof*(*state*) | *proof*(*chain*) \rightarrow *proof*(*prove*)

Generalized elimination means that additional elements with certain properties may be introduced in the current context, by virtue of a locally proven “soundness statement”. Technically speaking, the **obtain** language element is like a declaration of **fix** and **assume** (see also §6.2.1), together with a soundness proof of its additional claim. According to the nature of existential reasoning, assumptions get eliminated from any result exported from the context later, provided that the corresponding parameters do *not* occur in the conclusion.



The derived Isar command **obtain** is defined as follows (where b_1, \dots, b_k shall refer to (optional) facts indicated for forward chaining).

$$\begin{aligned}
 \langle \text{using } b_1 \dots b_k \rangle \text{ **obtain** } x_1 \dots x_m \text{ **where** } a: \varphi_1 \dots \varphi_n \langle \text{proof} \rangle &\equiv \\
 \text{have } \bigwedge thesis. (\bigwedge x_1 \dots x_m. \varphi_1 \implies \dots \varphi_n \implies thesis) &\implies thesis \\
 \text{proof } succeed & \\
 \text{fix } thesis & \\
 \text{assume that } [Pure.intro?]: \bigwedge x_1 \dots x_m. \varphi_1 \implies \dots \varphi_n &\implies thesis \\
 \text{then show } thesis & \\
 \text{apply } - & \\
 \text{using } b_1 \dots b_k \langle \text{proof} \rangle & \\
 \text{qed} & \\
 \text{fix } x_1 \dots x_m \text{ **assume}^* a: \varphi_1 \dots \varphi_n &
 \end{aligned}**$$

Typically, the soundness proof is relatively straight-forward, often just by canonical automated tools such as “**by simp**” or “**by blast**”. Accordingly, the “*that*” reduction above is declared as simplification and introduction rule.

In a sense, **obtain** represents at the level of Isar proofs what would be meta-logical existential quantifiers and conjunctions. This concept has a broad range of useful applications, ranging from plain elimination (or introduction) of object-level existential and conjunctions, to elimination over results of symbolic evaluation of recursive definitions, for example. Also note that **obtain** without parameters acts much like **have**, where the result is treated as a genuine assumption.

An alternative name to be used instead of “*that*” above may be given in parentheses.

The improper variant **guess** is similar to **obtain**, but derives the obtained statement from the course of reasoning! The proof starts with a fixed goal *thesis*. The subsequent proof may refine this to anything of the form like $\bigwedge x_1 \dots x_m. \varphi_1 \implies \dots \varphi_n \implies thesis$, but must not introduce new subgoals. The final goal state is then used as reduction rule for the obtain scheme described above. Obtained parameters x_1, \dots, x_m are marked as internal by default, which prevents the proof context from being polluted by ad-hoc variables. The variable names and type constraints given as arguments for **guess** specify a prefix of obtained parameters explicitly in the text.

It is important to note that the facts introduced by **obtain** and **guess** may not be polymorphic: any type-variables occurring here are fixed in the present context!

6.5 Calculational reasoning

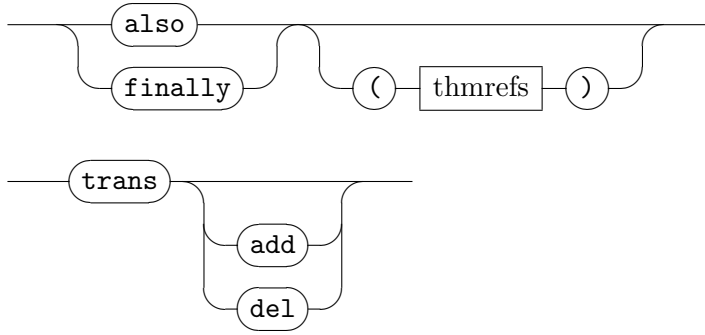
also	:	$proof(state) \rightarrow proof(state)$
finally	:	$proof(state) \rightarrow proof(chain)$
moreover	:	$proof(state) \rightarrow proof(state)$
ultimately	:	$proof(state) \rightarrow proof(chain)$
print_trans_rules*	:	$context \rightarrow$
		$trans : attribute$
		$sym : attribute$
		$symmetric : attribute$

Calculational proof is forward reasoning with implicit application of transitivity rules (such those of $=$, \leq , $<$). Isabelle/Isar maintains an auxiliary fact register *calculation* for accumulating results obtained by transitivity composed with the current result. Command **also** updates *calculation* involving *this*, while **finally** exhibits the final *calculation* by forward chaining towards the next goal statement. Both commands require valid current facts, i.e. may occur only after commands that produce theorems such as **assume**, **note**, or some finished proof of **have**, **show** etc. The **moreover** and **ultimately** commands are similar to **also** and **finally**, but only collect further results in *calculation* without applying any rules yet.

Also note that the implicit term abbreviation “ \dots ” has its canonical application with calculational proofs. It refers to the argument of the preceding statement. (The argument of a curried infix expression happens to be its right-hand side.)

Isabelle/Isar calculations are implicitly subject to block structure in the sense that new threads of calculational reasoning are commenced for any new block (as opened by a local goal, for example). This means that, apart from being able to nest calculations, there is no separate *begin-calculation* command required.

The Isar calculation proof commands may be defined as follows:²

$$\begin{aligned} \mathbf{also}_0 &\equiv \mathbf{note} \text{ calculation} = \text{this} \\ \mathbf{also}_n+1 &\equiv \mathbf{note} \text{ calculation} = \text{trans } [OF \text{ calculation this}] \\ \mathbf{finally} &\equiv \mathbf{also} \mathbf{from} \text{ calculation} \\ \mathbf{moreover} &\equiv \mathbf{note} \text{ calculation} = \text{calculation this} \\ \mathbf{ultimately} &\equiv \mathbf{moreover} \mathbf{from} \text{ calculation} \end{aligned}$$


also ($a_1 \dots a_n$) maintains the auxiliary *calculation* register as follows. The first occurrence of **also** in some calculational thread initializes *calculation* by *this*. Any subsequent **also** on the same level of block-structure updates *calculation* by some transitivity rule applied to *calculation* and *this* (in that order). Transitivity rules are picked from the current context, unless alternative rules are given as explicit arguments.

finally ($a_1 \dots a_n$) maintaining *calculation* in the same way as **also**, and concludes the current calculational thread. The final result is exhibited as fact for forward chaining towards the next goal. Basically, **finally** just abbreviates **also from calculation**. Typical idioms for concluding calculational proofs are “**finally show** *?thesis* .” and “**finally have** φ .”.

moreover and **ultimately** are analogous to **also** and **finally**, but collect results only, without applying rules.

²We suppress internal bookkeeping such as proper handling of block-structure.

print_trans_rules prints the list of transitivity rules (for calculational commands **also** and **finally**) and symmetry rules (for the *symmetric* operation and single step elimination patters) of the current context.

trans declares theorems as transitivity rules.

sym declares symmetry rules, as well as *Pure.elim?* rules.

symmetric resolves a theorem with some rule declared as *sym* in the current context. For example, “**assume** [*symmetric*]: $x = y$ ” produces a swapped fact derived from that assumption.

In structured proof texts it is often more appropriate to use an explicit single-step elimination proof, such as “**assume** $x = y$ **then have** $y = x$..”.

6.6 Proof by cases and induction

6.6.1 Rule contexts

```

      case      : proof(state) → proof(state)
  print_cases*  : context →
    case_names  : attribute
  case_conclusion : attribute
      params    : attribute
    consumes   : attribute

```

The puristic way to build up Isar proof contexts is by explicit language elements like **fix**, **assume**, **let** (see §6.2.1). This is adequate for plain natural deduction, but easily becomes unwieldy in concrete verification tasks, which typically involve big induction rules with several cases.

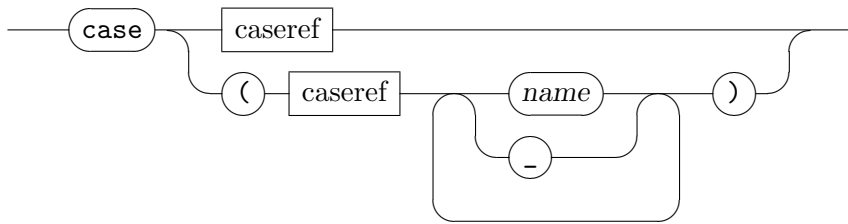
The **case** command provides a shorthand to refer to a local context symbolically: certain proof methods provide an environment of named “cases” of the form $c: x_1, \dots, x_m, \varphi_1, \dots, \varphi_n$; the effect of “**case** c ” is then equivalent to “**fix** $x_1 \dots x_m$ **assume** $c: \varphi_1 \dots \varphi_n$ ”. Term bindings may be covered as well, notably *?case* for the main conclusion.

By default, the “terminology” x_1, \dots, x_m of a case value is marked as hidden, i.e. there is no way to refer to such parameters in the subsequent proof text. After all, original rule parameters stem from somewhere outside of the current proof text. By using the explicit form “**case** (c $y_1 \dots y_m$)” instead, the proof author is able to chose local names that fit nicely into the current context.

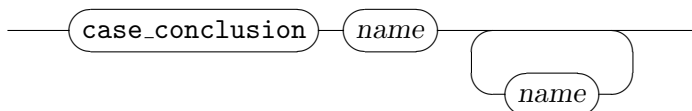
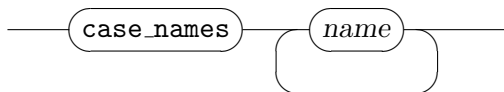
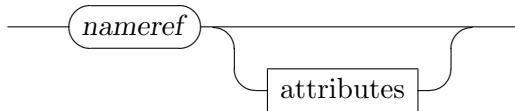
It is important to note that proper use of **case** does not provide means to peek at the current goal state, which is not directly observable in Isar! Nonetheless, goal refinement commands do provide named cases $goal_i$ for each subgoal $i = 1, \dots, n$ of the resulting goal state. Using this extra feature requires great care, because some bits of the internal tactical machinery intrude the proof text. In particular, parameter names stemming from the left-over of automated reasoning tools are usually quite unpredictable.

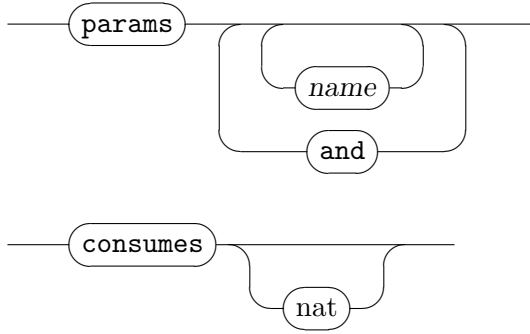
Under normal circumstances, the text of cases emerge from standard elimination or induction rules, which in turn are derived from previous theory specifications in a canonical way (say from **inductive** definitions).

Proper cases are only available if both the proof method and the rules involved support this. By using appropriate attributes, case names, conclusions, and parameters may be also declared by hand. Thus variant versions of rules that have been derived manually become ready to use in advanced case analysis later.



caseref





case ($c\ x_1 \dots x_m$) invokes a named local context $c: x_1, \dots, x_m, \varphi_1, \dots, \varphi_m$, as provided by an appropriate proof method (such as *cases* and *induct*). The command “**case** ($c\ x_1 \dots x_m$)” abbreviates “**fix** $x_1 \dots x_m$ **assume** $c: \varphi_1 \dots \varphi_m$ ”.

print_cases prints all local contexts of the current state, using Isar proof language notation.

case_names $c_1 \dots c_k$ declares names for the local contexts of premises of a theorem; c_1, \dots, c_k refers to the *suffix* of the list of premises.

case_conclusion $c\ d_1 \dots d_k$ declares names for the conclusions of a named premise c ; here d_1, \dots, d_k refers to the prefix of arguments of a logical formula built by nesting a binary connective (e.g. \vee).

Note that proof methods such as *induct* and *coinduct* already provide a default name for the conclusion as a whole. The need to name sub-formulas only arises with cases that split into several sub-cases, as in common co-induction rules.

params $p_1 \dots p_m$ **and** $\dots q_1 \dots q_n$ renames the innermost parameters of premises 1, \dots , n of some theorem. An empty list of names may be given to skip positions, leaving the present parameters unchanged.

Note that the default usage of case rules does *not* directly expose parameters to the proof context.

consumes n declares the number of “major premises” of a rule, i.e. the number of facts to be consumed when it is applied by an appropriate proof method. The default value of *consumes* is $n = 1$, which is appropriate for the usual kind of cases and induction rules for inductive sets (cf. §10.7). Rules without any *consumes* declaration given are treated as if *consumes* 0 had been specified.

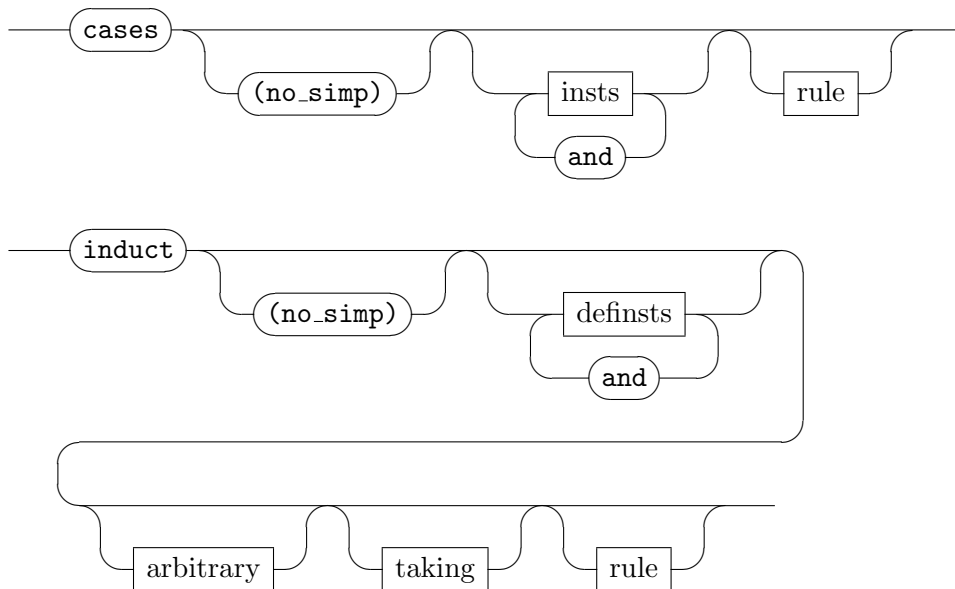
Note that explicit *consumes* declarations are only rarely needed; this is already taken care of automatically by the higher-level *cases*, *induct*, and *coinduct* declarations.

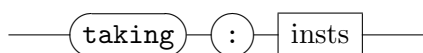
6.6.2 Proof methods

cases : *method*
induct : *method*
coinduct : *method*

The *cases*, *induct*, and *coinduct* methods provide a uniform interface to common proof techniques over datatypes, inductive predicates (or sets), recursive functions etc. The corresponding rules may be specified and instantiated in a casual manner. Furthermore, these methods provide named local contexts that may be invoked via the **case** proof command within the subsequent proof text. This accommodates compact proof texts even when reasoning about large specifications.

The *induct* method also provides some additional infrastructure in order to be applicable to structure statements (either using explicit meta-level connectives, or including facts and parameters separately). This avoids cumbersome encoding of “strengthened” inductive statements within the object-logic.





cases insts R applies method *rule* with an appropriate case distinction theorem, instantiated to the subjects *insts*. Symbolic case names are bound according to the rule's local contexts.

The rule is determined as follows, according to the facts and arguments passed to the *cases* method:

facts	arguments	rule
	<i>cases</i>	classical case split
	<i>cases</i> t	datatype exhaustion (type of t)
$\vdash A \ t$	<i>cases</i> \dots	inductive predicate/set elimination (of A)
\dots	<i>cases</i> \dots rule: R	explicit rule R

Several instantiations may be given, referring to the *suffix* of premises of the case rule; within each premise, the *prefix* of variables is instantiated. In most situations, only a single term needs to be specified; this refers to the first variable of the last premise (it is usually the same for all cases). The (*no_simp*) option can be used to disable pre-simplification of cases (see the description of *induct* below for details).

induct insts R is analogous to the *cases* method, but refers to induction rules, which are determined as follows:

facts	arguments	rule
	<i>induct</i> $P \ x$	datatype induction (type of x)
$\vdash A \ x$	<i>induct</i> \dots	predicate/set induction (of A)
\dots	<i>induct</i> \dots rule: R	explicit rule R

Several instantiations may be given, each referring to some part of a mutual inductive definition or datatype — only related partial induction rules may be used together, though. Any of the lists of terms P, x, \dots refers to the *suffix* of variables present in the induction rule. This enables the writer to specify only induction variables, or both predicates and variables, for example.

Instantiations may be definitional: equations $x \equiv t$ introduce local definitions, which are inserted into the claim and discharged after applying the induction rule. Equalities reappear in the inductive cases, but have been transformed according to the induction principle being involved here. In order to achieve practically useful induction hypotheses, some variables occurring in t need to be fixed (see below). Instantiations of the form t , where t is not a variable, are taken as a shorthand for $x \equiv t$, where x is a fresh variable. If this is not intended, t has to be enclosed in parentheses. By default, the equalities generated by definitional instantiations are pre-simplified using a specific set of rules, usually consisting of distinctness and injectivity theorems for datatypes. This pre-simplification may cause some of the parameters of an inductive

case to disappear, or may even completely delete some of the inductive cases, if one of the equalities occurring in their premises can be simplified to *False*. The (*no_simp*) option can be used to disable pre-simplification. Additional rules to be used in pre-simplification can be declared using the *induct_simp* attribute.

The optional “*arbitrary: x₁ ... x_m*” specification generalizes variables x_1, \dots, x_m of the original goal before applying induction. Thus induction hypotheses may become sufficiently general to get the proof through. Together with definitional instantiations, one may effectively perform induction over expressions of a certain structure.

The optional “*taking: t₁ ... t_n*” specification provides additional instantiations of a prefix of pending variables in the rule. Such schematic induction rules rarely occur in practice, though.

coinduct inst R is analogous to the *induct* method, but refers to coinduction rules, which are determined as follows:

goal		arguments	rule
	<i>coinduct</i>	x	type coinduction (type of x)
$A \ x$	<i>coinduct</i>	\dots	predicate/set coinduction (of A)
\dots	<i>coinduct</i>	$\dots \text{ rule: } R$	explicit rule R

Coinduction is the dual of induction. Induction essentially eliminates $A \ x$ towards a generic result $P \ x$, while coinduction introduces $A \ x$ starting with $B \ x$, for a suitable “bisimulation” B . The cases of a coinduct rule are typically named after the predicates or sets being covered, while the conclusions consist of several alternatives being named after the individual destructor patterns.

The given instantiation refers to the *suffix* of variables occurring in the rule’s major premise, or conclusion if unavailable. An additional “*taking: t₁ ... t_n*” specification may be required in order to specify the bisimulation to be used in the coinduction step.

Above methods produce named local contexts, as determined by the instantiated rule as given in the text. Beyond that, the *induct* and *coinduct* methods guess further instantiations from the goal specification itself. Any persisting unresolved schematic variables of the resulting rule will render the the corresponding case invalid. The term binding *?case* for the conclusion will be provided with each case, provided that term is fully specified.

The **print_cases** command prints all named cases present in the current proof state.

Despite the additional infrastructure, both *cases* and *coinduct* merely apply a certain rule, after instantiation, while conforming due to the usual way of monotonic natural deduction: the context of a structured statement $\bigwedge x_1 \dots x_m. \varphi_1 \implies \dots \varphi_n \implies \dots$ reappears unchanged after the case split.

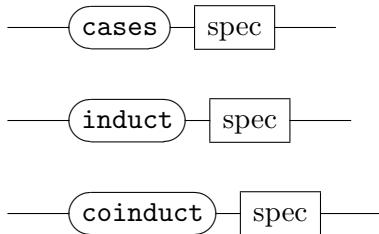
The *induct* method is fundamentally different in this respect: the meta-level structure is passed through the “recursive” course involved in the induction. Thus the original statement is basically replaced by separate copies, corresponding to the induction hypotheses and conclusion; the original goal context is no longer available. Thus local assumptions, fixed parameters and definitions effectively participate in the inductive rephrasing of the original statement.

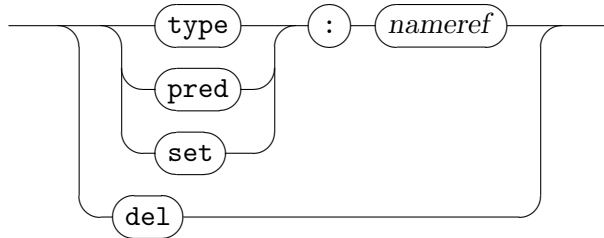
In induction proofs, local assumptions introduced by cases are split into two different kinds: *hyps* stemming from the rule and *prems* from the goal statement. This is reflected in the extracted cases accordingly, so invoking “**case** *c*” will provide separate facts *c.hyps* and *c.prems*, as well as fact *c* to hold the all-inclusive list.

Facts presented to either method are consumed according to the number of “major premises” of the rule involved, which is usually 0 for plain cases and induction rules of datatypes etc. and 1 for rules of inductive predicates or sets and the like. The remaining facts are inserted into the goal verbatim before the actual *cases*, *induct*, or *coinduct* rule is applied.

6.6.3 Declaring rules

```
print_induct_rules* : context →
    cases : attribute
    induct : attribute
    coinduct : attribute
```



spec

print_induct_rules prints cases and induct rules for predicates (or sets) and types of the current context.

cases, *induct*, and *coinduct* (as attributes) declare rules for reasoning about (co)inductive predicates (or sets) and types, using the corresponding methods of the same name. Certain definitional packages of object-logics usually declare emerging cases and induction rules as expected, so users rarely need to intervene.

Rules may be deleted via the *del* specification, which covers all of the *type/pred/set* sub-categories simultaneously. For example, *cases del* removes any *cases* rules declared for some type, predicate, or set.

Manual rule declarations usually refer to the *case_names* and *params* attributes to adjust names of cases and parameters of a rule; the *consumes* declaration is taken care of automatically: *consumes* 0 is specified for “type” rules and *consumes* 1 for “predicate” / “set” rules.

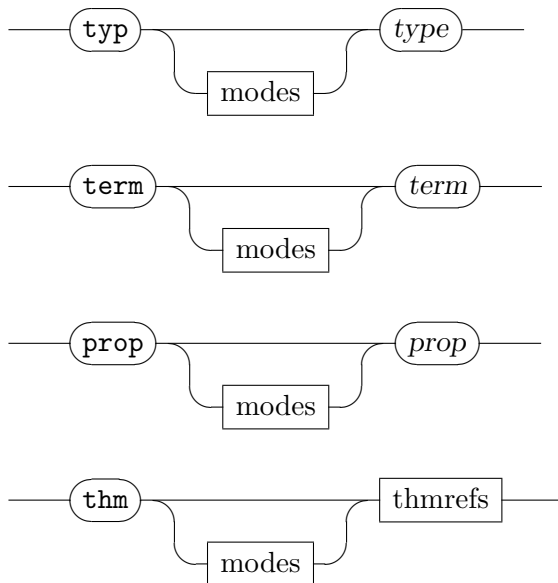
Inner syntax — the term language

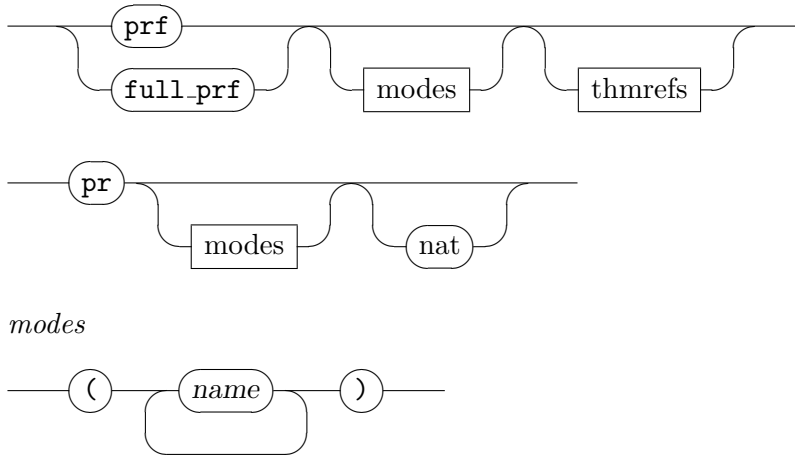
7.1 Printing logical entities

7.1.1 Diagnostic commands

`typ*` : *context* →
`term*` : *context* →
`prop*` : *context* →
`thm*` : *context* →
`prf*` : *context* →
`full_prf*` : *context* →
`pr*` : *any* →

These diagnostic commands assist interactive development by printing internal logical entities in a human-readable fashion.





typ τ reads and prints types of the meta-logic according to the current theory or proof context.

term t and **prop** φ read, type-check and print terms or propositions according to the current theory or proof context; the inferred type of t is output as well. Note that these commands are also useful in inspecting the current environment of term abbreviations.

thm $a_1 \dots a_n$ retrieves theorems from the current theory or proof context. Note that any attributes included in the theorem specifications are applied to a temporary context derived from the current theory or proof; the result is discarded, i.e. attributes involved in a_1, \dots, a_n do not have any permanent effect.

prf displays the (compact) proof term of the current proof state (if present), or of the given theorems. Note that this requires proof terms to be switched on for the current object logic (see the “Proof terms” section of the Isabelle reference manual for information on how to do this).

full_prf is like **prf**, but displays the full proof term, i.e. also displays information omitted in the compact proof term, which is denoted by “_” placeholders there.

pr goals prints the current proof state (if present), including current facts and goals. The optional limit arguments affect the number of goals to be displayed, which is initially 10. Omitting limit value leaves the current setting unchanged.

All of the diagnostic commands above admit a list of *modes* to be specified, which is appended to the current print mode (see also [24]). Thus the output behavior may be modified according particular print mode features. For example, **pr** (*latex xsymbols*) would print the current proof state with mathematical symbols and special characters represented in L^AT_EX source, according to the Isabelle style [39].

Note that antiquotations (cf. §4.2) provide a more systematic way to include formal items into the printed text document.

7.1.2 Details of printed content

<code>show_types</code>	: bool Config.T	default false
<code>show_sorts</code>	: bool Config.T	default false
<code>show_consts</code>	: bool Config.T	default false
<code>show_abbrevs</code>	: bool Config.T	default true
<code>long_names</code>	: bool Unsynchronized.ref	default false
<code>short_names</code>	: bool Unsynchronized.ref	default false
<code>unique_names</code>	: bool Unsynchronized.ref	default true
<code>show_brackets</code>	: bool Config.T	default false
<code>eta_contract</code>	: bool Config.T	default true
<code>Goal_Display.goals_limit</code>	: int Config.T	default 10
<code>Goal_Display.show_main_goal</code>	: bool Config.T	default false
<code>show_hyps</code>	: bool Config.T	default false
<code>show_tags</code>	: bool Config.T	default false
<code>show_question_marks</code>	: bool Config.T	default true

These global ML variables control the detail of information that is displayed for types, terms, theorems, goals etc.

In interactive sessions, the user interface usually manages these global parameters of the Isabelle process, even with some concept of persistence. Nonetheless it is occasionally useful to manipulate ML variables directly, e.g. using **ML_val** or **ML_command**.

Batch-mode logic sessions may be configured by putting appropriate ML text directly into the `ROOT.ML` file.

`show_types` and `show_sorts` control printing of type constraints for term variables, and sort constraints for type variables. By default, neither of these are shown in output. If `show_sorts` is set to `true`, types are always shown as well.

Note that displaying types and sorts may explain why a polymorphic inference rule fails to resolve with some goal, or why a rewrite rule does not apply as expected.

`show_consts` controls printing of types of constants when displaying a goal state.

Note that the output can be enormous, because polymorphic constants often occur at several different type instances.

`show_abbrevs` controls folding of constant abbreviations.

`long_names`, `short_names`, and `unique_names` control the way of printing fully qualified internal names in external form. See also §4.2 for the document antiquotation options of the same names.

`show_brackets` controls bracketing in pretty printed output. If set to `true`, all sub-expressions of the pretty printing tree will be parenthesized, even if this produces malformed term syntax! This crude way of showing the internal structure of pretty printed entities may occasionally help to diagnose problems with operator priorities, for example.

`eta_contract` controls η -contracted printing of terms.

The η -contraction law asserts $(\lambda x. f\ x) \equiv f$, provided x is not free in f . It asserts *extensionality* of functions: $f \equiv g$ if $f\ x \equiv g\ x$ for all x . Higher-order unification frequently puts terms into a fully η -expanded form. For example, if F has type $(\tau \Rightarrow \tau) \Rightarrow \tau$ then its expanded form is $\lambda h. F\ (\lambda x. h\ x)$.

Setting `eta_contract` makes Isabelle perform η -contractions before printing, so that $\lambda h. F\ (\lambda x. h\ x)$ appears simply as F .

Note that the distinction between a term and its η -expanded form occasionally matters. While higher-order resolution and rewriting operate modulo $\alpha\beta\eta$ -conversion, some other tools might look at terms more discretely.

`Goal_Display.goals_limit` controls the maximum number of subgoals to be shown in goal output.

`Goal_Display.show_main_goal` controls whether the main result to be proven should be displayed. This information might be relevant for schematic goals, to inspect the current claim that has been synthesized so far.

`show_hyps` controls printing of implicit hypotheses of local facts. Normally, only those hypotheses are displayed that are *not* covered by the assumptions of the current context: this situation indicates a fault in some tool being used.

By setting `show_hyps` to `true`, output of *all* hypotheses can be enforced, which is occasionally useful for diagnostic purposes.

`show_tags` controls printing of extra annotations within theorems, such as internal position information, or the case names being attached by the attribute `case_names`.

Note that the *tagged* and *untagged* attributes provide low-level access to the collection of tags associated with a theorem.

`show_question_marks` controls printing of question marks for schematic variables, such as `?x`. Only the leading question mark is affected, the remaining text is unchanged (including proper markup for schematic variables that might be relevant for user interfaces).

7.1.3 Printing limits

```
Pretty.margin_default: int Unsynchronized.ref
print_depth: int -> unit
```

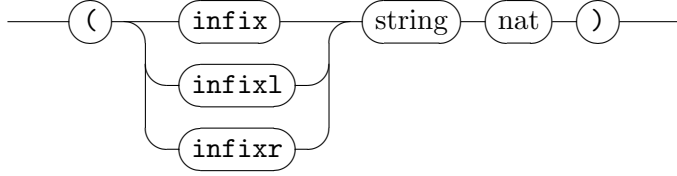
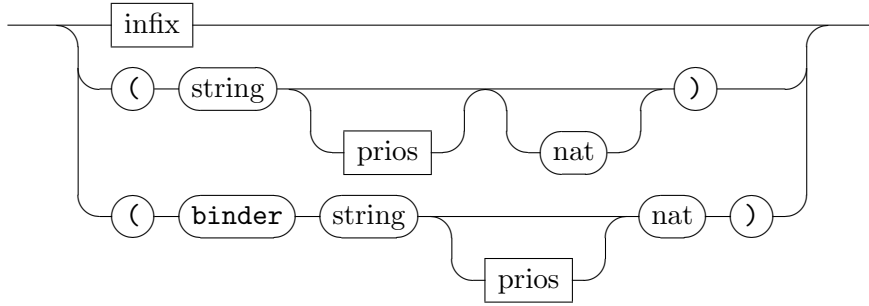
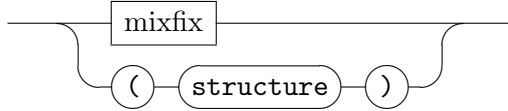
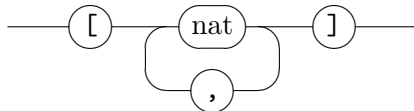
These ML functions set limits for pretty printed text.

`Pretty.margin_default` indicates the global default for the right margin of the built-in pretty printer, with initial value 76. Note that user-interfaces typically control margins automatically when resizing windows, or even bypass the formatting engine of Isabelle/ML altogether and do it within the front end via Isabelle/Scala.

`print_depth` *n* limits the printing depth of the ML toplevel pretty printer; the precise effect depends on the ML compiler and run-time system. Typically *n* should be less than 10. Bigger values such as 100–1000 are useful for debugging.

7.2 Mixfix annotations

Mixfix annotations specify concrete *inner syntax* of Isabelle types and terms. Locally fixed parameters in toplevel theorem statements, locale specifications etc. also admit mixfix annotations.

infix*mixfix**structmixfix**prios*

Here the string specifications refer to the actual mixfix template, which may include literal text, spacing, blocks, and arguments (denoted by “_”); the special symbol “\<index>” (printed as “i”) represents an index argument that specifies an implicit structure reference (see also §5.5). Infix and binder declarations provide common abbreviations for particular mixfix declarations. So in practice, mixfix templates mostly degenerate to literal text for concrete syntax, such as “++” for an infix symbol.

In full generality, mixfix declarations work as follows. Suppose a constant $c :: \tau_1 \Rightarrow \dots \tau_n \Rightarrow \tau$ is annotated by $(\text{mixfix } [p_1, \dots, p_n] p)$, where *mixfix* is a string $d_0 - d_1 - \dots - d_n$ consisting of delimiters that surround argument positions as indicated by underscores.

Altogether this determines a production for a context-free priority grammar, where for each argument i the syntactic category is determined by τ_i

(with priority p_i), and the result category is determined from τ (with priority p). Priority specifications are optional, with default 0 for arguments and 1000 for the result.

Since τ may be again a function type, the constant type scheme may have more argument positions than the mixfix pattern. Printing a nested application $c\ t_1 \dots t_m$ for $m > n$ works by attaching concrete notation only to the innermost part, essentially by printing $(c\ t_1 \dots t_n) \dots t_m$ instead. If a term has fewer arguments than specified in the mixfix template, the concrete syntax is ignored.

A mixfix template may also contain additional directives for pretty printing, notably spaces, blocks, and breaks. The general template format is a sequence over any of the following entities.

d is a delimiter, namely a non-empty sequence of characters other than the following special characters:

- ' single quote
- _ underscore
- 1 index symbol
- (open parenthesis
-) close parenthesis
- / slash

' escapes the special meaning of these meta-characters, producing a literal version of the following character, unless that is a blank.

A single quote followed by a blank separates delimiters, without affecting printing, but input tokens may have additional white space here.

_ is an argument position, which stands for a certain syntactic category in the underlying grammar.

1 is an indexed argument position; this is the place where implicit structure arguments can be attached.

s is a non-empty sequence of spaces for printing. This and the following specifications do not affect parsing at all.

$(n$ opens a pretty printing block. The optional number specifies how much indentation to add when a line break occurs within the block. If the parenthesis is not followed by digits, the indentation defaults to 0. A block specified via $(00$ is unbreakable.

) closes a pretty printing block.

// forces a line break.

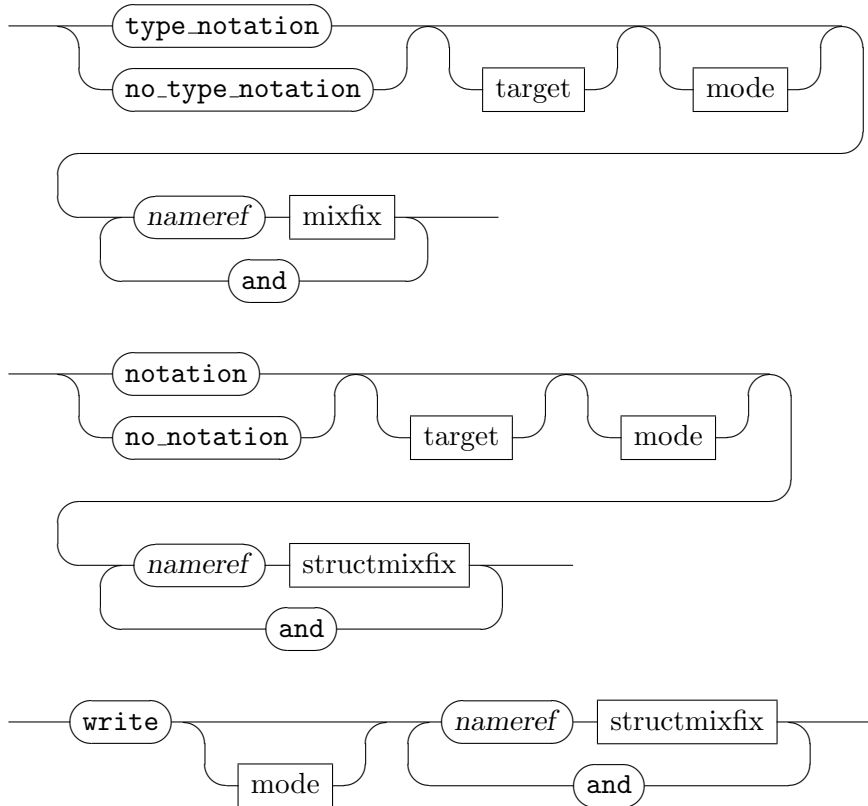
/s allows a line break. Here *s* stands for the string of spaces (zero or more) right after the slash. These spaces are printed if the break is *not* taken.

For example, the template (`_ +/ _`) specifies an infix operator. There are two argument positions; the delimiter `+` is preceded by a space and followed by a space or line break; the entire phrase is a pretty printing block.

The general idea of pretty printing with blocks and breaks is also described in [29].

7.3 Explicit notation

type_notation : *local_theory* → *local_theory*
no_type_notation : *local_theory* → *local_theory*
notation : *local_theory* → *local_theory*
no_notation : *local_theory* → *local_theory*
write : *proof(state)* → *proof(state)*



type_notation c (mx) associates mixfix syntax with an existing type constructor. The arity of the constructor is retrieved from the context.

no_type_notation is similar to **type_notation**, but removes the specified syntax annotation from the present context.

notation c (mx) associates mixfix syntax with an existing constant or fixed variable. The type declaration of the given entity is retrieved from the context.

no_notation is similar to **notation**, but removes the specified syntax annotation from the present context.

write is similar to **notation**, but works within an Isar proof body.

Note that the more primitive commands **syntax** and **no_syntax** (§7.6) provide raw access to the syntax tables of a global theory.

7.4 The Pure syntax

7.4.1 Priority grammars

A context-free grammar consists of a set of *terminal symbols*, a set of *non-terminal symbols* and a set of *productions*. Productions have the form $A = \gamma$, where A is a nonterminal and γ is a string of terminals and nonterminals. One designated nonterminal is called the *root symbol*. The language defined by the grammar consists of all strings of terminals that can be derived from the root symbol by applying productions as rewrite rules.

The standard Isabelle parser for inner syntax uses a *priority grammar*. Each nonterminal is decorated by an integer priority: $A^{(p)}$. In a derivation, $A^{(p)}$ may be rewritten using a production $A^{(q)} = \gamma$ only if $p \leq q$. Any priority grammar can be translated into a normal context-free grammar by introducing new nonterminals and productions.

Formally, a set of context free productions G induces a derivation relation \longrightarrow_G as follows. Let α and β denote strings of terminal or nonterminal symbols. Then $\alpha A^{(p)} \beta \longrightarrow_G \alpha \gamma \beta$ holds if and only if G contains some production $A^{(q)} = \gamma$ for $p \leq q$.

The following grammar for arithmetic expressions demonstrates how binding power and associativity of operators can be enforced by priorities.

$$\begin{aligned}
A^{(1000)} &= (A^{(0)}) \\
A^{(1000)} &= 0 \\
A^{(0)} &= A^{(0)} + A^{(1)} \\
A^{(2)} &= A^{(3)} * A^{(2)} \\
A^{(3)} &= - A^{(3)}
\end{aligned}$$

The choice of priorities determines that $-$ binds tighter than $*$, which binds tighter than $+$. Furthermore $+$ associates to the left and $*$ to the right.

For clarity, grammars obey these conventions:

- All priorities must lie between 0 and 1000.
- Priority 0 on the right-hand side and priority 1000 on the left-hand side may be omitted.
- The production $A^{(p)} = \alpha$ is written as $A = \alpha \ (p)$, i.e. the priority of the left-hand side actually appears in a column on the far right.
- Alternatives are separated by $|$.
- Repetition is indicated by dots (...) in an informal but obvious way.

Using these conventions, the example grammar specification above takes the form:

$$\begin{array}{rcl}
A & = & (A) \\
| & & 0 \\
| & & A + A^{(1)} \quad (0) \\
| & & A^{(3)} * A^{(2)} \quad (2) \\
| & & - A^{(3)} \quad (3)
\end{array}$$

7.4.2 The Pure grammar

The priority grammar of the *Pure* theory is defined as follows:

$$any = prop \mid logic$$

$$\begin{array}{rcl}
prop & = & (prop) \\
| & & prop^{(4)} :: type \quad (3) \\
| & & any^{(3)} =?= any^{(2)} \quad (2) \\
| & & any^{(3)} == any^{(2)} \quad (2)
\end{array}$$

		$any^{(3)} \equiv any^{(2)}$	(2)
		$prop^{(3)} \&\&\& prop^{(2)}$	(2)
		$prop^{(2)} ==> prop^{(1)}$	(1)
		$prop^{(2)} \implies prop^{(1)}$	(1)
		$[prop ; \dots ; prop] ==> prop^{(1)}$	(1)
		$\llbracket prop ; \dots ; prop \rrbracket \implies prop^{(1)}$	(1)
		$!! idts . prop$	(0)
		$\bigwedge idts . prop$	(0)
		OFCLASS (<i>type</i> , <i>logic</i>)	
		SORT_CONSTRAINT (<i>type</i>)	
		TERM <i>logic</i>	
		PROP <i>aprop</i>	
<i>aprop</i>	=	(<i>aprop</i>)	
		<i>id</i> <i>longid</i> <i>var</i> <i>_</i> ...	
		CONST <i>id</i> CONST <i>longid</i>	
		$logic^{(1000)}$ $any^{(1000)}$... $any^{(1000)}$	(999)
<i>logic</i>	=	(<i>logic</i>)	
		$logic^{(4)} :: type$	(3)
		<i>id</i> <i>longid</i> <i>var</i> <i>_</i> ...	
		CONST <i>id</i> CONST <i>longid</i>	
		$logic^{(1000)}$ $any^{(1000)}$... $any^{(1000)}$	(999)
		% <i>pttrns</i> . $any^{(3)}$	(3)
		$\lambda pttrns . any^{(3)}$	(3)
		TYPE (<i>type</i>)	
<i>idt</i>	=	(<i>idt</i>) <i>id</i> <i>_</i>	
		<i>id</i> :: <i>type</i>	(0)
		<i>_</i> :: <i>type</i>	(0)
<i>idts</i>	=	<i>idt</i> $idt^{(1)}$ <i>idts</i>	(0)
<i>pttrn</i>	=	<i>idt</i>	
<i>pttrns</i>	=	<i>pttrn</i> $pttrn^{(1)}$ <i>pttrns</i>	(0)

$$\begin{aligned}
type &= (type) \\
&| tid \mid tvar \mid _ \\
&| tid :: sort \mid tvar :: sort \mid _ :: sort \\
&| id \mid type^{(1000)} id \mid (type , \dots , type) id \\
&| longid \mid type^{(1000)} longid \\
&| (type , \dots , type) longid \\
&| type^{(1)} \Rightarrow type & (0) \\
&| type^{(1)} \Rightarrow type & (0) \\
&| [type , \dots , type] \Rightarrow type & (0) \\
&| [type , \dots , type] \Rightarrow type & (0) \\
\\
sort &= id \mid longid \mid \{ \} \\
&| \{ (id \mid longid) , \dots , (id \mid longid) \}
\end{aligned}$$

Here literal terminals are printed **verbatim**; see also §7.5 for further token categories of the inner syntax. The meaning of the nonterminals defined by the above grammar is as follows:

any denotes any term.

prop denotes meta-level propositions, which are terms of type *prop*. The syntax of such formulae of the meta-logic is carefully distinguished from usual conventions for object-logics. In particular, plain λ -term notation is *not* recognized as *prop*.

aprop denotes atomic propositions, which are embedded into regular *prop* by means of an explicit **PROP** token.

Terms of type *prop* with non-constant head, e.g. a plain variable, are printed in this form. Constants that yield type *prop* are expected to provide their own concrete syntax; otherwise the printed version will appear like *logic* and cannot be parsed again as *prop*.

logic denotes arbitrary terms of a logical type, excluding type *prop*. This is the main syntactic category of object-logic entities, covering plain λ -term notation (variables, abstraction, application), plus anything defined by the user.

When specifying notation for logical entities, all logical types (excluding *prop*) are *collapsed* to this single category of *logic*.

idt denotes identifiers, possibly constrained by types.

idts denotes a sequence of *idt*. This is the most basic category for variables in iterated binders, such as λ or \bigwedge .

pitrn and *pitrns* denote patterns for abstraction, cases bindings etc. In Pure, these categories start as a merely copy of *idt* and *idts*, respectively. Object-logics may add additional productions for binding forms.

type denotes types of the meta-logic.

sort denotes meta-level sorts.

Here are some further explanations of certain syntax features.

- In *idts*, note that $x :: \text{nat } y$ is parsed as $x :: (\text{nat } y)$, treating y like a type constructor applied to *nat*. To avoid this interpretation, write $(x :: \text{nat}) \ y$ with explicit parentheses.
- Similarly, $x :: \text{nat } y :: \text{nat}$ is parsed as $x :: (\text{nat } y :: \text{nat})$. The correct form is $(x :: \text{nat}) (y :: \text{nat})$, or $(x :: \text{nat}) \ y :: \text{nat}$ if y is last in the sequence of identifiers.
- Type constraints for terms bind very weakly. For example, $x < y :: \text{nat}$ is normally parsed as $(x < y) :: \text{nat}$, unless $<$ has a very low priority, in which case the input is likely to be ambiguous. The correct form is $x < (y :: \text{nat})$.
- Constraints may be either written with two literal colons “: :” or the double-colon symbol `\<Colon>`, which actually looks exactly the same in some L^AT_EX styles.
- Dummy variables (written as underscore) may occur in different roles.

A type “_” or “_ :: *sort*” acts like an anonymous inference parameter, which is filled-in according to the most general type produced by the type-checking phase.

A bound “_” refers to a vacuous abstraction, where the body does not refer to the binding introduced here. As in the term $\lambda x _ . x$, which is α -equivalent to $\lambda x y . x$.

A free “_” refers to an implicit outer binding. Higher definitional packages usually allow forms like $f \ x _ = x$.

A schematic “ $_$ ” (within a term pattern, see §3.2.6) refers to an anonymous variable that is implicitly abstracted over its context of locally bound variables. For example, this allows pattern matching of $\{x. f\ x = g\ x\}$ against $\{x. _ = _\}$, or even $\{_ _ = _\}$ by using both bound and schematic dummies.

- The three literal dots “ \dots ” may be also written as ellipsis symbol `\<dots>`. In both cases this refers to a special schematic variable, which is bound in the context. This special term abbreviation works nicely with calculational reasoning (§6.5).

7.5 Lexical matters

The inner lexical syntax vaguely resembles the outer one (§3.1), but some details are different. There are two main categories of inner syntax tokens:

1. *delimiters* — the literal tokens occurring in productions of the given priority grammar (cf. §7.4.1);
2. *named tokens* — various categories of identifiers etc.

Delimiters override named tokens and may thus render certain identifiers inaccessible. Sometimes the logical context admits alternative ways to refer to the same entity, potentially via qualified names.

The categories for named tokens are defined once and for all as follows, reusing some categories of the outer token syntax (§3.1).

<i>id</i>	=	<i>ident</i>
<i>longid</i>	=	<i>longident</i>
<i>var</i>	=	<i>var</i>
<i>tid</i>	=	<i>typefree</i>
<i>tvar</i>	=	<i>typevar</i>
<i>num</i>	=	<i>nat</i> <i>-nat</i>
<i>float_token</i>	=	<i>nat.nat</i> <i>-nat.nat</i>
<i>xnum</i>	=	<i>#nat</i> <i>#-nat</i>
<i>xstr</i>	=	<i>'' ... ''</i>

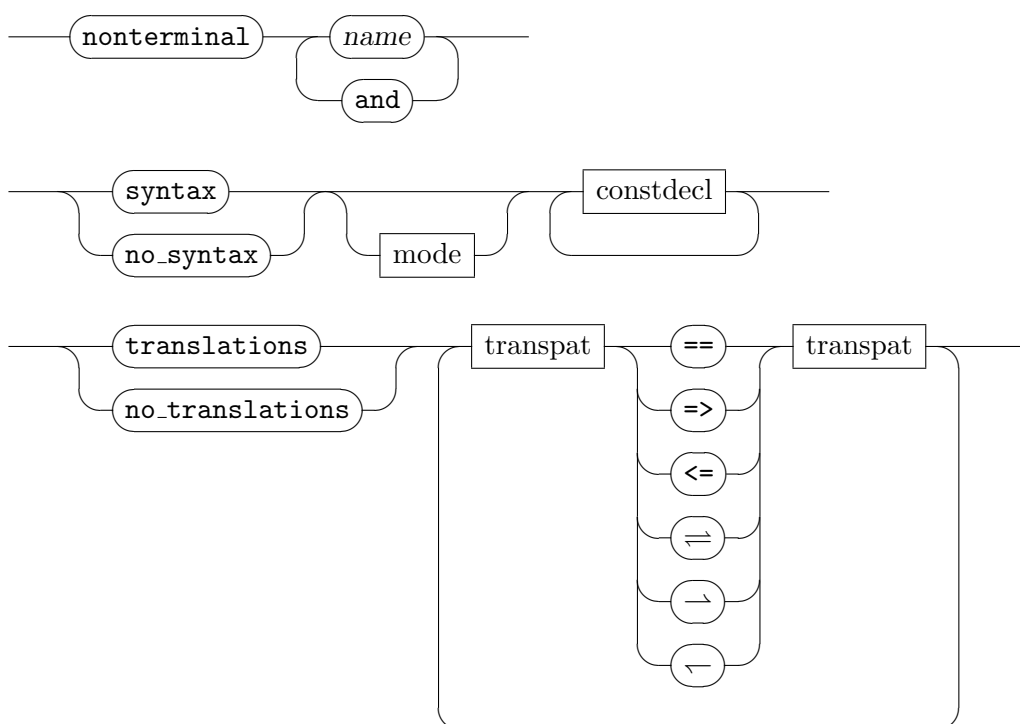
The token categories *num*, *float_token*, *xnum*, and *xstr* are not used in Pure. Object-logics may implement numerals and string constants by adding

appropriate syntax declarations, together with some translation functions (e.g. see Isabelle/HOL).

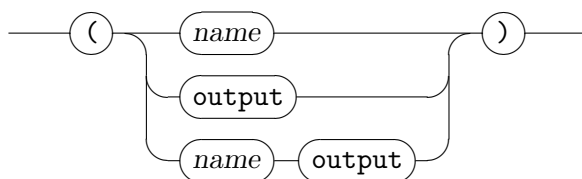
The derived categories *num_const* and *float_const* provide robust access to *num*, and *float_token*, respectively: the syntax tree holds a syntactic constant instead of a free variable.

7.6 Syntax and translations

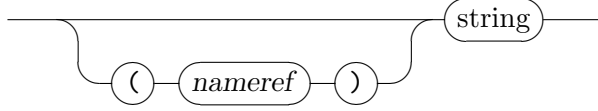
nonterminal : *theory* \rightarrow *theory*
syntax : *theory* \rightarrow *theory*
no_syntax : *theory* \rightarrow *theory*
translations : *theory* \rightarrow *theory*
no_translations : *theory* \rightarrow *theory*



mode



transpat



nonterminal *c* declares a type constructor *c* (without arguments) to act as purely syntactic type: a nonterminal symbol of the inner syntax.

syntax (*mode*) *decls* is similar to **consts** *decls*, except that the actual logical signature extension is omitted. Thus the context free grammar of Isabelle’s inner syntax may be augmented in arbitrary ways, independently of the logic. The *mode* argument refers to the print mode that the grammar rules belong; unless the **output** indicator is given, all productions are added both to the input and output grammar.

no_syntax (*mode*) *decls* removes grammar declarations (and translations) resulting from *decls*, which are interpreted in the same manner as for **syntax** above.

translations *rules* specifies syntactic translation rules (i.e. macros): parse / print rules (\rightleftharpoons), parse rules (\rightarrow), or print rules (\leftarrow). Translation patterns may be prefixed by the syntactic category to be used for parsing; the default is *logic*.

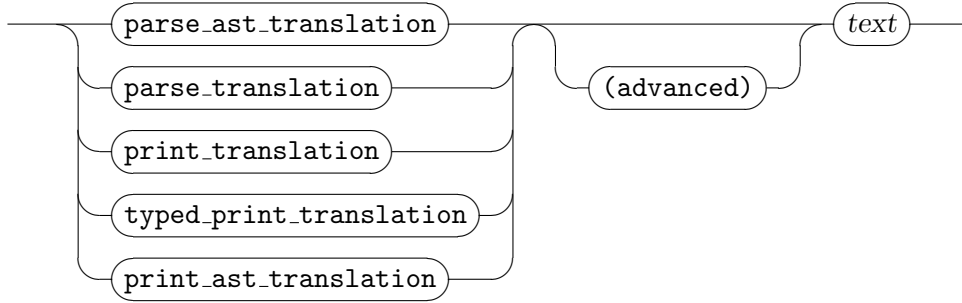
no_translations *rules* removes syntactic translation rules, which are interpreted in the same manner as for **translations** above.

7.7 Syntax translation functions

```

parse_ast_translation : theory  $\rightarrow$  theory
parse_translation    : theory  $\rightarrow$  theory
print_translation    : theory  $\rightarrow$  theory
typed_print_translation : theory  $\rightarrow$  theory
print_ast_translation : theory  $\rightarrow$  theory

```

Syntax translation functions written in ML admit almost arbitrary manipulations of Isabelle’s inner syntax. Any of the above commands have a single *text* argument that refers to an ML expression of appropriate type, which are as follows by default:

```
val parse_ast_translation  : (string * (ast list -> ast)) list
val parse_translation     : (string * (term list -> term)) list
val print_translation     : (string * (term list -> term)) list
val typed_print_translation :
  (string * (bool -> typ -> term list -> term)) list
val print_ast_translation : (string * (ast list -> ast)) list
```

If the (*advanced*) option is given, the corresponding translation functions may depend on the current theory or proof context. This allows to implement advanced syntax mechanisms, as translations functions may refer to specific theory declarations or auxiliary proof data.

See also [24] for more information on the general concept of syntax transformations in Isabelle.

```
val parse_ast_translation:
  (string * (Proof.context -> ast list -> ast)) list
val parse_translation:
  (string * (Proof.context -> term list -> term)) list
val print_translation:
  (string * (Proof.context -> term list -> term)) list
val typed_print_translation:
  (string * (Proof.context -> bool -> typ -> term list -> term)) list
val print_ast_translation:
  (string * (Proof.context -> ast list -> ast)) list
```

7.8 Inspecting the syntax

`print_syntax*` : *context* →

print_syntax prints the inner syntax of the current context. The output can be quite large; the most important sections are explained below.

lexicon lists the delimiters of the inner token language; see §7.5.

prods lists the productions of the underlying priority grammar; see §7.4.1.

The nonterminal $A^{(p)}$ is rendered in plain text as $A[p]$; delimiters are quoted. Many productions have an extra $\dots \Rightarrow name$. These names later become the heads of parse trees; they also guide the pretty printer.

Productions without such parse tree names are called *copy productions*. Their right-hand side must have exactly one nonterminal symbol (or named token). The parser does not create a new parse tree node for copy productions, but simply returns the parse tree of the right-hand symbol.

If the right-hand side of a copy production consists of a single nonterminal without any delimiters, then it is called a *chain production*. Chain productions act as abbreviations: conceptually, they are removed from the grammar by adding new productions. Priority information attached to chain productions is ignored; only the dummy value -1 is displayed.

print_modes lists the alternative print modes provided by this grammar; see §??.

parse_rules and *print_rules* relate to syntax translations (macros); see §7.6.

parse_ast_translation and *print_ast_translation* list sets of constants that invoke translation functions for abstract syntax trees, which are only required in very special situations; see §7.7.

parse_translation and *print_translation* list the sets of constants that invoke regular translation functions; see §7.7.

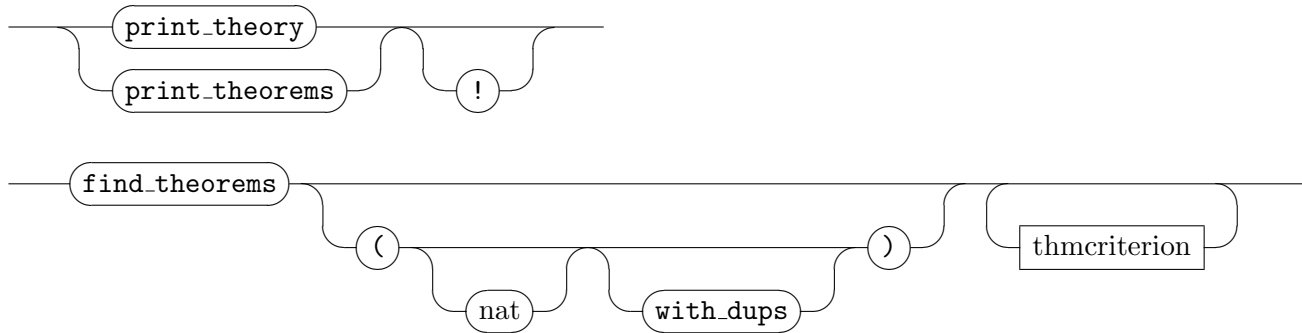
Other commands

8.1 Inspecting the context

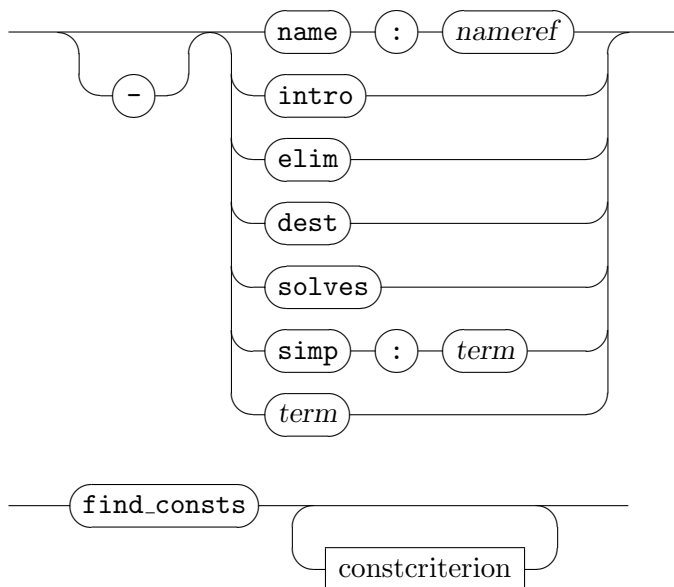
```

print_commands* : any →
  print_theory* : context →
  print_methods* : context →
  print_attributes* : context →
  print_theorems* : context →
  find_theorems* : context →
  find_consts* : context →
  thm_deps* : context →
  unused_thms* : context →
  print_facts* : context →
  print_binds* : context →

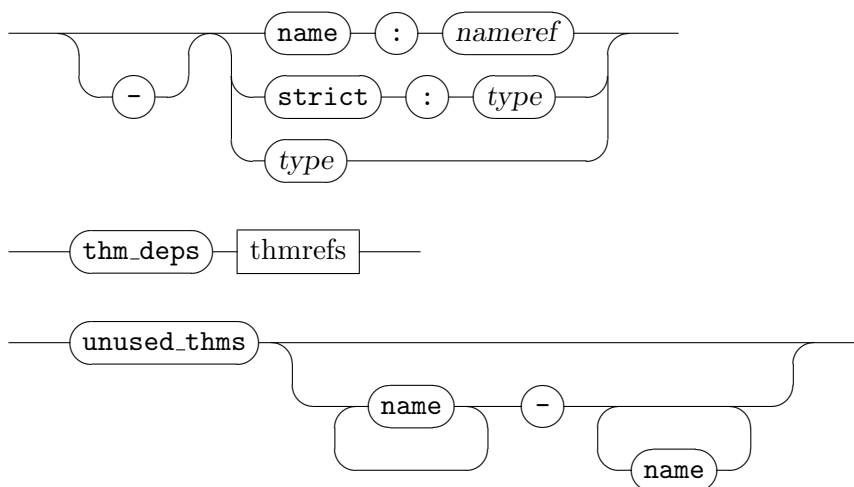
```



thm criterion



const criterion



These commands print certain parts of the theory and proof context. Note that there are some further ones available, such as for the set of rules declared for simplifications.

print_commands prints Isabelle's outer theory syntax, including keywords and command.

print_theory prints the main logical content of the theory context; the "!" option indicates extra verbosity.

print_methods prints all proof methods available in the current theory context.

print_attributes prints all attributes available in the current theory context.

print_theorems prints theorems resulting from the last command; the “!” option indicates extra verbosity.

find_theorems *criteria* retrieves facts from the theory or proof context matching all of given search criteria. The criterion *name: p* selects all theorems whose fully qualified name matches pattern *p*, which may contain “*” wildcards. The criteria *intro*, *elim*, and *dest* select theorems that match the current goal as introduction, elimination or destruction rules, respectively. The criterion *solves* returns all rules that would directly solve the current goal. The criterion *simp: t* selects all rewrite rules whose left-hand side matches the given term. The criterion *term t* selects all theorems that contain the pattern *t* – as usual, patterns may contain occurrences of the dummy “_”, schematic variables, and type constraints.

Criteria can be preceded by “–” to select theorems that do *not* match. Note that giving the empty list of criteria yields *all* currently known facts. An optional limit for the number of printed facts may be given; the default is 40. By default, duplicates are removed from the search result. Use *with_dups* to display duplicates.

find_consts *criteria* prints all constants whose type meets all of the given criteria. The criterion *strict: ty* is met by any type that matches the type pattern *ty*. Patterns may contain both the dummy type “_” and sort constraints. The criterion *ty* is similar, but it also matches against subtypes. The criterion *name: p* and the prefix “–” function as described for **find_theorems**.

thm_deps $a_1 \dots a_n$ visualizes dependencies of facts, using Isabelle’s graph browser tool (see also [39]).

unused_thms $A_1 \dots A_m - B_1 \dots B_n$ displays all unused theorems in theories $B_1 \dots B_n$ or their parents, but not in $A_1 \dots A_m$ or their parents. If n is 0, the end of the range of theories defaults to the current theory. If no range is specified, only the unused theorems in the current theory are displayed.

print_facts prints all local facts of the current context, both named and unnamed ones.

print_binds prints all term abbreviations present in the context.

8.2 History commands

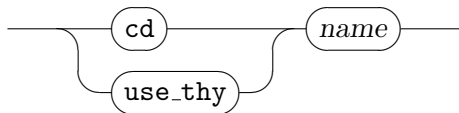
$\mathbf{undo}^{**} : any \rightarrow any$
 $\mathbf{linear_undo}^{**} : any \rightarrow any$
 $\mathbf{kill}^{**} : any \rightarrow any$

The Isabelle/Isar top-level maintains a two-stage history, for theory and proof state transformation. Basically, any command can be undone using **undo**, excluding mere diagnostic elements. Note that a theorem statement with a *finished* proof is treated as a single unit by **undo**. In contrast, the variant **linear_undo** admits to step back into the middle of a proof. The **kill** command aborts the current history node altogether, discontinuing a proof or even the whole theory. This operation is *not* undo-able.

! History commands should never be used with user interfaces such as Proof General [1, 2], which takes care of stepping forth and back itself. Interfering by manual **undo**, **linear_undo**, or even **kill** commands would quickly result in utter confusion.

8.3 System commands

$\mathbf{cd}^* : any \rightarrow$
 $\mathbf{pwd}^* : any \rightarrow$
 $\mathbf{use_thy}^* : any \rightarrow$



cd path changes the current directory of the Isabelle process.

pwd prints the current working directory.

use_thy *A* preload theory *A*. These system commands are scarcely used when working interactively, since loading of theories is done automatically as required.

Isabelle file specification may contain path variables (e.g. `$ISABELLE_HOME`) that are expanded accordingly. Note that `~` abbreviates `$HOME`, and `~~` abbreviates `$ISABELLE_HOME`. The general syntax for path specifications follows POSIX conventions.

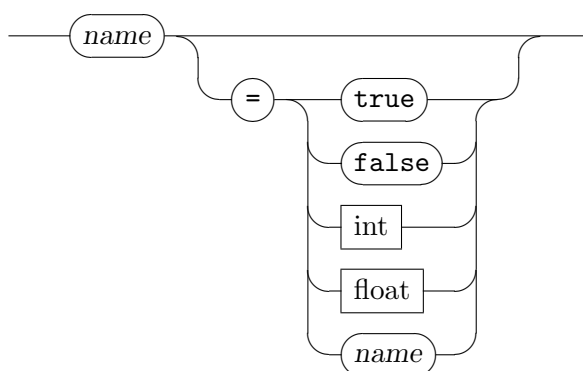
Generic tools and packages

9.1 Configuration options

Isabelle/Pure maintains a record of named configuration options within the theory or proof context, with values of type `bool`, `int`, `real`, or `string`. Tools may declare options in ML, and then refer to these values (relative to the context). Thus global reference variables are easily avoided. The user may change the value of a configuration option by means of an associated attribute of the same name. This form of context declaration works particularly well with commands such as **declare** or **using**.

For historical reasons, some tools cannot take the full proof context into account and merely refer to the background theory. This is accommodated by configuration options being declared as “global”, which may not be changed within a local context.

print_configs : *context* \rightarrow



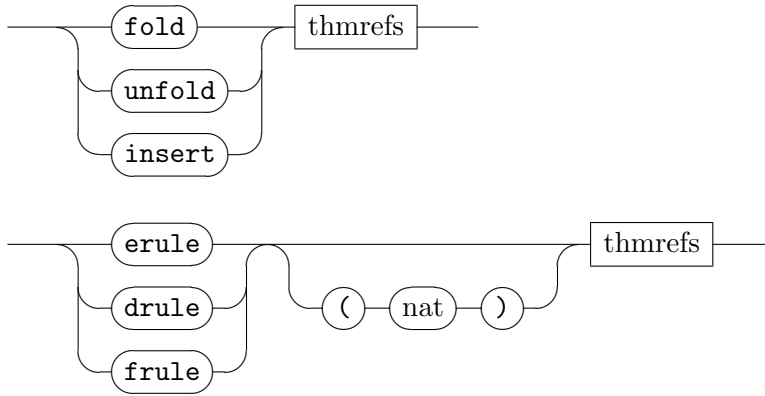
print_configs prints the available configuration options, with names, types, and current values.

name = *value* as an attribute expression modifies the named option, with the syntax of the value depending on the option’s type. For `bool` the default value is *true*. Any attempt to change a global option in a local context is ignored.

9.2 Basic proof tools

9.2.1 Miscellaneous methods and attributes

unfold : method
fold : method
insert : method
*erule** : method
*drule** : method
*frule** : method
succeed : method
fail : method



unfold $a_1 \dots a_n$ and *fold* $a_1 \dots a_n$ expand (or fold back) the given definitions throughout all goals; any chained facts provided are inserted into the goal and subject to rewriting as well.

insert $a_1 \dots a_n$ inserts theorems as facts into all goals of the proof state. Note that current facts indicated for forward chaining are ignored.

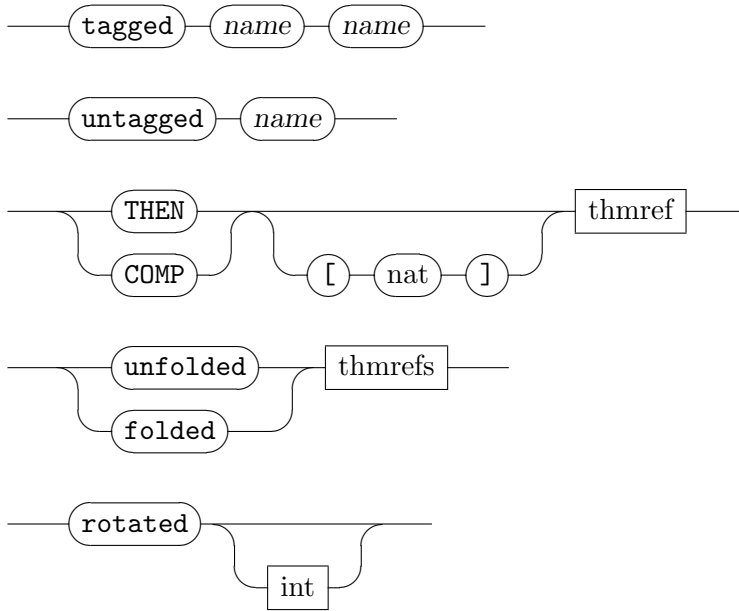
erule $a_1 \dots a_n$, *drule* $a_1 \dots a_n$, and *frule* $a_1 \dots a_n$ are similar to the basic *rule* method (see §6.3.3), but apply rules by elim-resolution, destruct-resolution, and forward-resolution, respectively [34]. The optional natural number argument (default 0) specifies additional assumption steps to be performed here.

Note that these methods are improper ones, mainly serving for experimentation and tactic script emulation. Different modes of basic rule application are usually expressed in Isar at the proof language level, rather than via implicit proof state manipulations. For example, a proper single-step elimination would be done using the plain *rule* method, with forward chaining of current facts.

succeed yields a single (unchanged) result; it is the identity of the “,” method combinator (cf. §6.3.1).

fail yields an empty result sequence; it is the identity of the “|” method combinator (cf. §6.3.1).

tagged : attribute
untagged : attribute
THEN : attribute
COMP : attribute
unfolded : attribute
folded : attribute
rotated : attribute
elim_format : attribute
*standard** : attribute
*no_vars** : attribute



tagged name value and *untagged name* add and remove *tags* of some theorem. Tags may be any list of string pairs that serve as formal comment. The first string is considered the tag name, the second its value. Note that *untagged* removes any tags of the same name.

THEN a and *COMP a* compose rules by resolution. *THEN* resolves with the first premise of *a* (an alternative position may be also specified);

the *COMP* version skips the automatic lifting process that is normally intended (cf. *op RS* and *op COMP* in [34]).

unfolded $a_1 \dots a_n$ and *folded* $a_1 \dots a_n$ expand and fold back again the given definitions throughout a rule.

rotated n rotate the premises of a theorem by n (default 1).

elim_format turns a destruction rule into elimination rule format, by resolving with the rule $PROP A \Rightarrow (PROP A \Rightarrow PROP B) \Rightarrow PROP B$.

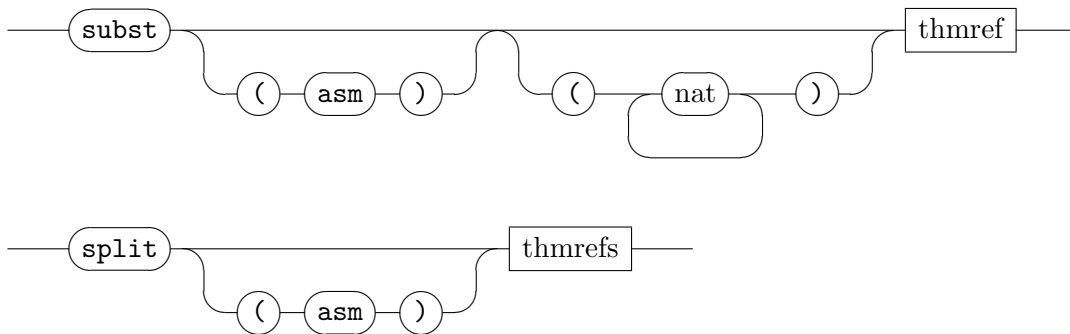
Note that the Classical Reasoner (§9.4) provides its own version of this operation.

standard puts a theorem into the standard form of object-rules at the outermost theory level. Note that this operation violates the local proof context (including active locales).

no_vars replaces schematic variables by free ones; this is mainly for tuning output of pretty printed theorems.

9.2.2 Low-level equational reasoning

subst : method
hypsubst : method
split : method



These methods provide low-level facilities for equational reasoning that are intended for specialized applications only. Normally, single step calculations would be performed in a structured text (see also §6.5), while the Simplifier methods provide the canonical way for automated normalization (see §9.3).

subst eq performs a single substitution step using rule *eq*, which may be either a meta or object equality.

subst (asm) eq substitutes in an assumption.

subst (i ... j) eq performs several substitutions in the conclusion. The numbers *i* to *j* indicate the positions to substitute at. Positions are ordered from the top of the term tree moving down from left to right. For example, in $(a + b) + (c + d)$ there are three positions where commutativity of $+$ is applicable: 1 refers to $a + b$, 2 to the whole term, and 3 to $c + d$.

If the positions in the list $(i \dots j)$ are non-overlapping (e.g. (2 3) in $(a + b) + (c + d)$) you may assume all substitutions are performed simultaneously. Otherwise the behaviour of *subst* is not specified.

subst (asm) (i ... j) eq performs the substitutions in the assumptions. The positions refer to the assumptions in order from left to right. For example, given in a goal of the form $P(a + b) \implies P(c + d) \implies \dots$, position 1 of commutativity of $+$ is the subterm $a + b$ and position 2 is the subterm $c + d$.

hypsubst performs substitution using some assumption; this only works for equations of the form $x = t$ where x is a free or bound variable.

split a₁ ... a_n performs single-step case splitting using the given rules. By default, splitting is performed in the conclusion of a goal; the *(asm)* option indicates to operate on assumptions instead.

Note that the *simp* method already involves repeated application of split rules as declared in the current context.

9.2.3 Further tactic emulations

The following improper proof methods emulate traditional tactics. These admit direct access to the goal state, which is normally considered harmful! In particular, this may involve both numbered goal addressing (default 1), and dynamic instantiation within the scope of some subgoal.

- ! Dynamic instantiations refer to universally quantified parameters of a subgoal
- (the dynamic context) rather than fixed variables and term abbreviations of a (static) Isar context.

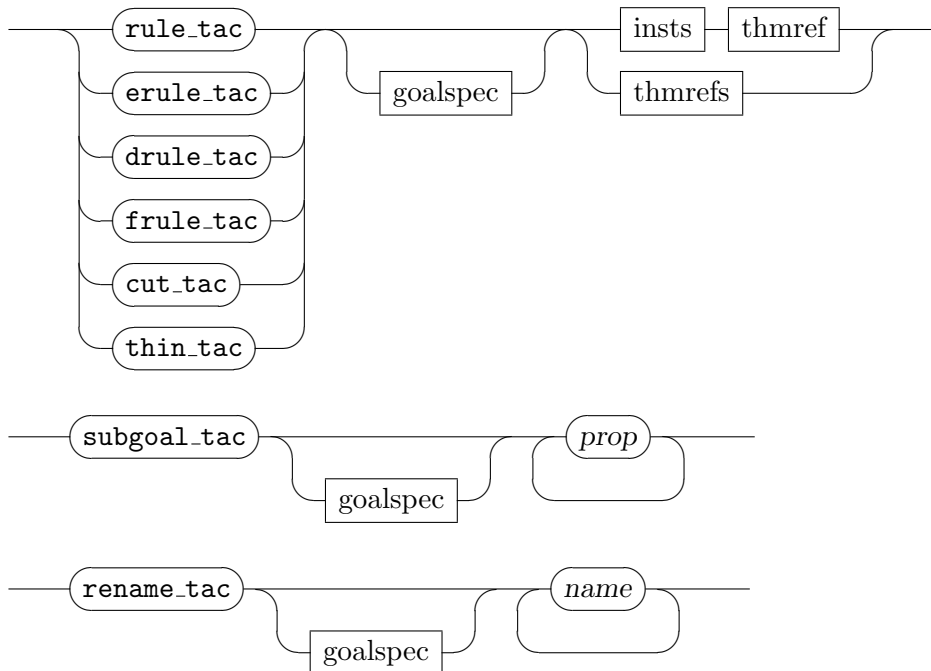
Tactic emulation methods, unlike their ML counterparts, admit simultaneous instantiation from both dynamic and static contexts. If names occur in both contexts goal parameters hide locally fixed variables. Likewise, schematic variables refer to term abbreviations, if present in the static context. Otherwise the schematic variable is interpreted as a schematic variable and left to be solved by unification with certain parts of the subgoal.

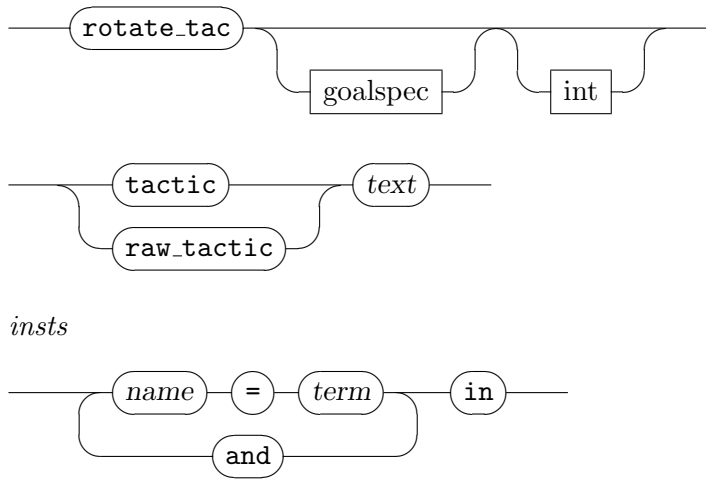
Note that the tactic emulation proof methods in Isabelle/Isar are consistently named *foo_tac*. Note also that variable names occurring on left hand sides of instantiations must be preceded by a question mark if they coincide with a keyword or contain dots. This is consistent with the attribute *where* (see §6.3.3).

```

    rule_tac*   : method
    erule_tac*  : method
    drule_tac*  : method
    frule_tac*  : method
    cut_tac*    : method
    thin_tac*   : method
    subgoal_tac* : method
    rename_tac* : method
    rotate_tac* : method
    tactic*     : method
    raw_tactic* : method

```





rule_tac etc. do resolution of rules with explicit instantiation. This works the same way as the ML tactics `res_inst_tac` etc. (see [34])

Multiple rules may be only given if there is no instantiation; then *rule_tac* is the same as `resolve_tac` in ML (see [34]).

cut_tac inserts facts into the proof state as assumption of a subgoal, see also `Tactic.cut_facts_tac` in [34]. Note that the scope of schematic variables is spread over the main goal statement. Instantiations may be given as well, see also ML tactic `cut_inst_tac` in [34].

thin_tac φ deletes the specified assumption from a subgoal; note that φ may contain schematic variables. See also `thin_tac` in [34].

subgoal_tac φ adds φ as an assumption to a subgoal. See also `subgoal_tac` and `subgoals_tac` in [34].

rename_tac $x_1 \dots x_n$ renames parameters of a goal according to the list x_1, \dots, x_n , which refers to the *suffix* of variables.

rotate_tac n rotates the assumptions of a goal by n positions: from right to left if n is positive, and from left to right if n is negative; the default value is 1. See also `rotate_tac` in [34].

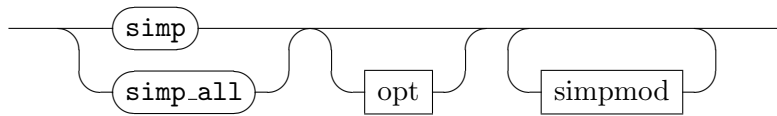
tactic text produces a proof method from any ML text of type `tactic`. Apart from the usual ML environment and the current proof context, the ML code may refer to the locally bound values `facts`, which indicates any current facts used for forward-chaining.

raw_tactic is similar to *tactic*, but presents the goal state in its raw internal form, where simultaneous subgoals appear as conjunction of the logical framework instead of the usual split into several subgoals. While feature this is useful for debugging of complex method definitions, it should not never appear in production theories.

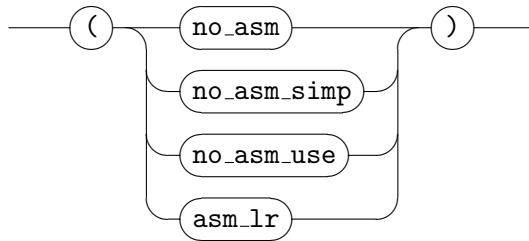
9.3 The Simplifier

9.3.1 Simplification methods

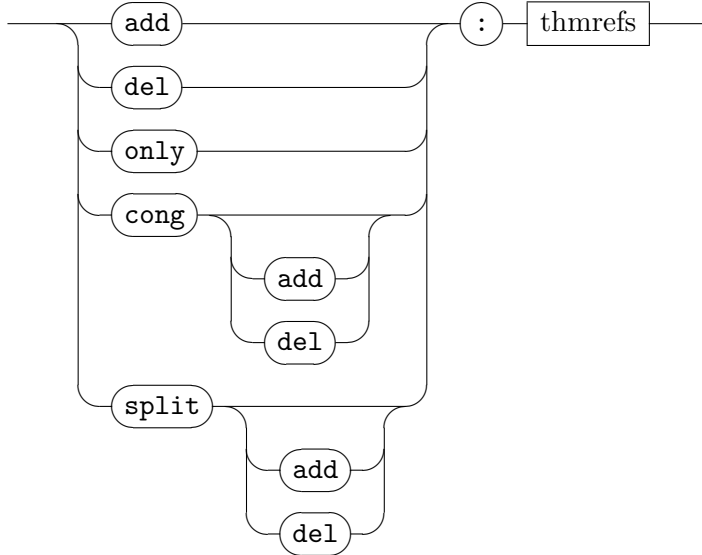
simp : method
simp_all : method



opt



simpmod



simp invokes the Simplifier, after declaring additional rules according to the arguments given. Note that the **only** modifier first removes all other rewrite rules, congruences, and looper tactics (including splits), and then behaves like **add**.

The **cong** modifiers add or delete Simplifier congruence rules (see also [24]), the default is to add.

The **split** modifiers add or delete rules for the Splitter (see also [24]), the default is to add. This works only if the Simplifier method has been properly setup to include the Splitter (all major object logics such HOL, HOLCF, FOL, ZF do this already).

simp_all is similar to *simp*, but acts on all goals (backwards from the last to the first one).

By default the Simplifier methods take local assumptions fully into account, using equational assumptions in the subsequent normalization process, or simplifying assumptions themselves (cf. **asm_full_simp_tac** in [24]). In structured proofs this is usually quite well behaved in practice: just the local premises of the actual goal are involved, additional facts may be inserted via explicit forward-chaining (via **then**, **from**, **using** etc.).

Additional Simplifier options may be specified to tune the behavior further (mostly for unstructured scripts with many accidental local facts): “(*no_asm*)” means assumptions are ignored completely (cf. **simp_tac**), “(*no_asm_simp*)” means assumptions are used in the simplification of the conclusion but

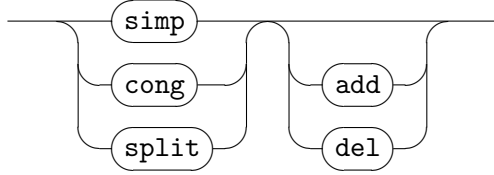
are not themselves simplified (cf. `asm_simp_tac`), and “*(no_asm_use)*” means assumptions are simplified but are not used in the simplification of each other or the conclusion (cf. `full_simp_tac`). For compatibility reasons, there is also an option “*(asm_lr)*”, which means that an assumption is only used for simplifying assumptions which are to the right of it (cf. `asm_lr_simp_tac`).

The configuration option *depth_limit* limits the number of recursive invocations of the simplifier during conditional rewriting.

The Splitter package is usually configured to work as part of the Simplifier. The effect of repeatedly applying `split_tac` can be simulated by “*(simp only: split: a₁ ... a_n)*”. There is also a separate *split* method available for single-step case splitting.

9.3.2 Declaring rules

```
print_simpset* : context →
    simp : attribute
    cong : attribute
    split : attribute
```



`print_simpset` prints the collection of rules declared to the Simplifier, which is also known as “simpset” internally [24].

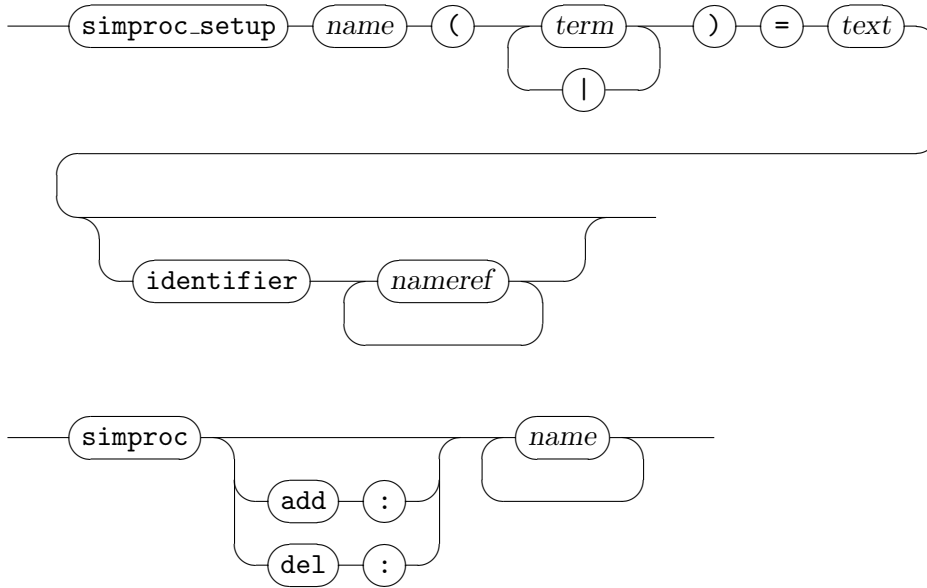
simp declares simplification rules.

cong declares congruence rules.

split declares case split rules.

9.3.3 Simplification procedures

```
simproc_setup : local_theory → local_theory
    simproc : attribute
```



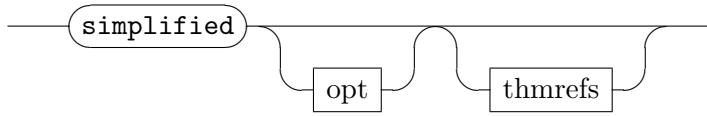
simproc_setup defines a named simplification procedure that is invoked by the Simplifier whenever any of the given term patterns match the current redex. The implementation, which is provided as ML source text, needs to be of type `morphism -> simpset -> cterm -> thm option`, where the `cterm` represents the current redex r and the result is supposed to be some proven rewrite rule $r \equiv r'$ (or a generalized version), or `NONE` to indicate failure. The `simpset` argument holds the full context of the current Simplifier invocation, including the actual Isar proof context. The `morphism` informs about the difference of the original compilation context wrt. the one of the actual application later on. The optional **identifier** specifies theorems that represent the logical content of the abstract theory of this simproc.

Morphisms and identifiers are only relevant for simprocs that are defined within a local target context, e.g. in a locale.

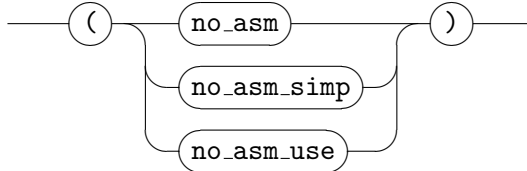
simproc add: name and *simproc del: name* add or delete named simprocs to the current Simplifier context. The default is to add a simproc. Note that **simproc_setup** already adds the new simproc to the subsequent context.

9.3.4 Forward simplification

simplified : *attribute*



opt



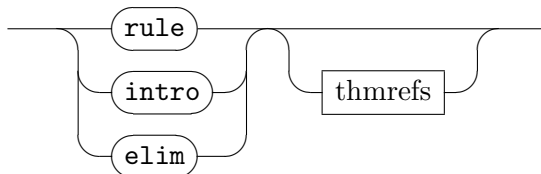
simplified $a_1 \dots a_n$ causes a theorem to be simplified, either by exactly the specified rules a_1, \dots, a_n , or the implicit Simplifier context if no arguments are given. The result is fully simplified by default, including assumptions and conclusion; the options *no_asm* etc. tune the Simplifier in the same way as the for the *simp* method.

Note that forward simplification restricts the simplifier to its most basic operation of term rewriting; solver and loopier tactics [24] are *not* involved here. The *simplified* attribute should be only rarely required under normal circumstances.

9.4 The Classical Reasoner

9.4.1 Basic methods

rule : method
contradiction : method
intro : method
elim : method



rule as offered by the Classical Reasoner is a refinement over the primitive one (see §6.3.3). Both versions essentially work the same, but the classical version observes the classical rule context in addition to that of Isabelle/Pure.

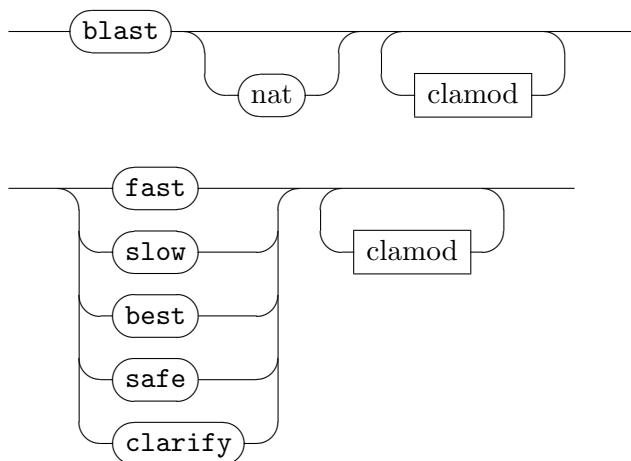
Common object logics (HOL, ZF, etc.) declare a rich collection of classical rules (even if these would qualify as intuitionistic ones), but only few declarations to the rule context of Isabelle/Pure (§6.3.3).

contradiction solves some goal by contradiction, deriving any result from both $\neg A$ and A . Chained facts, which are guaranteed to participate, may appear in either order.

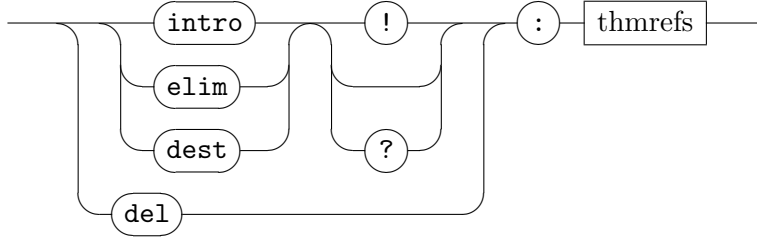
intro and *elim* repeatedly refine some goal by intro- or elim-resolution, after having inserted any chained facts. Exactly the rules given as arguments are taken into account; this allows fine-tuned decomposition of a proof problem, in contrast to common automated tools.

9.4.2 Automated methods

blast : method
fast : method
slow : method
best : method
safe : method
clarify : method



clamod



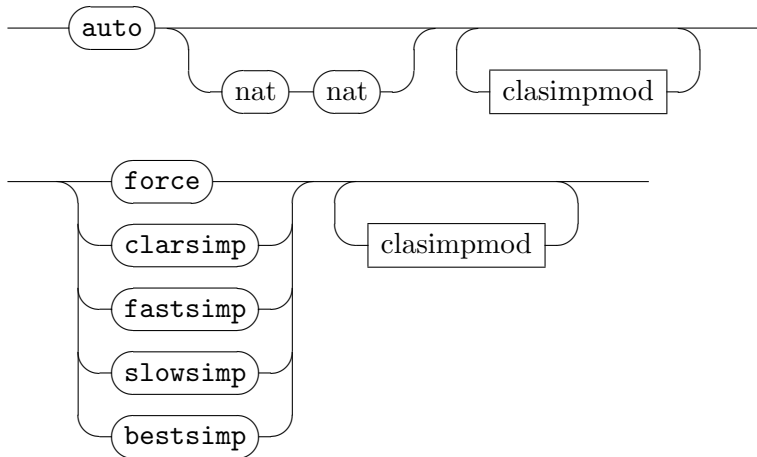
blast refers to the classical tableau prover (see `blast_tac` in [24]). The optional argument specifies a user-supplied search bound (default 20).

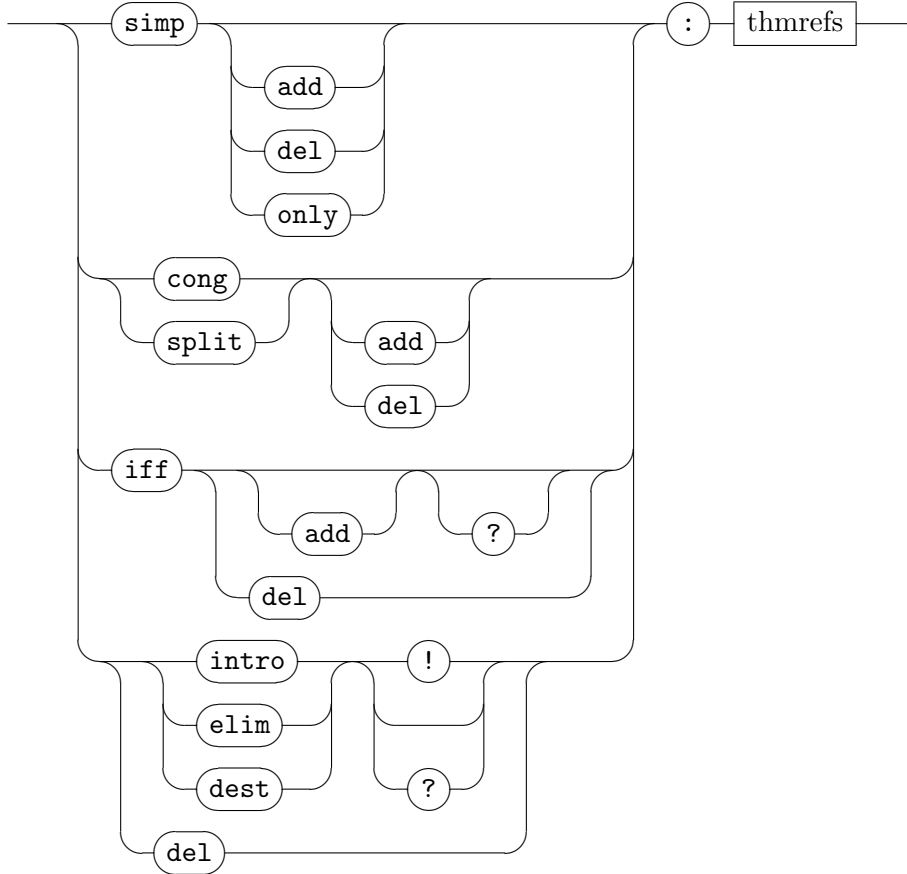
fast, *slow*, *best*, *safe*, and *clarify* refer to the generic classical reasoner. See `fast_tac`, `slow_tac`, `best_tac`, `safe_tac`, and `clarify_tac` in [24] for more information.

Any of the above methods support additional modifiers of the context of classical rules. Their semantics is analogous to the attributes given before. Facts provided by forward chaining are inserted into the goal before commencing proof search.

9.4.3 Combined automated methods

auto : method
force : method
clarsimp : method
fastsimp : method
slowsimp : method
bestsimp : method



clasimpmod

auto, *force*, *clarsimp*, *fastsimp*, *slowsimp*, and *bestsimp* provide access to Isabelle's combined simplification and classical reasoning tactics. These correspond to *auto_tac*, *force_tac*, *clarsimp_tac*, and Classical Reasoner tactics with the Simplifier added as wrapper, see [24] for more information. The modifier arguments correspond to those given in §9.3 and §9.4. Just note that the ones related to the Simplifier are prefixed by **simp** here.

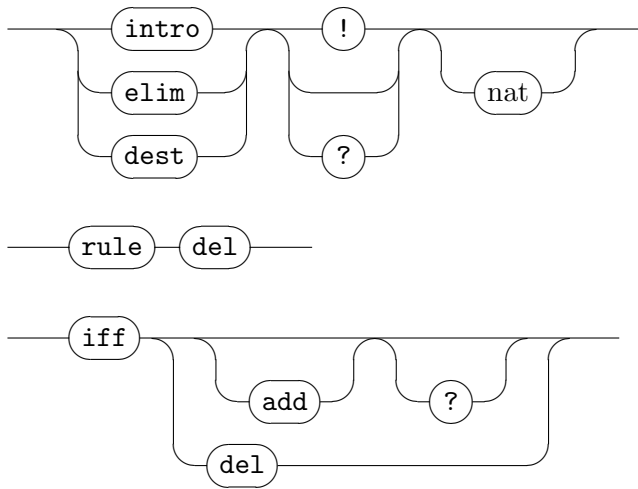
Facts provided by forward chaining are inserted into the goal before doing the search.

9.4.4 Declaring rules

```

print_claset* : context →
    intro : attribute
    elim  : attribute
    dest  : attribute
    rule  : attribute
    iff   : attribute

```



print_claset prints the collection of rules declared to the Classical Reasoner, which is also known as “claset” internally [24].

intro, *elim*, and *dest* declare introduction, elimination, and destruction rules, respectively. By default, rules are considered as *unsafe* (i.e. not applied blindly without backtracking), while “!” classifies as *safe*. Rule declarations marked by “?” coincide with those of Isabelle/Pure, cf. §6.3.3 (i.e. are only applied in single steps of the *rule* method). The optional natural number specifies an explicit weight argument, which is ignored by automated tools, but determines the search order of single rule steps.

rule del deletes introduction, elimination, or destruction rules from the context.

iff declares logical equivalences to the Simplifier and the Classical reasoner at the same time. Non-conditional rules result in a “safe” introduction and elimination pair; conditional ones are considered “unsafe”. Rules with negative conclusion are automatically inverted (using \neg -elimination internally).

The “?” version of *iff* declares rules to the Isabelle/Pure context only, and omits the Simplifier declaration.

9.4.5 Classical operations

swapped : *attribute*

swapped turns an introduction rule into an elimination, by resolving with the classical swap principle $(\neg B \implies A) \implies (\neg A \implies B)$.

9.5 Object-logic setup

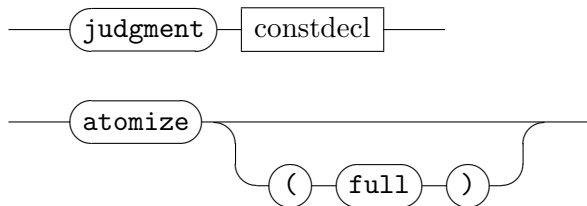
judgment : *theory* \rightarrow *theory*
atomize : *method*
atomize : *attribute*
rule_format : *attribute*
rulify : *attribute*

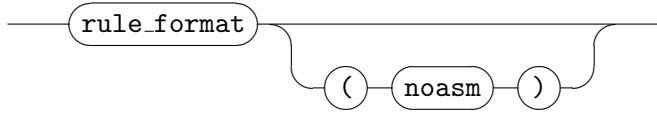
The very starting point for any Isabelle object-logic is a “truth judgment” that links object-level statements to the meta-logic (with its minimal language of *prop* that covers universal quantification \bigwedge and implication \implies).

Common object-logics are sufficiently expressive to internalize rule statements over \bigwedge and \implies within their own language. This is useful in certain situations where a rule needs to be viewed as an atomic statement from the meta-level perspective, e.g. $\bigwedge x. x \in A \implies P x$ versus $\forall x \in A. P x$.

From the following language elements, only the *atomize* method and *rule_format* attribute are occasionally required by end-users, the rest is for those who need to setup their own object-logic. In the latter case existing formulations of Isabelle/FOL or Isabelle/HOL may be taken as realistic examples.

Generic tools may refer to the information provided by object-logic declarations internally.





judgment $c :: \sigma \ (mx)$ declares constant c as the truth judgment of the current object-logic. Its type σ should specify a coercion of the category of object-level propositions to *prop* of the Pure meta-logic; the mixfix annotation (mx) would typically just link the object language (internally of syntactic category *logic*) with that of *prop*. Only one **judgment** declaration may be given in any theory development.

atomize (as a method) rewrites any non-atomic premises of a sub-goal, using the meta-level equations declared via *atomize* (as an attribute) beforehand. As a result, heavily nested goals become amenable to fundamental operations such as resolution (cf. the *rule* method). Giving the “(full)” option here means to turn the whole subgoal into an object-statement (if possible), including the outermost parameters and assumptions as well.

A typical collection of *atomize* rules for a particular object-logic would provide an internalization for each of the connectives of \bigwedge , \implies , and \equiv . Meta-level conjunction should be covered as well (this is particularly important for locales, see §5.5).

rule_format rewrites a theorem by the equalities declared as *rulify* rules in the current object-logic. By default, the result is fully normalized, including assumptions and conclusions at any depth. The *(no_asm)* option restricts the transformation to the conclusion of a rule.

In common object-logics (HOL, FOL, ZF), the effect of *rule_format* is to replace (bounded) universal quantification (\forall) and implication (\longrightarrow) by the corresponding rule statements over \bigwedge and \implies .

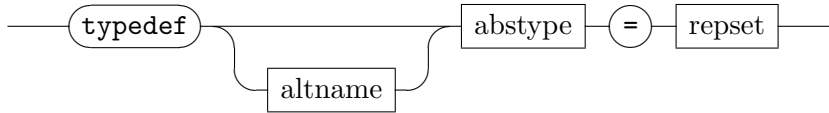
Part III

Object-Logics

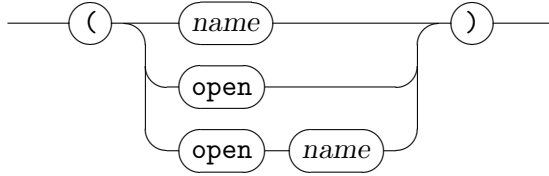
Isabelle/HOL

10.1 Typedef axiomatization

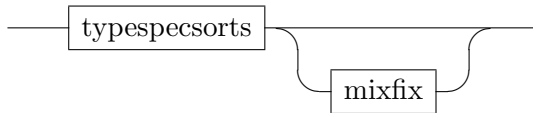
typedef : *local_theory* \rightarrow *proof(prove)*



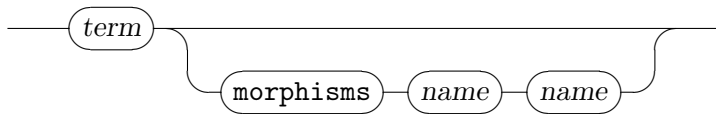
altname



abstype



reset



typedef $(\alpha_1, \dots, \alpha_n) \ t = A$ axiomatizes a Gordon/HOL-style type definition in the background theory of the current context, depending on a non-emptiness result of the set A (which needs to be proven interactively).

The raw type may not depend on parameters or assumptions of the context — this is logically impossible in Isabelle/HOL — but the non-emptiness property can be local, potentially resulting in multiple interpretations in target contexts. Thus the established bijection between

the representing set A and the new type t may semantically depend on local assumptions.

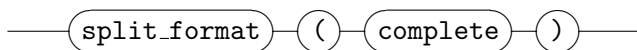
By default, **typedef** defines both a type t and a set (term constant) of the same name, unless an alternative base name is given in parentheses, or the “(*open*)” declaration is used to suppress a separate constant definition altogether. The injection from type to set is called Rep_t , its inverse Abs_t — this may be changed via an explicit **morphisms** declaration.

Theorems Rep_t , $Rep_t_inverse$, and $Abs_t_inverse$ provide the most basic characterization as a corresponding injection/surjection pair (in both directions). Rules Rep_t_inject and Abs_t_inject provide a slightly more convenient view on the injectivity part, suitable for automated proof tools (e.g. in *simp* or *iff* declarations). Rules Rep_t_cases/Rep_t_induct , and Abs_t_cases/Abs_t_induct provide alternative views on surjectivity; these are already declared as set or type rules for the generic *cases* and *induct* methods.

An alternative name for the set definition (and other derived entities) may be specified in parentheses; the default is to use t as indicated before.

10.2 Adhoc tuples

*split_format** : *attribute*



split_format (*complete*) causes arguments in function applications to be represented canonically according to their tuple type structure.

Note that this operation tends to invent funny names for new local parameters introduced.

10.3 Records

In principle, records merely generalize the concept of tuples, where components may be addressed by labels instead of just position. The logical

infrastructure of records in Isabelle/HOL is slightly more advanced, though, supporting truly extensible record schemes. This admits operations that are polymorphic with respect to record extension, yielding “object-oriented” effects like (single) inheritance. See also [16] for more details on object-oriented verification and record subtyping in HOL.

10.3.1 Basic concepts

Isabelle/HOL supports both *fixed* and *schematic* records at the level of terms and types. The notation is as follows:

	record terms	record types
fixed	$\langle x = a, y = b \rangle$	$\langle x :: A, y :: B \rangle$
schematic	$\langle x = a, y = b, \dots = m \rangle$	$\langle x :: A, y :: B, \dots :: M \rangle$

The ASCII representation of $\langle x = a \rangle$ is $(| x = a |)$.

A fixed record $\langle x = a, y = b \rangle$ has field x of value a and field y of value b . The corresponding type is $\langle x :: A, y :: B \rangle$, assuming that $a :: A$ and $b :: B$.

A record scheme like $\langle x = a, y = b, \dots = m \rangle$ contains fields x and y as before, but also possibly further fields as indicated by the “...” notation (which is actually part of the syntax). The improper field “...” of a record scheme is called the *more part*. Logically it is just a free variable, which is occasionally referred to as “row variable” in the literature. The more part of a record scheme may be instantiated by zero or more further components. For example, the previous scheme may get instantiated to $\langle x = a, y = b, z = c, \dots = m' \rangle$, where m' refers to a different more part. Fixed records are special instances of record schemes, where “...” is properly terminated by the $() :: \text{unit}$ element. In fact, $\langle x = a, y = b \rangle$ is just an abbreviation for $\langle x = a, y = b, \dots = () \rangle$.

Two key observations make extensible records in a simply typed language like HOL work out:

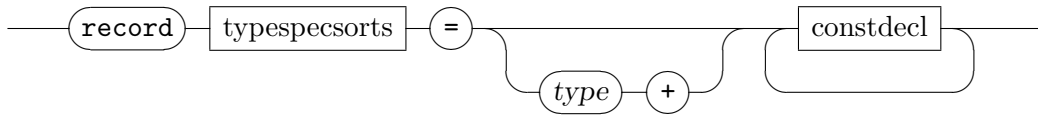
1. the more part is internalized, as a free term or type variable,
2. field names are externalized, they cannot be accessed within the logic as first-class values.

In Isabelle/HOL record types have to be defined explicitly, fixing their field names and types, and their (optional) parent record. Afterwards, records may be formed using above syntax, while obeying the canonical

order of fields as given by their declaration. The record package provides several standard operations like selectors and updates. The common setup for various generic proof tools enable succinct reasoning patterns. See also the Isabelle/HOL tutorial [19] for further instructions on using records in practice.

10.3.2 Record specifications

record : *theory* \rightarrow *theory*



record $(\alpha_1, \dots, \alpha_m) \ t = \tau + c_1 :: \sigma_1 \dots c_n :: \sigma_n$ defines extensible record type $(\alpha_1, \dots, \alpha_m) \ t$, derived from the optional parent record τ by adding new field components $c_i :: \sigma_i$ etc.

The type variables of τ and σ_i need to be covered by the (distinct) parameters $\alpha_1, \dots, \alpha_m$. Type constructor t has to be new, while τ needs to specify an instance of an existing record type. At least one new field c_i has to be specified. Basically, field names need to belong to a unique record. This is not a real restriction in practice, since fields are qualified by the record name internally.

The parent record specification τ is optional; if omitted t becomes a root record. The hierarchy of all records declared within a theory context forms a forest structure, i.e. a set of trees starting with a root record each. There is no way to merge multiple parent records!

For convenience, $(\alpha_1, \dots, \alpha_m) \ t$ is made a type abbreviation for the fixed record type $\langle c_1 :: \sigma_1, \dots, c_n :: \sigma_n \rangle$, likewise is $(\alpha_1, \dots, \alpha_m, \zeta) \ t_scheme$ made an abbreviation for $\langle c_1 :: \sigma_1, \dots, c_n :: \sigma_n, \dots :: \zeta \rangle$.

10.3.3 Record operations

Any record definition of the form presented above produces certain standard operations. Selectors and updates are provided for any field, including the improper one “*more*”. There are also cumulative record constructor functions. To simplify the presentation below, we assume for now that $(\alpha_1, \dots, \alpha_m) \ t$ is a root record with fields $c_1 :: \sigma_1, \dots, c_n :: \sigma_n$.

Selectors and **updates** are available for any field (including “more”):

$$\begin{aligned} c_i &:: \langle \bar{c} :: \bar{\sigma}, \dots :: \zeta \rangle \Rightarrow \sigma_i \\ c_i_update &:: \sigma_i \Rightarrow \langle \bar{c} :: \bar{\sigma}, \dots :: \zeta \rangle \Rightarrow \langle \bar{c} :: \bar{\sigma}, \dots :: \zeta \rangle \end{aligned}$$

There is special syntax for application of updates: $r(x := a)$ abbreviates term $x_update\ a\ r$. Further notation for repeated updates is also available: $r(x := a)(y := b)(z := c)$ may be written $r(x := a, y := b, z := c)$. Note that because of postfix notation the order of fields shown here is reverse than in the actual term. Since repeated updates are just function applications, fields may be freely permuted in $(x := a, y := b, z := c)$, as far as logical equality is concerned. Thus commutativity of independent updates can be proven within the logic for any two fields, but not as a general theorem.

The **make** operation provides a cumulative record constructor function:

$$t.make :: \sigma_1 \Rightarrow \dots \sigma_n \Rightarrow \langle \bar{c} :: \bar{\sigma} \rangle$$

We now reconsider the case of non-root records, which are derived of some parent. In general, the latter may depend on another parent as well, resulting in a list of *ancestor records*. Appending the lists of fields of all ancestors results in a certain field prefix. The record package automatically takes care of this by lifting operations over this context of ancestor fields. Assuming that $(\alpha_1, \dots, \alpha_m)$ t has ancestor fields $b_1 :: \varrho_1, \dots, b_k :: \varrho_k$, the above record operations will get the following types:

$$\begin{aligned} c_i &:: \langle \bar{b} :: \bar{\varrho}, \bar{c} :: \bar{\sigma}, \dots :: \zeta \rangle \Rightarrow \sigma_i \\ c_i_update &:: \sigma_i \Rightarrow \langle \bar{b} :: \bar{\varrho}, \bar{c} :: \bar{\sigma}, \dots :: \zeta \rangle \Rightarrow \langle \bar{b} :: \bar{\varrho}, \bar{c} :: \bar{\sigma}, \dots :: \zeta \rangle \\ t.make &:: \varrho_1 \Rightarrow \dots \varrho_k \Rightarrow \sigma_1 \Rightarrow \dots \sigma_n \Rightarrow \langle \bar{b} :: \bar{\varrho}, \bar{c} :: \bar{\sigma} \rangle \end{aligned}$$

Some further operations address the extension aspect of a derived record scheme specifically: $t.fields$ produces a record fragment consisting of exactly the new fields introduced here (the result may serve as a more part elsewhere); $t.extend$ takes a fixed record and adds a given more part; $t.truncate$ restricts a record scheme to a fixed record.

$$\begin{aligned} t.fields &:: \sigma_1 \Rightarrow \dots \sigma_n \Rightarrow \langle \bar{c} :: \bar{\sigma} \rangle \\ t.extend &:: \langle \bar{b} :: \bar{\varrho}, \bar{c} :: \bar{\sigma} \rangle \Rightarrow \zeta \Rightarrow \langle \bar{b} :: \bar{\varrho}, \bar{c} :: \bar{\sigma}, \dots :: \zeta \rangle \\ t.truncate &:: \langle \bar{b} :: \bar{\varrho}, \bar{c} :: \bar{\sigma}, \dots :: \zeta \rangle \Rightarrow \langle \bar{b} :: \bar{\varrho}, \bar{c} :: \bar{\sigma} \rangle \end{aligned}$$

Note that $t.make$ and $t.fields$ coincide for root records.

10.3.4 Derived rules and proof tools

The record package proves several results internally, declaring these facts to appropriate proof tools. This enables users to reason about record structures quite conveniently. Assume that t is a record type as specified above.

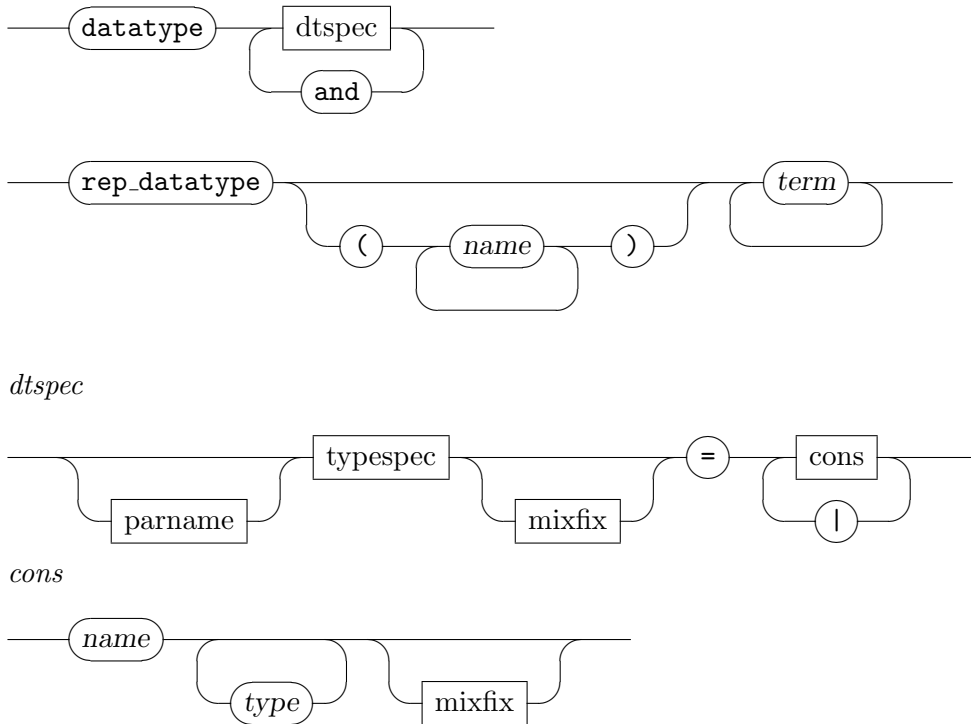
1. Standard conversions for selectors or updates applied to record constructor terms are made part of the default Simplifier context; thus proofs by reduction of basic operations merely require the *simp* method without further arguments. These rules are available as *t.simps*, too.
2. Selectors applied to updated records are automatically reduced by an internal simplification procedure, which is also part of the standard Simplifier setup.
3. Inject equations of a form analogous to $(x, y) = (x', y') \equiv x = x' \wedge y = y'$ are declared to the Simplifier and Classical Reasoner as *iff* rules. These rules are available as *t.iffs*.
4. The introduction rule for record equality analogous to $x\ r = x\ r' \implies y\ r = y\ r' \dots \implies r = r'$ is declared to the Simplifier, and as the basic rule context as “*intro?*”. The rule is called *t.equality*.
5. Representations of arbitrary record expressions as canonical constructor terms are provided both in *cases* and *induct* format (cf. the generic proof methods of the same name, §6.6). Several variations are available, for fixed records, record schemes, more parts etc.

The generic proof methods are sufficiently smart to pick the most sensible rule according to the type of the indicated record expression: users just need to apply something like “(*cases* r)” to a certain proof problem.

6. The derived record operations *t.make*, *t.fields*, *t.extend*, *t.truncate* are *not* treated automatically, but usually need to be expanded by hand, using the collective fact *t.defs*.

10.4 Datatypes

datatype : *theory* \rightarrow *theory*
rep_datatype : *theory* \rightarrow *proof*(*prove*)



datatype defines inductive datatypes in HOL.

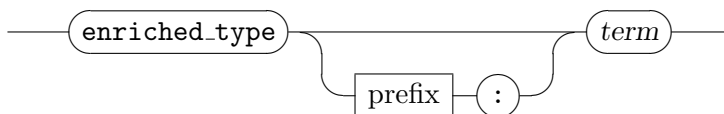
rep_datatype represents existing types as inductive ones, generating the standard infrastructure of derived concepts (primitive recursion etc.).

The induction and exhaustion theorems generated provide case names according to the constructors involved, while parameters are named after the types (see also §6.6).

See [18] for more details on datatypes, but beware of the old-style theory syntax being used there! Apart from proper proof methods for case-analysis and induction, there are also emulations of ML tactics *case_tac* and *induct_tac* available, see §10.12; these admit to refer directly to the internal structure of subgoals (including internally bound parameters).

10.5 Functorial structure of types

enriched_type : *local_theory* \rightarrow *proof*(*prove*)



enriched_type allows to prove and register properties about the functorial structure of type constructors; these properties then can be used by other packages to deal with those type constructors in certain type constructions. Characteristic theorems are noted in the current local theory; by default, they are prefixed with the base name of the type constructor, an explicit prefix can be given alternatively.

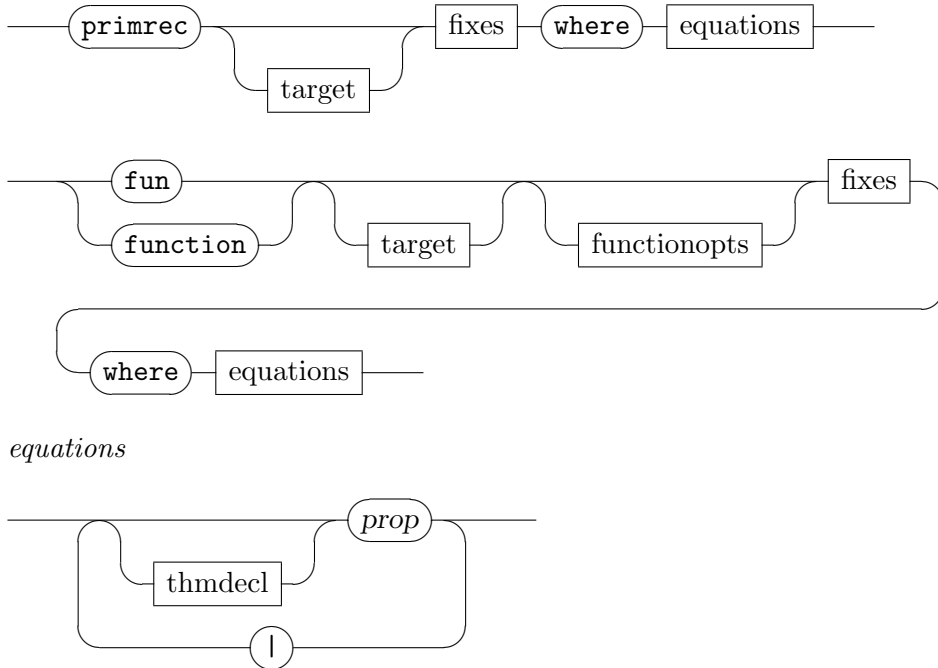
The given term m is considered as *mapper* for the corresponding type constructor and must conform to the following type pattern:

$$m :: \sigma_1 \Rightarrow \dots \sigma_k \Rightarrow (\bar{\alpha}_n) t \Rightarrow (\bar{\beta}_n) t$$

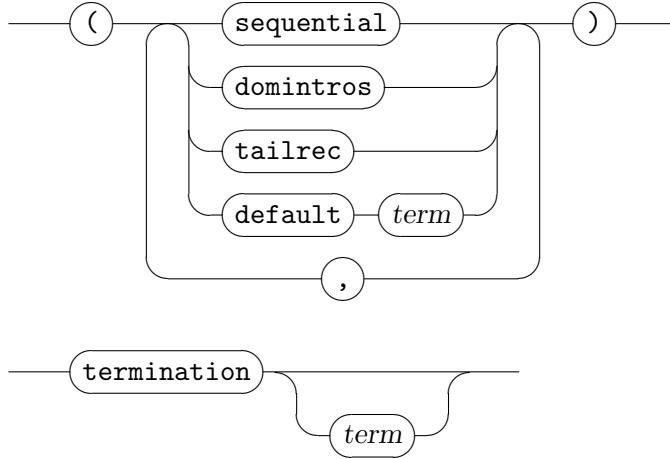
where t is the type constructor, $\bar{\alpha}_n$ and $\bar{\beta}_n$ are distinct type variables free in the local theory and $\sigma_1, \dots, \sigma_k$ is a subsequence of $\alpha_1 \Rightarrow \beta_1, \beta_1 \Rightarrow \alpha_1, \dots, \alpha_n \Rightarrow \beta_n, \beta_n \Rightarrow \alpha_n$.

10.6 Recursive functions

primrec : *local_theory* \rightarrow *local_theory*
fun : *local_theory* \rightarrow *local_theory*
function : *local_theory* \rightarrow *proof*(*prove*)
termination : *local_theory* \rightarrow *proof*(*prove*)



functionopts



primrec defines primitive recursive functions over datatypes, see also [18].

function defines functions by general wellfounded recursion. A detailed description with examples can be found in [10]. The function is specified by a set of (possibly conditional) recursive equations with arbitrary pattern matching. The command generates proof obligations for the completeness and the compatibility of patterns.

The defined function is considered partial, and the resulting simplification rules (named *f.psimps*) and induction rule (named *f.pinduct*) are guarded by a generated domain predicate *f_dom*. The **termination** command can then be used to establish that the function is total.

fun is a shorthand notation for “**function** (*sequential*)”, followed by automated proof attempts regarding pattern matching and termination. See [10] for further details.

termination *f* commences a termination proof for the previously defined function *f*. If this is omitted, the command refers to the most recent function definition. After the proof is closed, the recursive equations and the induction principle is established.

Recursive definitions introduced by the **function** command accommodate reasoning by induction (cf. §6.6): rule *c.induct* (where *c* is the name of the function definition) refers to a specific induction rule, with parameters named according to the user-specified equations. Cases are numbered (starting from 1).

For **primrec**, the induction principle coincides with structural recursion on the datatype the recursion is carried out.

The equations provided by these packages may be referred later as theorem list $f.simps$, where f is the (collective) name of the functions defined. Individual equations may be named explicitly as well.

The **function** command accepts the following options.

sequential enables a preprocessor which disambiguates overlapping patterns by making them mutually disjoint. Earlier equations take precedence over later ones. This allows to give the specification in a format very similar to functional programming. Note that the resulting simplification and induction rules correspond to the transformed specification, not the one given originally. This usually means that each equation given by the user may result in several theorems. Also note that this automatic transformation only works for ML-style datatype patterns.

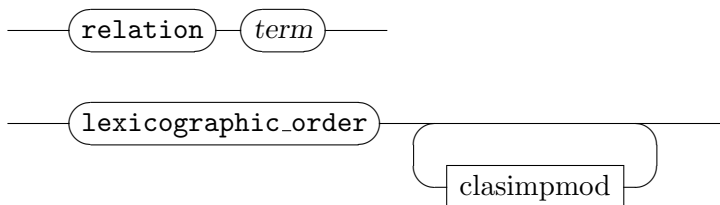
domintros enables the automated generation of introduction rules for the domain predicate. While mostly not needed, they can be helpful in some proofs about partial functions.

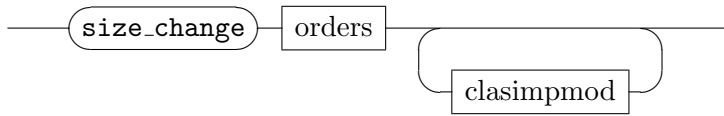
tailrec generates the unconstrained recursive equations even without a termination proof, provided that the function is tail-recursive. This currently only works

default d allows to specify a default value for a (partial) function, which will ensure that $f\ x = d\ x$ whenever $x \notin f.dom$.

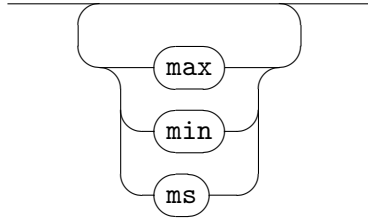
10.6.1 Proof methods related to recursive definitions

pat_completeness : method
 relation : method
lexicographic_order : method
 size_change : method





orders



pat_completeness is a specialized method to solve goals regarding the completeness of pattern matching, as required by the **function** package (cf. [10]).

relation R introduces a termination proof using the relation R . The resulting proof state will contain goals expressing that R is wellfounded, and that the arguments of recursive calls decrease with respect to R . Usually, this method is used as the initial proof step of manual termination proofs.

lexicographic_order attempts a fully automated termination proof by searching for a lexicographic combination of size measures on the arguments of the function. The method accepts the same arguments as the *auto* method, which it uses internally to prove local descents. The same context modifiers as for *auto* are accepted, see §9.4.3.

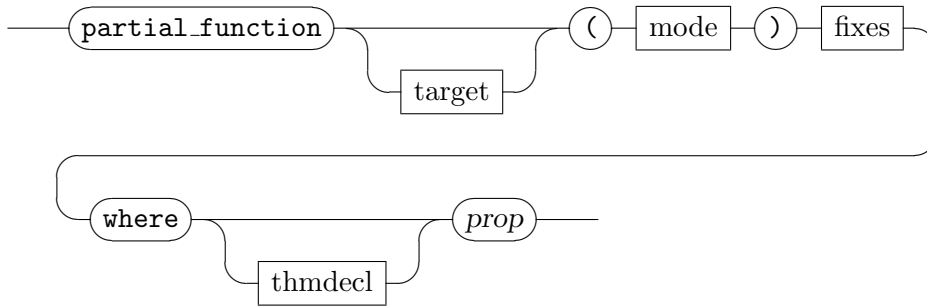
In case of failure, extensive information is printed, which can help to analyse the situation (cf. [10]).

size_change also works on termination goals, using a variation of the size-change principle, together with a graph decomposition technique (see [11] for details). Three kinds of orders are used internally: *max*, *min*, and *ms* (multiset), which is only available when the theory *Multiset* is loaded. When no order kinds are given, they are tried in order. The search for a termination proof uses SAT solving internally.

For local descent proofs, the same context modifiers as for *auto* are accepted, see §9.4.3.

10.6.2 Functions with explicit partiality

partial_function : $local_theory \rightarrow local_theory$
partial_function_mono : *attribute*



partial_function defines recursive functions based on fixpoints in complete partial orders. No termination proof is required from the user or constructed internally. Instead, the possibility of non-termination is modelled explicitly in the result type, which contains an explicit bottom element.

Pattern matching and mutual recursion are currently not supported. Thus, the specification consists of a single function described by a single recursive equation.

There are no fixed syntactic restrictions on the body of the function, but the induced functional must be provably monotonic wrt. the underlying order. The monotonicity proof is performed internally, and the definition is rejected when it fails. The proof can be influenced by declaring hints using the *partial_function_mono* attribute.

The mandatory *mode* argument specifies the mode of operation of the command, which directly corresponds to a complete partial order on the result type. By default, the following modes are defined:

option defines functions that map into the *option* type. Here, the value *None* is used to model a non-terminating computation. Monotonicity requires that if *None* is returned by a recursive call, then the overall result must also be *None*. This is best achieved through the use of the monadic operator *Option.bind*.

tailrec defines functions with an arbitrary result type and uses the slightly degenerated partial order where *undefined* is the bottom element. Now, monotonicity requires that if *undefined* is returned by a recursive call, then the overall result must also be *undefined*.

In practice, this is only satisfied when each recursive call is a tail call, whose result is directly returned. Thus, this mode of operation allows the definition of arbitrary tail-recursive functions.

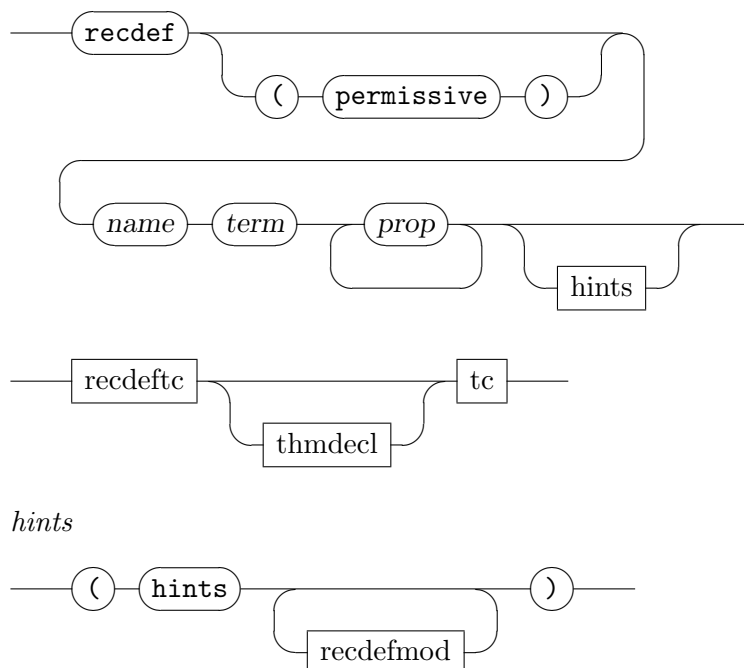
Experienced users may define new modes by instantiating the locale *partial_function_definitions* appropriately.

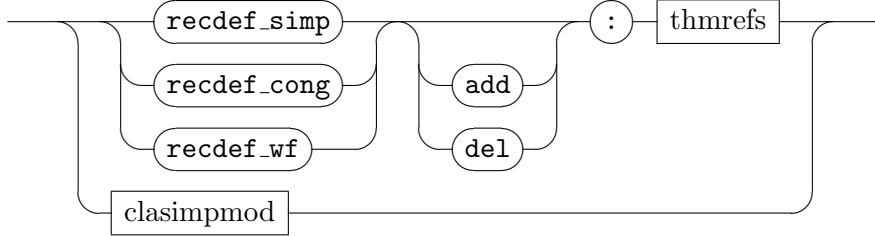
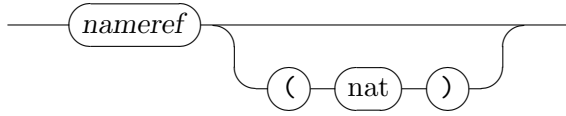
partial_function_mono declares rules for use in the internal mononicity proofs of partial function definitions.

10.6.3 Old-style recursive function definitions (TFL)

The old TFL commands **recdef** and **recdef_tc** for defining recursive are mostly obsolete; **function** or **fun** should be used instead.

recdef : *theory* \rightarrow *theory*)
recdef_tc* : *theory* \rightarrow *proof*(*prove*)



recdefmod*tc*

recdef defines general well-founded recursive functions (using the TFL package), see also [18]. The “(*permissive*)” option tells TFL to recover from failed proof attempts, returning unfinished results. The *recdef_simp*, *recdef_cong*, and *recdef_wf* hints refer to auxiliary rules to be used in the internal automated proof process of TFL. Additional *clasimpmod* declarations (cf. §9.4.3) may be given to tune the context of the Simplifier (cf. §9.3) and Classical reasoner (cf. §9.4).

recdef_tc *c* (*i*) recommences the proof for leftover termination condition number *i* (default 1) as generated by a **recdef** definition of constant *c*.

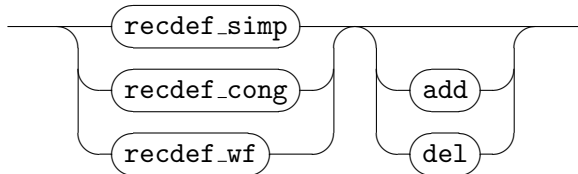
Note that in most cases, **recdef** is able to finish its internal proofs without manual intervention.

Hints for **recdef** may be also declared globally, using the following attributes.

```

recdef_simp  : attribute
recdef_cong : attribute
recdef_wf   : attribute

```



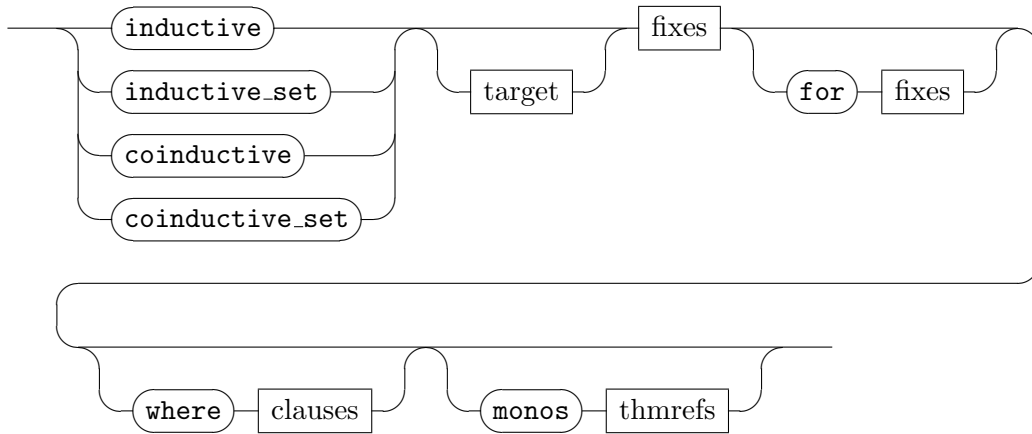
10.7 Inductive and coinductive definitions

An **inductive definition** specifies the least predicate (or set) R closed under given rules: applying a rule to elements of R yields a result within R . For example, a structural operational semantics is an inductive definition of an evaluation relation.

Dually, a **coinductive definition** specifies the greatest predicate / set R that is consistent with given rules: every element of R can be seen as arising by applying a rule to elements of R . An important example is using bisimulation relations to formalise equivalence of processes and infinite data structures.

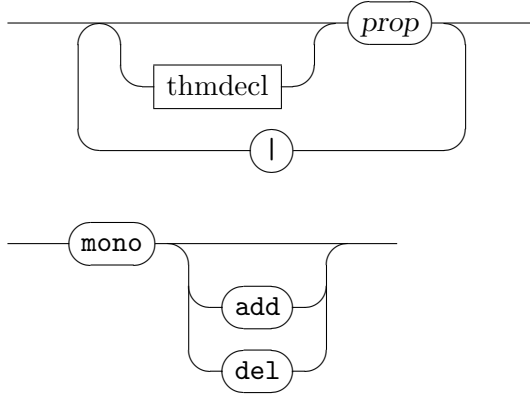
The HOL package is related to the ZF one, which is described in a separate paper,¹ which you should refer to in case of difficulties. The package is simpler than that of ZF thanks to implicit type-checking in HOL. The types of the (co)inductive predicates (or sets) determine the domain of the fixedpoint definition, and the package does not have to use inference rules for type-checking.

inductive : $local_theory \rightarrow local_theory$
inductive_set : $local_theory \rightarrow local_theory$
coinductive : $local_theory \rightarrow local_theory$
coinductive_set : $local_theory \rightarrow local_theory$
mono : *attribute*



¹It appeared in CADE [28]; a longer version is distributed with Isabelle.

clauses



inductive and **coinductive** define (co)inductive predicates from the introduction rules given in the **where** part. The optional **for** part contains a list of parameters of the (co)inductive predicates that remain fixed throughout the definition. The optional **monos** section contains *monotonicity theorems*, which are required for each operator applied to a recursive set in the introduction rules. There *must* be a theorem of the form $A \leq B \implies M A \leq M B$, for each premise $M R_i t$ in an introduction rule!

inductive_set and **coinductive_set** are wrappers for to the previous commands, allowing the definition of (co)inductive sets.

mono declares monotonicity rules. These rule are involved in the automated monotonicity proof of **inductive**.

10.7.1 Derived rules

Each (co)inductive definition R adds definitions to the theory and also proves some theorems:

$R.intros$ is the list of introduction rules as proven theorems, for the recursive predicates (or sets). The rules are also available individually, using the names given them in the theory file;

$R.cases$ is the case analysis (or elimination) rule;

$R.induct$ or $R.coinduct$ is the (co)induction rule.

When several predicates R_1, \dots, R_n are defined simultaneously, the list of introduction rules is called $R_1 \dots R_n.intros$, the case analysis rules are called $R_1.cases, \dots, R_n.cases$, and the list of mutual induction rules is called $R_1 \dots R_n.inducts$.

10.7.2 Monotonicity theorems

Each theory contains a default set of theorems that are used in monotonicity proofs. New rules can be added to this set via the *mono* attribute. The HOL theory *Inductive* shows how this is done. In general, the following monotonicity theorems may be added:

- Theorems of the form $A \leq B \implies M A \leq M B$, for proving monotonicity of inductive definitions whose introduction rules have premises involving terms such as $M R_i t$.
- Monotonicity theorems for logical operators, which are of the general form $(\dots \longrightarrow \dots) \implies \dots (\dots \longrightarrow \dots) \implies \dots \longrightarrow \dots$. For example, in the case of the operator \vee , the corresponding theorem is

$$\frac{P_1 \longrightarrow Q_1 \quad P_2 \longrightarrow Q_2}{P_1 \vee P_2 \longrightarrow Q_1 \vee Q_2}$$

- De Morgan style equations for reasoning about the “polarity” of expressions, e.g.

$$\neg \neg P \longleftrightarrow P \qquad \neg (P \wedge Q) \longleftrightarrow \neg P \vee \neg Q$$

- Equations for reducing complex operators to more primitive ones whose monotonicity can easily be proved, e.g.

$$(P \longrightarrow Q) \longleftrightarrow \neg P \vee Q \qquad Ball A P \equiv \forall x. x \in A \longrightarrow P x$$

10.8 Arithmetic proof support

arith : *method*
arith : *attribute*
arith_split : *attribute*

The *arith* method decides linear arithmetic problems (on types *nat*, *int*, *real*). Any current facts are inserted into the goal before running the procedure.

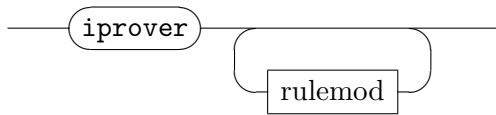
The *arith* attribute declares facts that are always supplied to the arithmetic provers implicitly.

The *arith_split* attribute declares case split rules to be expanded before *arith* is invoked.

Note that a simpler (but faster) arithmetic prover is already invoked by the Simplifier.

10.9 Intuitionistic proof search

iprover : *method*

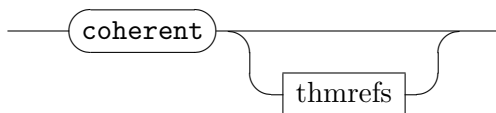


The *iprover* method performs intuitionistic proof search, depending on specifically declared rules from the context, or given as explicit arguments. Chained facts are inserted into the goal before commencing proof search.

Rules need to be classified as *intro*, *elim*, or *dest*; here the “!” indicator refers to “safe” rules, which may be applied aggressively (without considering back-tracking later). Rules declared with “?” are ignored in proof search (the single-step *rule* method still observes these). An explicit weight annotation may be given as well; otherwise the number of rule premises will be taken into account here.

10.10 Coherent Logic

coherent : *method*

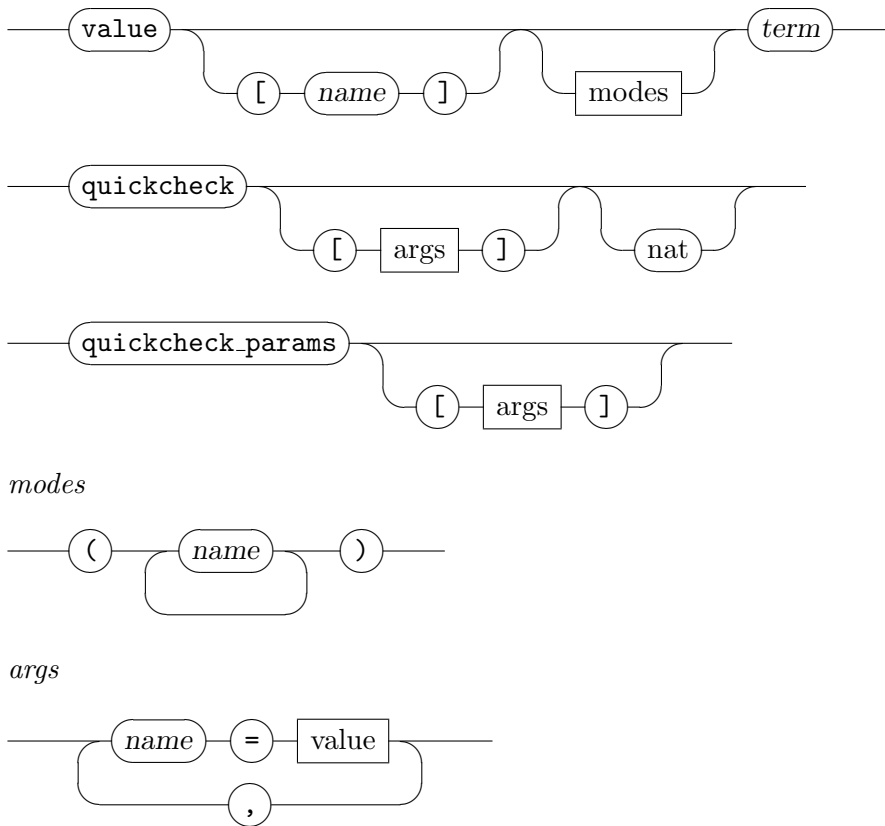


The *coherent* method solves problems of *Coherent Logic* [5], which covers applications in confluence theory, lattice theory and projective geometry. See `~~/src/HOL/ex/Coherent.thy` for some examples.

10.11 Checking and refuting propositions

Identifying incorrect propositions usually involves evaluation of particular assignments and systematic counter example search. This is supported by the following commands.

value* : *context* \rightarrow
quickcheck* : *proof* \rightarrow
quickcheck_params : *theory* \rightarrow *theory*



value *t* evaluates and prints a term; optionally *modes* can be specified, which are appended to the current print mode (see also [24]). Internally, the evaluation is performed by registered evaluators, which are invoked sequentially until a result is returned. Alternatively a specific evaluator can be selected using square brackets; typical evaluators use the current set of code equations to normalize and include *simp* for fully symbolic evaluation using the simplifier, *nbe* for *normalization by evaluation* and *code* for code generation in SML.

quickcheck tests the current goal for counter examples using a series of assignments for its free variables; by default the first subgoal is tested, an other can be selected explicitly using an optional goal index. Assignments can be chosen exhausting the search space upto a given size or using a fixed number of random assignments in the search space. By default, quickcheck uses exhaustive testing. A number of configuration options are supported for **quickcheck**, notably:

tester specifies how to explore the search space (e.g. exhaustive or random). An unknown configuration option is treated as an argument to tester, making *tester* = optional.

size specifies the maximum size of the search space for assignment values.

iterations sets how many sets of assignments are generated for each particular size.

no_assms specifies whether assumptions in structured proofs should be ignored.

timeout sets the time limit in seconds.

default_type sets the type(s) generally used to instantiate type variables.

report if set quickcheck reports how many tests fulfilled the preconditions.

quiet if not set quickcheck informs about the current size for assignment values.

expect can be used to check if the user's expectation was met (*no_expectation*, *no_counterexample*, or *counterexample*).

These option can be given within square brackets.

quickcheck_params changes quickcheck configuration options persistently.

10.12 Unstructured case analysis and induction

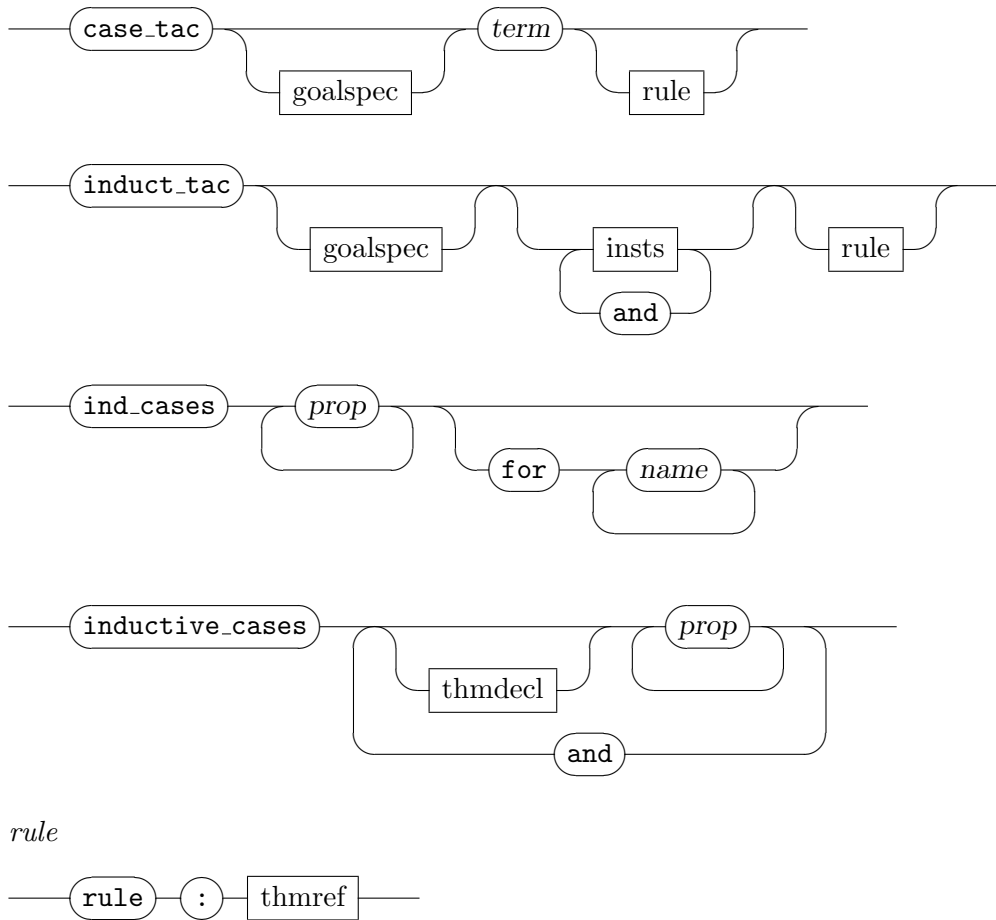
The following tools of Isabelle/HOL support cases analysis and induction in unstructured tactic scripts; see also §6.6 for proper Isar versions of similar

ideas.

```

    case_tac* : method
    induct_tac* : method
    ind_cases* : method
    inductive_cases* : local_theory → local_theory

```



`case_tac` and `induct_tac` admit to reason about inductive types. Rules are selected according to the declarations by the `cases` and `induct` attributes, cf. §6.6. The **datatype** package already takes care of this.

These unstructured tactics feature both goal addressing and dynamic instantiation. Note that named rule cases are *not* provided as would be by the proper `cases` and `induct` proof methods (see §6.6). Unlike the `induct` method, `induct_tac` does not handle structured rule statements, only the compact object-logic conclusion of the subgoal being addressed.

ind_cases and **inductive_cases** provide an interface to the internal **mk_cases** operation. Rules are simplified in an unrestricted forward manner.

While *ind_cases* is a proof method to apply the result immediately as elimination rules, **inductive_cases** provides case split theorems at the theory level for later use. The **for** argument of the *ind_cases* method allows to specify a list of variables that should be generalized before applying the resulting rule.

10.13 Executable code

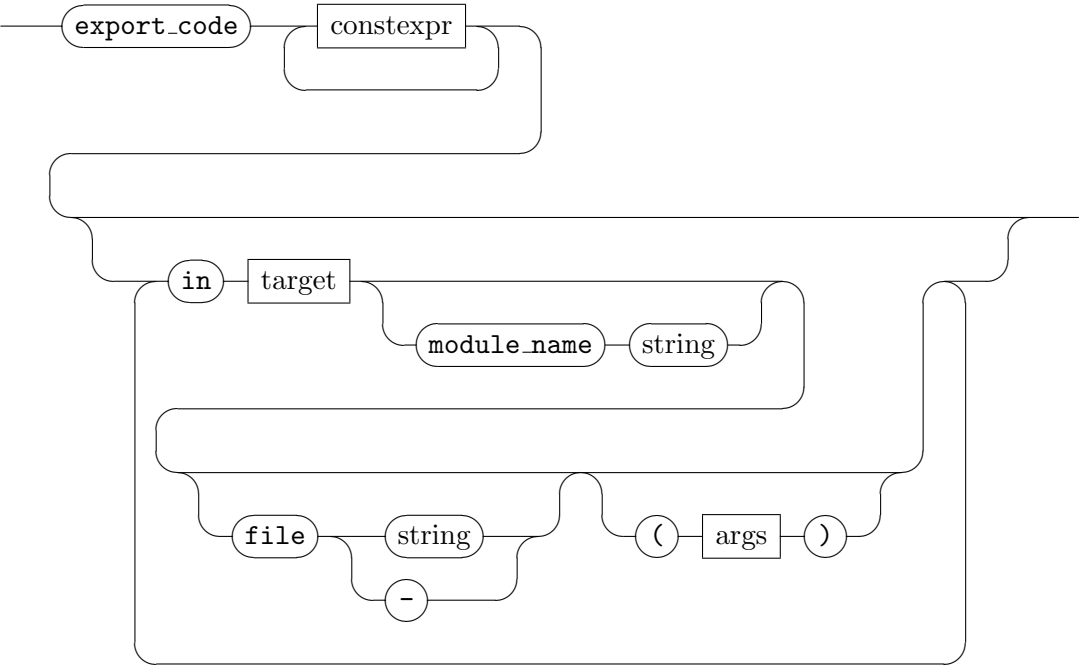
Isabelle/Pure provides two generic frameworks to support code generation from executable specifications. Isabelle/HOL instantiates these mechanisms in a way that is amenable to end-user applications.

One framework generates code from functional programs (including overloading using type classes) to SML [14], OCaml [12], Haskell [30] and Scala [22]. Conceptually, code generation is split up in three steps: *selection* of code theorems, *translation* into an abstract executable view and *serialization* to a specific *target language*. Inductive specifications can be executed using the predicate compiler which operates within HOL. See [8] for an introduction.

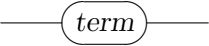
```

export_code*   : context →
                  code   : attribute
code_abort    : theory → theory
code_datatype : theory → theory
print_codesetup* : context →
                  code_inline : attribute
                  code_post   : attribute
print_codeproc* : context →
code_thms*      : context →
code_deps*      : context →
code_const     : theory → theory
code_type      : theory → theory
code_class     : theory → theory
code_instance  : theory → theory
code_reserved  : theory → theory
code_monad     : theory → theory
code_include   : theory → theory
code_modulename : theory → theory
code_reflect   : theory → theory

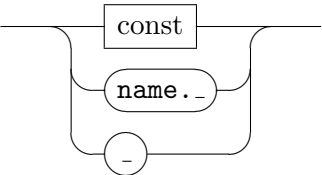
```

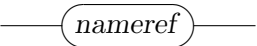
const



constexpr



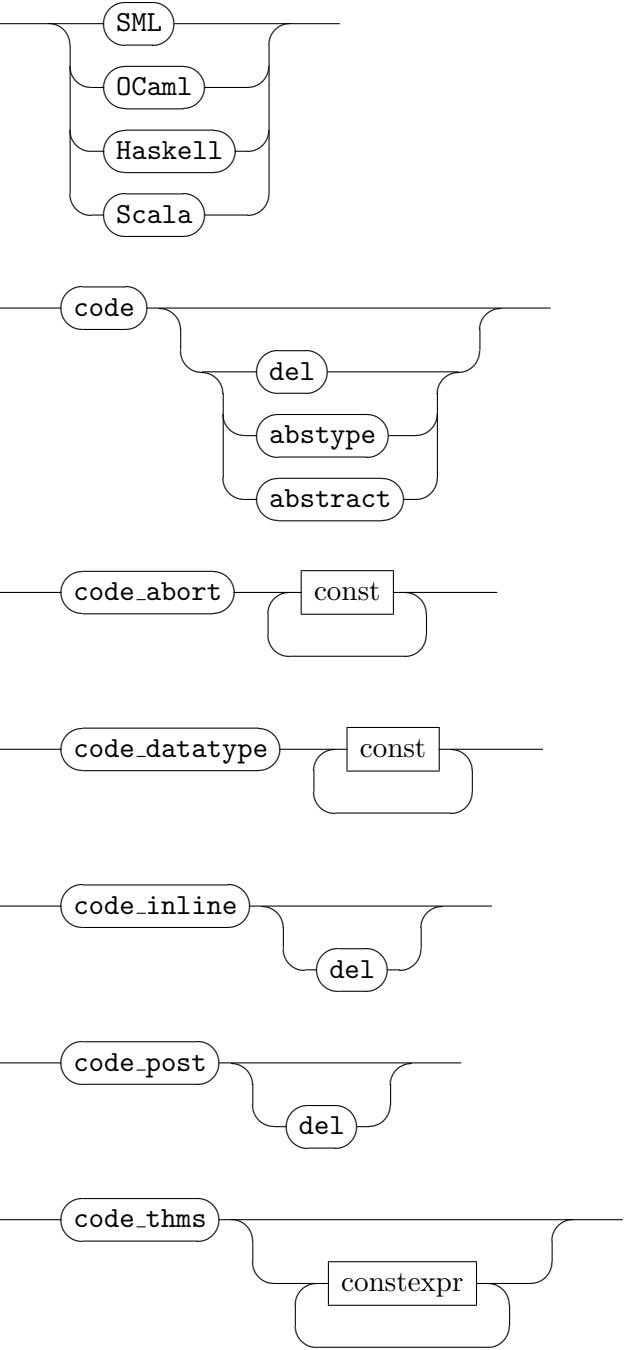
typeconstructor

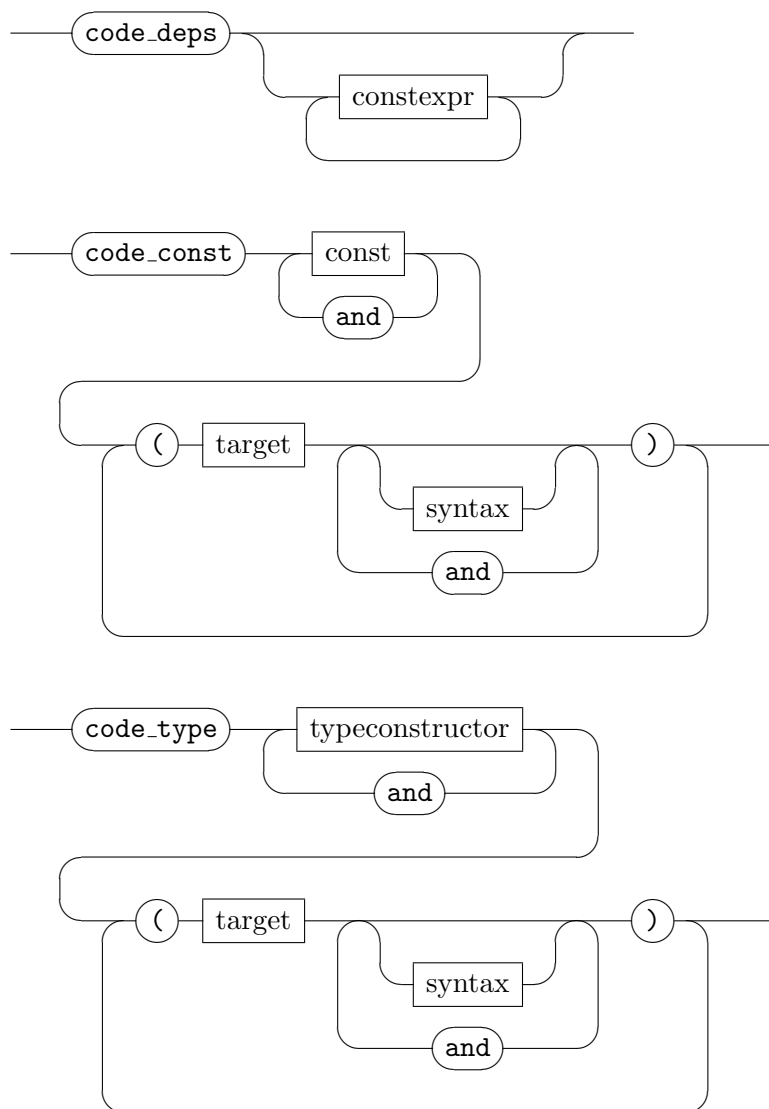


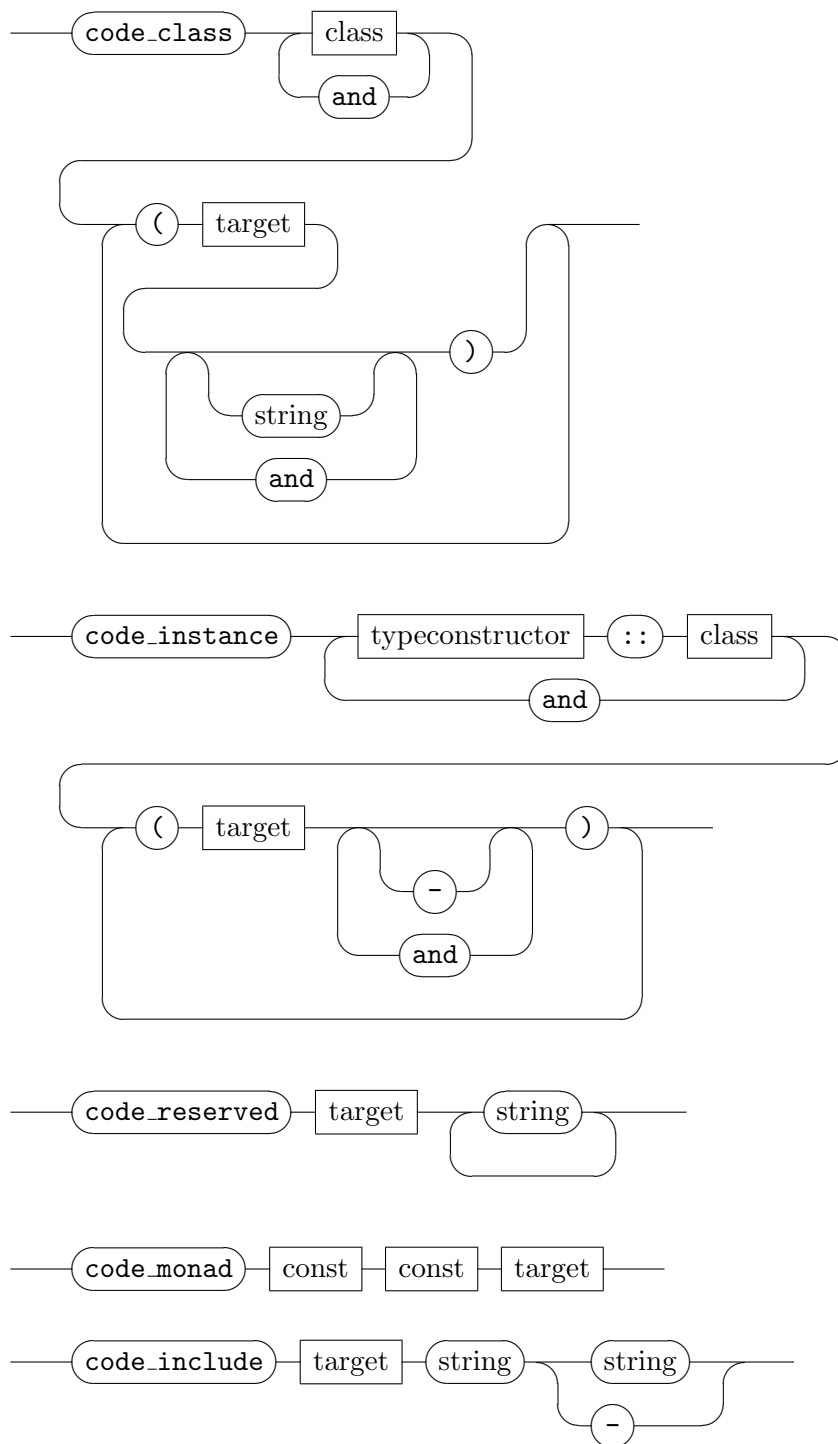
class

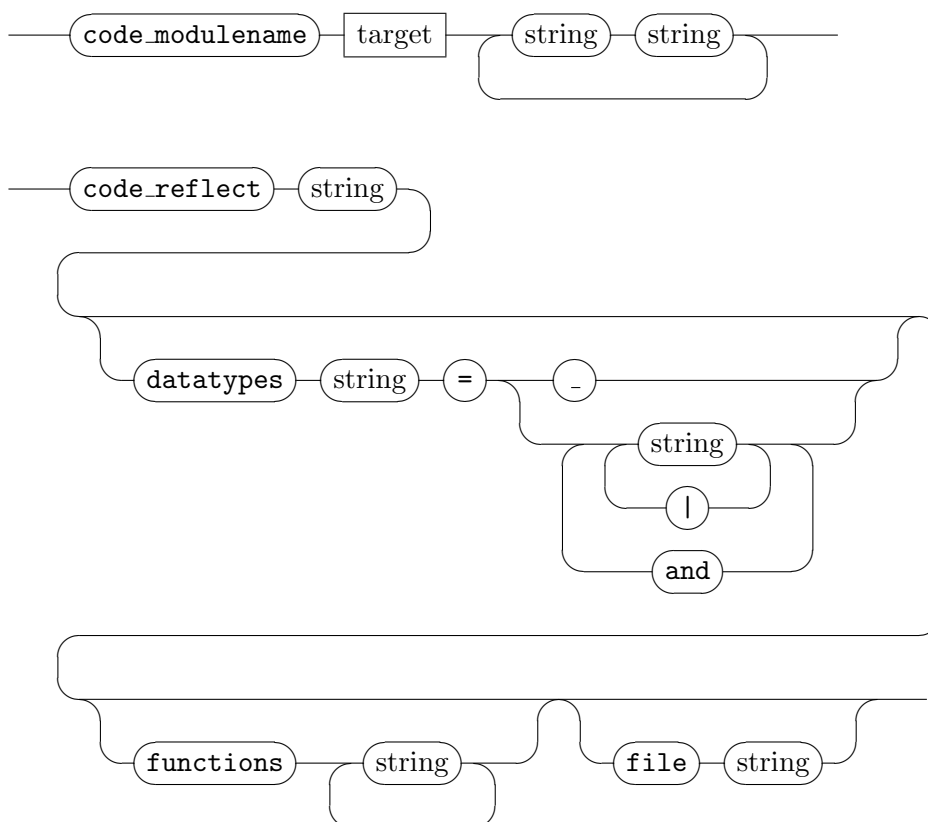


target

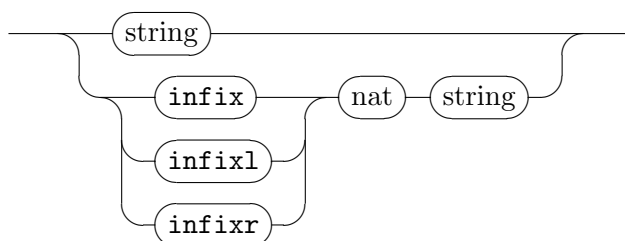








syntax



export_code generates code for a given list of constants in the specified target language(s). If no serialization instruction is given, only abstract code is generated internally.

Constants may be specified by giving them literally, referring to all executable constants within a certain theory by giving *name.**, or referring to *all* executable constants currently available by giving ***.

By default, for each involved theory one corresponding name space module is generated. Alternatively, a module name may be specified after the **module_name** keyword; then *all* code is placed in this module.

For *SML*, *OCaml* and *Scala* the file specification refers to a single file; for *Haskell*, it refers to a whole directory, where code is generated in multiple files reflecting the module hierarchy. Omitting the file specification denotes standard output.

Serializers take an optional list of arguments in parentheses. For *SML* and *OCaml*, “*no_signatures*” omits explicit module signatures.

For *Haskell* a module name prefix may be given using the “*root:*” argument; “*string_classes*” adds a “**deriving** (**Read**, **Show**)” clause to each appropriate datatype declaration.

code explicitly selects (or with option “*del*” deselects) a code equation for code generation. Usually packages introducing code equations provide a reasonable default setup for selection. Variants *code abstype* and *code abstract* declare abstract datatype certificates or code equations on abstract datatype representations respectively.

code_abort declares constants which are not required to have a definition by means of code equations; if needed these are implemented by program abort instead.

code_datatype specifies a constructor set for a logical type.

print_codesetup gives an overview on selected code equations and code generator datatypes.

code_inline declares (or with option “*del*” removes) inlining theorems which are applied as rewrite rules to any code equation during preprocessing.

code_post declares (or with option “*del*” removes) theorems which are applied as rewrite rules to any result of an evaluation.

print_codeproc prints the setup of the code generator preprocessor.

code_thms prints a list of theorems representing the corresponding program containing all given constants after preprocessing.

code_deps visualizes dependencies of theorems representing the corresponding program containing all given constants after preprocessing.

code_const associates a list of constants with target-specific serializations; omitting a serialization deletes an existing serialization.

code_type associates a list of type constructors with target-specific serializations; omitting a serialization deletes an existing serialization.

code_class associates a list of classes with target-specific class names; omitting a serialization deletes an existing serialization. This applies only to *Haskell*.

code_instance declares a list of type constructor / class instance relations as “already present” for a given target. Omitting a “–” deletes an existing “already present” declaration. This applies only to *Haskell*.

code_reserved declares a list of names as reserved for a given target, preventing it to be shadowed by any generated code.

code_monad provides an auxiliary mechanism to generate monadic code for Haskell.

code_include adds arbitrary named content (“include”) to generated code. A “–” as last argument will remove an already added “include”.

code_modulename declares aliasings from one module name onto another.

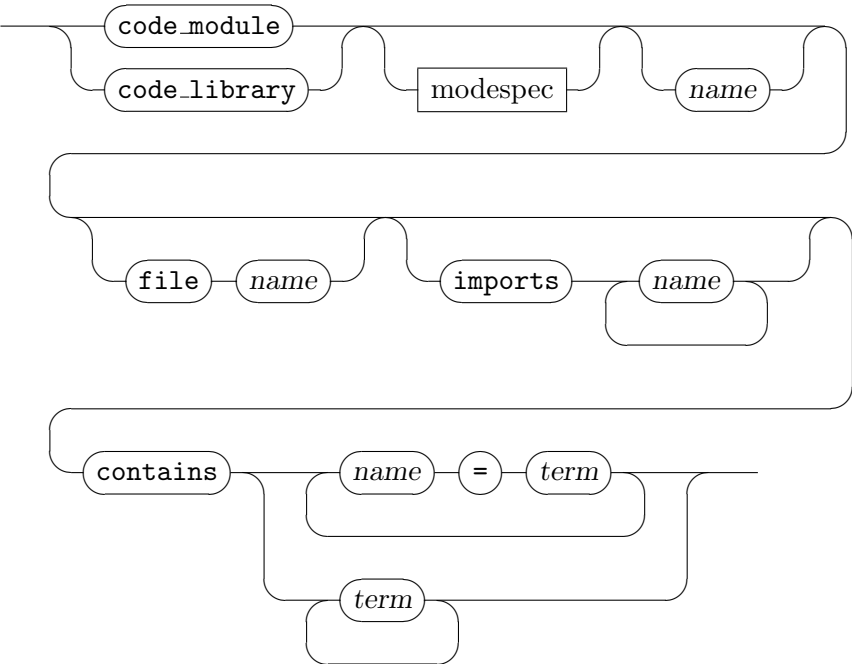
code_reflect without a “file” argument compiles code into the system runtime environment and modifies the code generator setup that future invocations of system runtime code generation referring to one of the “*datatypes*” or “*functions*” entities use these precompiled entities. With a “file” argument, the corresponding code is generated into that specified file without modifying the code generator setup.

The other framework generates code from both functional and relational programs to SML. See [18] for further information (this actually covers the new-style theory format as well).

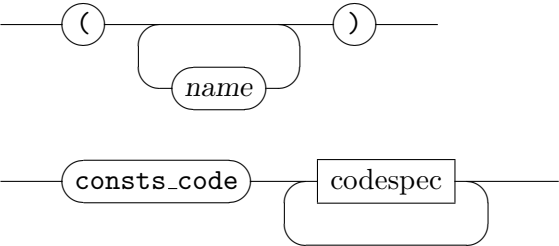
```

code_module  : theory → theory
code_library : theory → theory
consts_code  : theory → theory
types_code   : theory → theory
               code : attribute

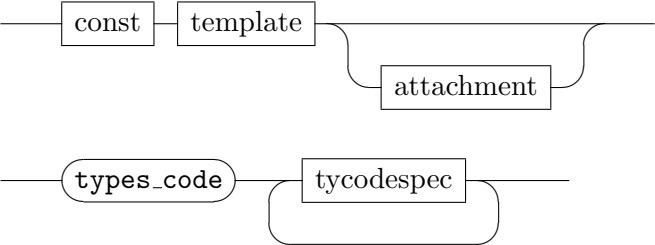
```

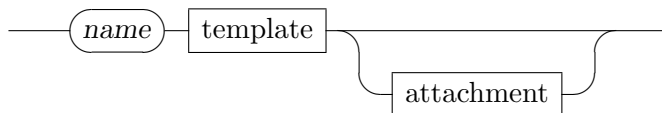
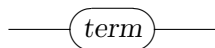
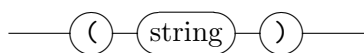
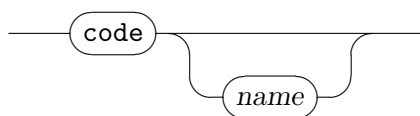
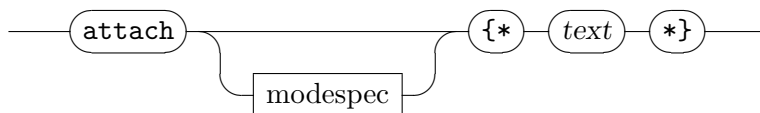


modespec



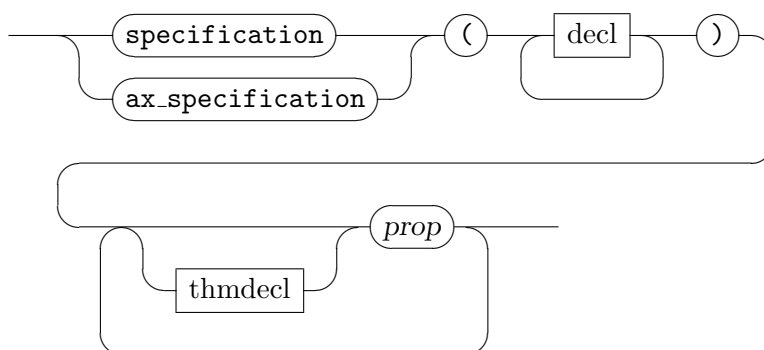
codespec



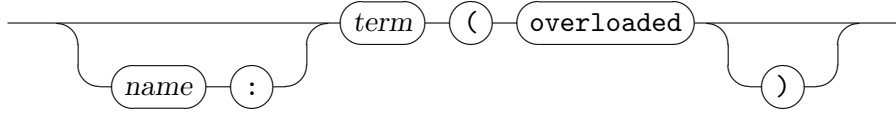
tycodespec*const**template**attachment*

10.14 Definition by specification

specification : *theory* \rightarrow *proof*(*prove*)
ax_specification : *theory* \rightarrow *proof*(*prove*)



decl



specification *decls* φ sets up a goal stating the existence of terms with the properties specified to hold for the constants given in *decls*. After finishing the proof, the theory will be augmented with definitions for the given constants, as well as with theorems stating the properties for these constants.

ax_specification *decls* φ sets up a goal stating the existence of terms with the properties specified to hold for the constants given in *decls*. After finishing the proof, the theory will be augmented with axioms expressing the properties given in the first place.

decl declares a constant to be defined by the specification given. The definition for the constant *c* is bound to the name *c_def* unless a theorem name is given in the declaration. Overloaded constants should be declared as such.

Whether to use **specification** or **ax_specification** is to some extent a matter of style. **specification** introduces no new axioms, and so by construction cannot introduce inconsistencies, whereas **ax_specification** does introduce axioms, but only after the user has explicitly proven it to be safe. A practical issue must be considered, though: After introducing two constants with the same properties using **specification**, one can prove that the two constants are, in fact, equal. If this might be a problem, one should use **ax_specification**.

Isabelle/HOLCF

11.1 Mixfix syntax for continuous operations

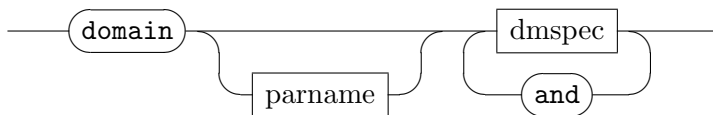
consts : *theory* \rightarrow *theory*

HOLCF provides a separate type for continuous functions $\alpha \rightarrow \beta$, with an explicit application operator $f \cdot x$. Isabelle mixfix syntax normally refers directly to the pure meta-level function type $\alpha \Rightarrow \beta$, with application $f x$.

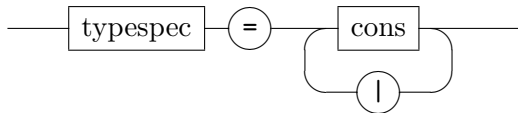
The HOLCF variant of **consts** modifies that of Pure Isabelle (cf. §5.9.3) such that declarations involving continuous function types are treated specifically. Any given syntax template is transformed internally, generating translation rules for the abstract and concrete representation of continuous application. Note that mixing of HOLCF and Pure application is *not* supported!

11.2 Recursive domains

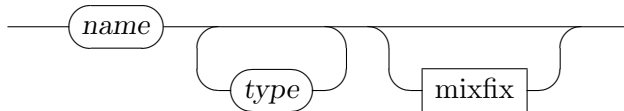
domain : *theory* \rightarrow *theory*



dmspec



cons



dtrules



Recursive domains in HOLCF are analogous to datatypes in classical HOL (cf. §10.4). Mutual recursion is supported, but no nesting nor arbitrary branching. Domain constructors may be strict (default) or lazy, the latter admits to introduce infinitary objects in the typical LCF manner (e.g. lazy lists). See also [15] for a general discussion of HOLCF domains.

Isabelle/ZF

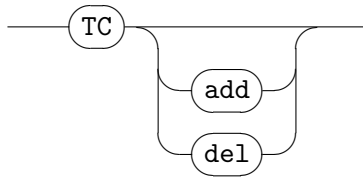
12.1 Type checking

The ZF logic is essentially untyped, so the concept of “type checking” is performed as logical reasoning about set-membership statements. A special method assists users in this task; a version of this is already declared as a “solver” in the standard Simplifier setup.

```

print_tcset*  : context →
    typecheck  : method
    TC        : attribute

```



print_tcset prints the collection of typechecking rules of the current context.

typecheck attempts to solve any pending type-checking problems in subgoals.

TC adds or deletes type-checking rules from the context.

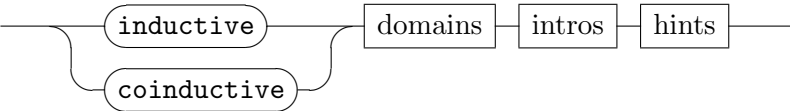
12.2 (Co)Inductive sets and datatypes

12.2.1 Set definitions

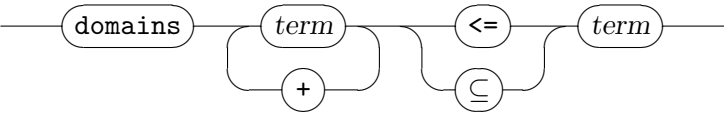
In ZF everything is a set. The generic inductive package also provides a specific view for “datatype” specifications. Coinductive definitions are available

in both cases, too.

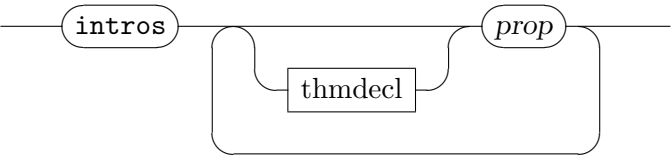
inductive : *theory* \rightarrow *theory*
coinductive : *theory* \rightarrow *theory*
datatype : *theory* \rightarrow *theory*
codatatype : *theory* \rightarrow *theory*



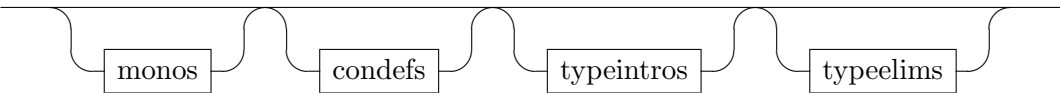
domains



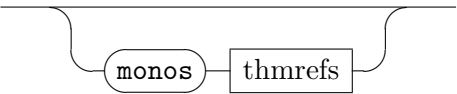
intros



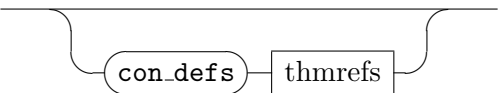
hints



monos



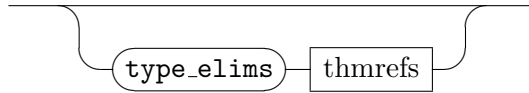
condefs



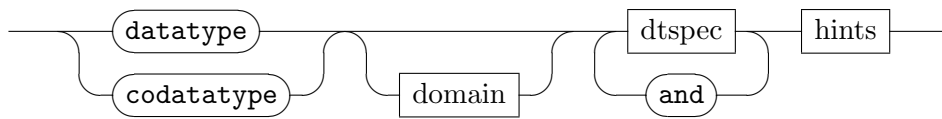
typeintros



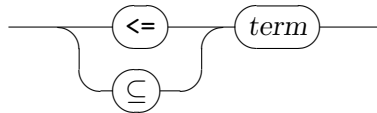
typeelims



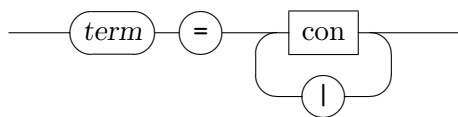
In the following syntax specification *monos*, *typeintros*, and *typeelims* are the same as above.



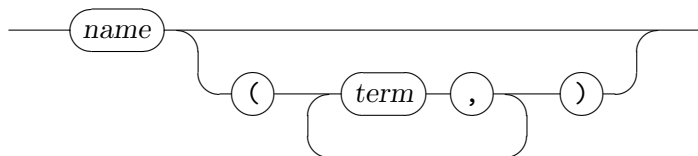
domain



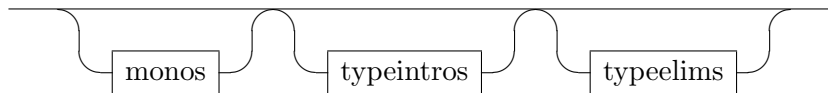
dtspec



con



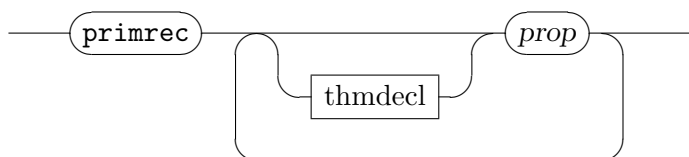
hints



See [23] for further information on inductive definitions in ZF, but note that this covers the old-style theory format.

12.2.2 Primitive recursive functions

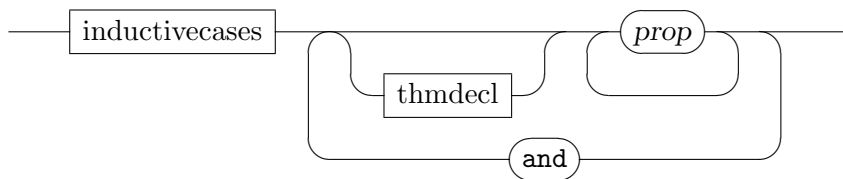
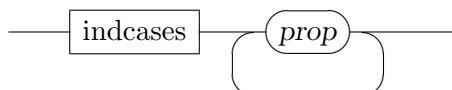
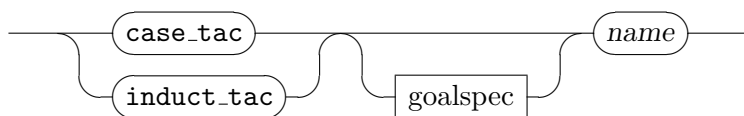
primrec : *theory* \rightarrow *theory*



12.2.3 Cases and induction: emulating tactic scripts

The following important tactical tools of Isabelle/ZF have been ported to Isar. These should not be used in proper proof texts.

*case_tac** : *method*
*induct_tac** : *method*
*ind_cases** : *method*
inductive_cases : *theory* \rightarrow *theory*



Part IV

Appendix

Isabelle/Isar quick reference

A.1 Proof commands

A.1.1 Primitives and basic syntax

fix x	augment context by $\bigwedge x. \square$
assume $a: \varphi$	augment context by $\varphi \implies \square$
then	indicate forward chaining of facts
have $a: \varphi$	prove local result
show $a: \varphi$	prove local result, refining some goal
using a	indicate use of additional facts
unfolding a	unfold definitional equations
proof $m_1 \dots$ qed m_2	indicate proof structure and refinements
{ ... }	indicate explicit blocks
next	switch blocks
note $a = b$	reconsider facts
let $p = t$	abbreviate terms by higher-order matching
write $c \ (mx)$	declare local mixfix syntax

$proof = prfx^* \mathbf{proof} \ method^? \ stmt^* \mathbf{qed} \ method^?$
 $\quad \quad \quad | \quad prfx^* \mathbf{done}$

$prfx = \mathbf{apply} \ method$
 $\quad \quad \quad | \quad \mathbf{using} \ facts$
 $\quad \quad \quad | \quad \mathbf{unfolding} \ facts$

$stmt = \{ \ stmt^* \}$
 $\quad \quad \quad | \quad \mathbf{next}$
 $\quad \quad \quad | \quad \mathbf{note} \ name = facts$
 $\quad \quad \quad | \quad \mathbf{let} \ term = term$
 $\quad \quad \quad | \quad \mathbf{write} \ name \ (mixfix)$
 $\quad \quad \quad | \quad \mathbf{fix} \ var^+$
 $\quad \quad \quad | \quad \mathbf{assume} \ name: props$
 $\quad \quad \quad | \quad \mathbf{then}^? \ goal$

$goal = \mathbf{have} \ name: props \ proof$
 $\quad \quad \quad | \quad \mathbf{show} \ name: props \ proof$

A.1.2 Abbreviations and synonyms

by m_1 m_2	\equiv	proof m_1 qed m_2
..	\equiv	by rule
.	\equiv	by this
hence	\equiv	then have
thus	\equiv	then show
from a	\equiv	note a then
with a	\equiv	from a and <i>this</i>
from <i>this</i>	\equiv	then
from <i>this</i> have	\equiv	hence
from <i>this</i> show	\equiv	thus

A.1.3 Derived elements

also ₀	\approx	note <i>calculation</i> = <i>this</i>
also _{$n+1$}	\approx	note <i>calculation</i> = <i>trans</i> [<i>OF calculation this</i>]
finally	\approx	also from <i>calculation</i>
moreover	\approx	note <i>calculation</i> = <i>calculation this</i>
ultimately	\approx	moreover from <i>calculation</i>
presume $a: \varphi$	\approx	assume $a: \varphi$
def $a: x \equiv t$	\approx	fix x assume $a: x \equiv t$
obtain x where $a: \varphi$	\approx	\dots fix x assume $a: \varphi$
case c	\approx	fix x assume $c: \varphi$
sorry	\approx	by cheating

A.1.4 Diagnostic commands

pr	print current state
thm a	print fact
prop φ	print proposition
term t	print term
typ τ	print type

A.2 Proof methods

Single steps (forward-chaining facts)

<i>assumption</i>	apply some assumption
<i>this</i>	apply current facts
<i>rule a</i>	apply some rule
<i>rule</i>	apply standard rule (default for proof)
<i>contradiction</i>	apply \neg elimination rule (any order)
<i>cases t</i>	case analysis (provides cases)
<i>induct x</i>	proof by induction (provides cases)

Repeated steps (inserting facts)

—	no rules
<i>intro a</i>	introduction rules
<i>intro_classes</i>	class introduction rules
<i>elim a</i>	elimination rules
<i>unfold a</i>	definitional rewrite rules

Automated proof tools (inserting facts)

<i>iprover</i>	intuitionistic proof search
<i>blast, fast</i>	Classical Reasoner
<i>simp, simp_all</i>	Simplifier (+ Splitter)
<i>auto, force</i>	Simplifier + Classical Reasoner
<i>arith</i>	Arithmetic procedures

A.3 Attributes

Operations

<i>OF a</i>	rule resolved with facts (skipping “_”)
<i>of t</i>	rule instantiated with terms (skipping “_”)
<i>where x = t</i>	rule instantiated with terms, by variable name
<i>symmetric</i>	resolution with symmetry rule
<i>THEN b</i>	resolution with another rule
<i>rule_format</i>	result put into standard rule format
<i>elim_format</i>	destruct rule turned into elimination rule format

Declarations

<i>simp</i>	Simplifier rule
<i>intro, elim, dest</i>	Pure or Classical Reasoner rule
<i>iff</i>	Simplifier + Classical Reasoner rule
<i>split</i>	case split rule
<i>trans</i>	transitivity rule
<i>sym</i>	symmetry rule

A.4 Rule declarations and methods

	<i>rule</i>	<i>iprover</i>	<i>blast</i> <i>fast</i>	<i>simp</i> <i>simp_all</i>	<i>auto</i> <i>force</i>
<i>Pure.elim! Pure.intro!</i>	×	×			
<i>Pure.elim Pure.intro</i>	×	×			
<i>elim! intro!</i>	×		×		×
<i>elim intro</i>	×		×		×
<i>iff</i>	×		×	×	×
<i>iff?</i>	×				
<i>elim? intro?</i>	×				
<i>simp</i>				×	×
<i>cong</i>				×	×
<i>split</i>				×	×

A.5 Emulating tactic scripts

A.5.1 Commands

apply <i>m</i>	apply proof method at initial position
apply_end <i>m</i>	apply proof method near terminal position
done	complete proof
defer <i>n</i>	move subgoal to end
prefer <i>n</i>	move subgoal to beginning
back	backtrack last command

A.5.2 Methods

<i>rule_tac insts</i>	resolution (with instantiation)
<i>erule_tac insts</i>	elim-resolution (with instantiation)
<i>drule_tac insts</i>	destruct-resolution (with instantiation)
<i>frule_tac insts</i>	forward-resolution (with instantiation)
<i>cut_tac insts</i>	insert facts (with instantiation)
<i>thin_tac</i> φ	delete assumptions
<i>subgoal_tac</i> φ	new claims
<i>rename_tac</i> <i>x</i>	rename innermost goal parameters
<i>rotate_tac</i> <i>n</i>	rotate assumptions of goal
<i>tactic text</i>	arbitrary ML tactic
<i>case_tac</i> <i>t</i>	exhaustion (datatypes)
<i>induct_tac</i> <i>x</i>	induction (datatypes)
<i>ind_cases</i> <i>t</i>	exhaustion + simplification (inductive predicates)

Predefined Isabelle symbols

Isabelle supports an infinite number of non-ASCII symbols, which are represented in source text as `\<name>` (where *name* may be any identifier). It is left to front-end tools how to present these symbols to the user. The collection of predefined standard symbols given below is available by default for Isabelle document output, due to appropriate definitions of `\isasymname` for each `\<name>` in the `isabellesym.sty` file. Most of these symbols are displayed properly in Proof General if used with the X-Symbol package.

Moreover, any single symbol (or ASCII character) may be prefixed by `\<^sup>`, for superscript and `\<^sub>`, for subscript, such as $A\<^sup>\<star>$, for A^* the alternative versions `\<^isub>` and `\<^isup>` are considered as quasi letters and may be used within identifiers. Sub- and superscripts that span a region of text are marked up with `\<^bsub>...\<^esub>`, and `\<^bsup>...\<^esup>` respectively. Furthermore, all ASCII characters and most other symbols may be printed in bold by prefixing `\<^bold>` such as `\<^bold>\<alpha>` for α . Note that `\<^bold>`, may *not* be combined with sub- or superscripts for single symbols.

Further details of Isabelle document preparation are covered in chapter 4.

<code>\<zero></code>	0	<code>\<one></code>	1
<code>\<two></code>	2	<code>\<three></code>	3
<code>\<four></code>	4	<code>\<five></code>	5
<code>\<six></code>	6	<code>\<seven></code>	7
<code>\<eight></code>	8	<code>\<nine></code>	9
<code>\<A></code>	<i>A</i>	<code>\</code>	<i>B</i>
<code>\<C></code>	<i>C</i>	<code>\<D></code>	<i>D</i>
<code>\<E></code>	<i>E</i>	<code>\<F></code>	<i>F</i>
<code>\<G></code>	<i>G</i>	<code>\<H></code>	<i>H</i>
<code>\<I></code>	<i>I</i>	<code>\<J></code>	<i>J</i>
<code>\<K></code>	<i>K</i>	<code>\<L></code>	<i>L</i>
<code>\<M></code>	<i>M</i>	<code>\<N></code>	<i>N</i>
<code>\<O></code>	<i>O</i>	<code>\<P></code>	<i>P</i>
<code>\<Q></code>	<i>Q</i>	<code>\<R></code>	<i>R</i>

<code>\<S></code>	\mathcal{S}	<code>\<T></code>	\mathcal{T}
<code>\<U></code>	\mathcal{U}	<code>\<V></code>	\mathcal{V}
<code>\<W></code>	\mathcal{W}	<code>\<X></code>	\mathcal{X}
<code>\<Y></code>	\mathcal{Y}	<code>\<Z></code>	\mathcal{Z}
<code>\<a></code>	a	<code>\</code>	b
<code>\<c></code>	c	<code>\<d></code>	d
<code>\<e></code>	e	<code>\<f></code>	f
<code>\<g></code>	g	<code>\<h></code>	h
<code>\<i></code>	i	<code>\<j></code>	j
<code>\<k></code>	k	<code>\<l></code>	l
<code>\<m></code>	m	<code>\<n></code>	n
<code>\<o></code>	o	<code>\<p></code>	p
<code>\<q></code>	q	<code>\<r></code>	r
<code>\<s></code>	s	<code>\<t></code>	t
<code>\<u></code>	u	<code>\<v></code>	v
<code>\<w></code>	w	<code>\<x></code>	x
<code>\<y></code>	y	<code>\<z></code>	z
<code>\<AA></code>	\mathfrak{A}	<code>\<BB></code>	\mathfrak{B}
<code>\<CC></code>	\mathfrak{C}	<code>\<DD></code>	\mathfrak{D}
<code>\<EE></code>	\mathfrak{E}	<code>\<FF></code>	\mathfrak{F}
<code>\<GG></code>	\mathfrak{G}	<code>\<HH></code>	\mathfrak{H}
<code>\<II></code>	\mathfrak{I}	<code>\<JJ></code>	\mathfrak{J}
<code>\<KK></code>	\mathfrak{K}	<code>\<LL></code>	\mathfrak{L}
<code>\<MM></code>	\mathfrak{M}	<code>\<NN></code>	\mathfrak{N}
<code>\<OO></code>	\mathfrak{O}	<code>\<PP></code>	\mathfrak{P}
<code>\<QQ></code>	\mathfrak{Q}	<code>\<RR></code>	\mathfrak{R}
<code>\<SS></code>	\mathfrak{S}	<code>\<TT></code>	\mathfrak{T}
<code>\<UU></code>	\mathfrak{U}	<code>\<VV></code>	\mathfrak{V}
<code>\<WW></code>	\mathfrak{W}	<code>\<XX></code>	\mathfrak{X}
<code>\<YY></code>	\mathfrak{Y}	<code>\<ZZ></code>	\mathfrak{Z}
<code>\<aa></code>	a	<code>\<bb></code>	b
<code>\<cc></code>	c	<code>\<dd></code>	d
<code>\<ee></code>	e	<code>\<ff></code>	f
<code>\<gg></code>	g	<code>\<hh></code>	h
<code>\<ii></code>	i	<code>\<jj></code>	j
<code>\<kk></code>	k	<code>\<ll></code>	l
<code>\<mm></code>	m	<code>\<nn></code>	n
<code>\<oo></code>	o	<code>\<pp></code>	p
<code>\<qq></code>	q	<code>\<rr></code>	r

<code>\<ss></code>	\mathfrak{s}	<code>\<tt></code>	\mathfrak{t}
<code>\<uu></code>	\mathfrak{u}	<code>\<vv></code>	\mathfrak{v}
<code>\<ww></code>	\mathfrak{w}	<code>\<xx></code>	\mathfrak{x}
<code>\<yy></code>	\mathfrak{y}	<code>\<zz></code>	\mathfrak{z}
<code>\<alpha></code>	α	<code>\<beta></code>	β
<code>\<gamma></code>	γ	<code>\<delta></code>	δ
<code>\<epsilon></code>	ε	<code>\<zeta></code>	ζ
<code>\<eta></code>	η	<code>\<theta></code>	ϑ
<code>\<iota></code>	ι	<code>\<kappa></code>	κ
<code>\<lambda></code>	λ	<code>\<mu></code>	μ
<code>\<nu></code>	ν	<code>\<xi></code>	ξ
<code>\<pi></code>	π	<code>\<rho></code>	ϱ
<code>\<sigma></code>	σ	<code>\<tau></code>	τ
<code>\<upsilon></code>	υ	<code>\<phi></code>	φ
<code>\<chi></code>	χ	<code>\<psi></code>	ψ
<code>\<omega></code>	ω	<code>\<Gamma></code>	Γ
<code>\<Delta></code>	Δ	<code>\<Theta></code>	Θ
<code>\<Lambda></code>	Λ	<code>\<Xi></code>	Ξ
<code>\<Pi></code>	Π	<code>\<Sigma></code>	Σ
<code>\<Upsilon></code>	Υ	<code>\<Phi></code>	Φ
<code>\<Psi></code>	Ψ	<code>\<Omega></code>	Ω
<code>\<bool></code>	\mathbb{B}	<code>\<complex></code>	\mathbb{C}
<code>\<nat></code>	\mathbb{N}	<code>\<rat></code>	\mathbb{Q}
<code>\<real></code>	\mathbb{R}	<code>\<int></code>	\mathbb{Z}
<code>\<leftarrow></code>	\leftarrow	<code>\<longleftarrow></code>	\longleftarrow
<code>\<rightarrow></code>	\rightarrow	<code>\<longrightarrow></code>	\longrightarrow
<code>\<Leftarrow></code>	\Lleftarrow	<code>\<Longleftarrow></code>	\Longleftarrow
<code>\<Rightarrow></code>	\Rightarrow	<code>\<Longrightarrow></code>	\Longrightarrow
<code>\<leftrightarrow></code>	\leftrightarrow	<code>\<longleftrightarrow></code>	\longleftrightarrow
<code>\<Leftrightarrow></code>	\Leftrightarrow	<code>\<Longleftrightarrow></code>	\Longleftrightarrow
<code>\<mapsto></code>	\mapsto	<code>\<longmapsto></code>	\longmapsto
<code>\<midarrow></code>	\mid	<code>\<Midarrow></code>	$=$
<code>\<hookleftarrow></code>	\hookleftarrow	<code>\<hookrightarrow></code>	\hookrightarrow
<code>\<leftharpoondown></code>	\leftharpoonup	<code>\<rightharpoondown></code>	\rightharpoonup
<code>\<leftharpoonup></code>	\leftarrow	<code>\<rightharpoonup></code>	\rightarrow
<code>\<rightleftharpoons></code>	\rightleftharpoons	<code>\<leadsto></code>	\leadsto
<code>\<downharpoonleft></code>	\downarrow	<code>\<downharpoonright></code>	\downarrow
<code>\<upharpoonleft></code>	\uparrow	<code>\<upharpoonright></code>	\uparrow
<code>\<restriction></code>	\restriction	<code>\<Colon></code>	$::$

<code>\<up></code>	\uparrow	<code>\<Up></code>	\Uparrow
<code>\<down></code>	\downarrow	<code>\<Down></code>	\Downarrow
<code>\<updown></code>	\updownarrow	<code>\<Updown></code>	\Updownarrow
<code>\<langle></code>	\langle	<code>\<rangle></code>	\rangle
<code>\<lceil></code>	\lceil	<code>\<rceil></code>	\rceil
<code>\<lfloor></code>	\lfloor	<code>\<rfloor></code>	\rfloor
<code>\<lparr></code>	$\langle\!\!\langle$	<code>\<rparr></code>	$\rangle\!\!\rangle$
<code>\<lbrakk></code>	\llbracket	<code>\<rbrakk></code>	\rrbracket
<code>\<lbrace></code>	$\{$	<code>\<rbrace></code>	$\}$
<code>\<guillemotleft></code>	\ll	<code>\<guillemotright></code>	\gg
<code>\<bottom></code>	\perp	<code>\<top></code>	\top
<code>\<and></code>	\wedge	<code>\<And></code>	\bigwedge
<code>\<or></code>	\vee	<code>\<Or></code>	\bigvee
<code>\<forall></code>	\forall	<code>\<exists></code>	\exists
<code>\<nexists></code>	\nexists	<code>\<not></code>	\neg
<code>\<box></code>	\square	<code>\<diamond></code>	\diamond
<code>\<turnstile></code>	\vdash	<code>\<Turnstile></code>	\models
<code>\<tturnstile></code>	\Vdash	<code>\<TTurnstile></code>	\Vdash
<code>\<stileturn></code>	\dashv	<code>\<surd></code>	\surd
<code>\<le></code>	\leq	<code>\<ge></code>	\geq
<code>\<lless></code>	\ll	<code>\<ggreater></code>	\gg
<code>\<lesssim></code>	\lesssim	<code>\<greatersim></code>	\gtrsim
<code>\<lessapprox></code>	\approx	<code>\<greaterapprox></code>	\gtrapprox
<code>\<in></code>	\in	<code>\<notin></code>	\notin
<code>\<subset></code>	\subset	<code>\<supset></code>	\supset
<code>\<subseteq></code>	\subseteq	<code>\<supseteq></code>	\supseteq
<code>\<sqsubset></code>	\sqsubset	<code>\<sqsupset></code>	\sqsupset
<code>\<sqsubsetq></code>	\sqsubseteq	<code>\<sqsupseteq></code>	\sqsupseteq
<code>\<inter></code>	\cap	<code>\<Inter></code>	\bigcap
<code>\<union></code>	\cup	<code>\<Union></code>	\bigcup
<code>\<squnion></code>	\sqcup	<code>\<Squnion></code>	\bigsqcup
<code>\<sqinter></code>	\sqcap	<code>\<Sqinter></code>	\bigsqcap
<code>\<setminus></code>	\setminus	<code>\<propto></code>	\propto
<code>\<uplus></code>	\uplus	<code>\<Uplus></code>	$\mathop{\uplus}$
<code>\<noteq></code>	\neq	<code>\<sim></code>	\sim
<code>\<doteq></code>	\doteq	<code>\<simeq></code>	\simeq
<code>\<approx></code>	\approx	<code>\<asympt></code>	\asymp
<code>\<cong></code>	\cong	<code>\<smile></code>	\smile
<code>\<equiv></code>	\equiv	<code>\<frown></code>	\frown

<code>\<Join></code>	\bowtie	<code>\<bowtie></code>	\bowtie
<code>\<prec></code>	\prec	<code>\<succ></code>	\succ
<code>\<preceq></code>	\preceq	<code>\<succeq></code>	\succeq
<code>\<parallel></code>	\parallel	<code>\<bar></code>	$ $
<code>\<plusminus></code>	\pm	<code>\<minusplus></code>	\mp
<code>\<times></code>	\times	<code>\<div></code>	\div
<code>\<cdot></code>	\cdot	<code>\<star></code>	\star
<code>\<bullet></code>	\bullet	<code>\<circ></code>	\circ
<code>\<dagger></code>	\dagger	<code>\<ddagger></code>	\ddagger
<code>\<lhd></code>	\triangleleft	<code>\<rh></code>	\triangleright
<code>\<unlhd></code>	\trianglelefteq	<code>\<unrh></code>	\trianglerighteq
<code>\<triangleleft></code>	\triangleleft	<code>\<triangleright></code>	\triangleright
<code>\<triangle></code>	\triangle	<code>\<triangleq></code>	\triangleq
<code>\<oplus></code>	\oplus	<code>\<Oplus></code>	\bigoplus
<code>\<otimes></code>	\otimes	<code>\<Otimes></code>	\bigotimes
<code>\<odot></code>	\odot	<code>\<Odot></code>	\bigodot
<code>\<ominus></code>	\ominus	<code>\<oslash></code>	\oslash
<code>\<dots></code>	\dots	<code>\<cdots></code>	\cdots
<code>\<Sum></code>	\sum	<code>\<Prod></code>	\prod
<code>\<Coproduct></code>	\coprod	<code>\<infinity></code>	∞
<code>\<integral></code>	\int	<code>\<ointegral></code>	\oint
<code>\<clubsuit></code>	\clubsuit	<code>\<diamondsuit></code>	\diamondsuit
<code>\<heartsuit></code>	\heartsuit	<code>\<spadesuit></code>	\spadesuit
<code>\<aleph></code>	\aleph	<code>\<emptyset></code>	\emptyset
<code>\<nabla></code>	∇	<code>\<partial></code>	∂
<code>\<Re></code>	\Re	<code>\<Im></code>	\Im
<code>\<flat></code>	\flat	<code>\<natural></code>	\natural
<code>\<sharp></code>	\sharp	<code>\<angle></code>	\angle
<code>\<copyright></code>	\copyright	<code>\<registered></code>	$\text{\textcircled{R}}$
<code>\<hyphen></code>	$-$	<code>\<inverse></code>	$^{-1}$
<code>\<onesuperior></code>	1	<code>\<onequarter></code>	$\frac{1}{4}$
<code>\<twosuperior></code>	2	<code>\<onehalf></code>	$\frac{1}{2}$
<code>\<threesuperior></code>	3	<code>\<threequarters></code>	$\frac{3}{4}$
<code>\<ordfeminine></code>	a	<code>\<ordmasculine></code>	o
<code>\<section></code>	\S	<code>\<paragraph></code>	\P
<code>\<exclamdown></code>	i	<code>\<questiondown></code>	?
<code>\<euro></code>	€	<code>\<pounds></code>	£
<code>\<yen></code>	¥	<code>\<cent></code>	¢
<code>\<currency></code>	¤	<code>\<degree></code>	$^\circ$

<code>\<amalg></code>	\amalg	<code>\<mho></code>	\mho
<code>\<lozenge></code>	\lozenge	<code>\<wp></code>	\wp
<code>\<wrong></code>	\wr	<code>\<struct></code>	\struct
<code>\<acute></code>	$\acute{}$	<code>\<index></code>	\index
<code>\<dieresis></code>	$\ddot{}$	<code>\<cedilla></code>	¸
<code>\<hungarumlaut></code>	¨	<code>\<spacespace></code>	 ^{space}
<code>\<module></code>	$\langle \mathbf{module} \rangle$	<code>\<some></code>	ϵ

ML tactic expressions

Isar Proof methods closely resemble traditional tactics, when used in unstructured sequences of **apply** commands. Isabelle/Isar provides emulations for all major ML tactics of classic Isabelle — mostly for the sake of easy porting of existing developments, as actual Isar proof texts would demand much less diversity of proof methods.

Unlike tactic expressions in ML, Isar proof methods provide proper concrete syntax for additional arguments, options, modifiers etc. Thus a typical method text is usually more concise than the corresponding ML tactic. Furthermore, the Isar versions of classic Isabelle tactics often cover several variant forms by a single method with separate options to tune the behavior. For example, method *simp* replaces all of *simp_tac* / *asm_simp_tac* / *full_simp_tac* / *asm_full_simp_tac*, there is also concrete syntax for augmenting the Simplifier context (the current “simpset”) in a convenient way.

C.1 Resolution tactics

Classic Isabelle provides several variant forms of tactics for single-step rule applications (based on higher-order resolution). The space of resolution tactics has the following main dimensions.

1. The “mode” of resolution: *intro*, *elim*, *destruct*, or *forward* (e.g. *resolve_tac*, *eresolve_tac*, *dresolve_tac*, *forward_tac*).
2. Optional explicit instantiation (e.g. *resolve_tac* vs. *res_inst_tac*).
3. Abbreviations for singleton arguments (e.g. *resolve_tac* vs. *rtac*).

Basically, the set of Isar tactic emulations *rule_tac*, *erule_tac*, *drule_tac*, *frule_tac* (see §9.2.3) would be sufficient to cover the four modes, either with or without instantiation, and either with single or multiple arguments. Although it is more convenient in most cases to use the plain *rule* method (see §6.3.3), or any of its “improper” variants *erule*, *drule*, *frule* (see §9.2.1). Note that explicit goal addressing is only supported by the actual *rule_tac* version.

With this in mind, plain resolution tactics correspond to Isar methods as follows.

<code>rtac a 1</code>	<code>rule a</code>
<code>resolve_tac [a₁, ...] 1</code>	<code>rule a₁ ...</code>
<code>res_inst_tac ctxt [(x₁, t₁), ...] a 1</code>	<code>rule_tac x₁ = t₁ and ... in a</code>
<code>rtac a i</code>	<code>rule_tac [i] a</code>
<code>resolve_tac [a₁, ...] i</code>	<code>rule_tac [i] a₁ ...</code>
<code>res_inst_tac ctxt [(x₁, t₁), ...] a i</code>	<code>rule_tac [i] x₁ = t₁ and ... in a</code>

Note that explicit goal addressing may be usually avoided by changing the order of subgoals with **defer** or **prefer** (see §6.3.4).

C.2 Simplifier tactics

The main Simplifier tactics `simp_tac` and variants (cf. [24]) are all covered by the *simp* and *simp_all* methods (see §9.3). Note that there is no individual goal addressing available, simplification acts either on the first goal (*simp*) or all goals (*simp_all*).

<code>asm_full_simp_tac @{simpset} 1</code>	<code>simp</code>
<code>ALLGOALS (asm_full_simp_tac @{simpset})</code>	<code>simp_all</code>
<code>simp_tac @{simpset} 1</code>	<code>simp (no_asm)</code>
<code>asm_simp_tac @{simpset} 1</code>	<code>simp (no_asm_simp)</code>
<code>full_simp_tac @{simpset} 1</code>	<code>simp (no_asm_use)</code>
<code>asm_lr_simp_tac @{simpset} 1</code>	<code>simp (asm_lr)</code>

C.3 Classical Reasoner tactics

The Classical Reasoner provides a rather large number of variations of automated tactics, such as `blast_tac`, `fast_tac`, `clarify_tac` etc. (see [24]). The corresponding Isar methods usually share the same base name, such as *blast*, *fast*, *clarify* etc. (see §9.4).

C.4 Miscellaneous tactics

There are a few additional tactics defined in various theories of Isabelle/HOL, some of these also in Isabelle/FOL or Isabelle/ZF. The most common ones of these may be ported to Isar as follows.

<code>stac a 1</code>		<code>subst a</code>
<code>hyp_subst_tac 1</code>		<code>hypsubst</code>
<code>strip_tac 1</code>	\approx	<code>intro strip</code>
<code>split_all_tac 1</code>		<code>simp (no_asm_simp) only: split_tupled_all</code>
	\approx	<code>simp only: split_tupled_all</code>
	\ll	<code>clarify</code>

C.5 Tacticals

Classic Isabelle provides a huge amount of tacticals for combination and modification of existing tactics. This has been greatly reduced in Isar, providing the bare minimum of combinators only: “;” (sequential composition), “[|” (alternative choices), “?” (try), “+” (repeat at least once). These are usually sufficient in practice; if all fails, arbitrary ML tactic code may be invoked via the *tactic* method (see §9.2.3).

Common ML tacticals may be expressed directly in Isar as follows:

<code>tac₁ THEN tac₂</code>	<code>meth₁, meth₂</code>
<code>tac₁ ORELSE tac₂</code>	<code>meth₁ meth₂</code>
<code>TRY tac</code>	<code>meth?</code>
<code>REPEAT1 tac</code>	<code>meth+</code>
<code>REPEAT tac</code>	<code>(meth+)?</code>
<code>EVERY [tac₁, ...]</code>	<code>meth₁, ...</code>
<code>FIRST [tac₁, ...]</code>	<code>meth₁ ...</code>

CHANGED (see [24]) is usually not required in Isar, since most basic proof methods already fail unless there is an actual change in the goal state. Nevertheless, “?” (try) may be used to accept *unchanged* results as well.

ALLGOALS, **SOMEGOAL** etc. (see [24]) are not available in Isar, since there is no direct goal addressing. Nevertheless, some basic methods address all goals internally, notably *simp_all* (see §9.3). Also note that **ALLGOALS** can be often replaced by “+” (repeat at least once), although this usually has a different operational behavior, such as solving goals in a different order.

Iterated resolution, such as **REPEAT (FIRSTGOAL**

(resolve_tac \<dots>)), is usually better expressed using the *intro* and *elim* methods of Isar (see §9.4).

Bibliography

- [1] D. Aspinall. Proof General. <http://proofgeneral.inf.ed.ac.uk/>.
- [2] D. Aspinall. Proof General: A generic tool for proof development. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 1785 of *Lecture Notes in Computer Science*, pages 38–42. Springer-Verlag, 2000.
- [3] G. Bauer and M. Wenzel. Calculational reasoning revisited — an Isabelle/Isar experience. In R. J. Boulton and P. B. Jackson, editors, *Theorem Proving in Higher Order Logics: TPHOLs 2001*, volume 2152 of *Lecture Notes in Computer Science*. Springer-Verlag, 2001.
- [4] S. Berghofer and T. Nipkow. Proof terms for simply typed higher order logic. In J. Harrison and M. Aagaard, editors, *Theorem Proving in Higher Order Logics: TPHOLs 2000*, volume 1869 of *Lecture Notes in Computer Science*, pages 38–52. Springer-Verlag, 2000.
- [5] M. Bezem and T. Coquand. Automating Coherent Logic. In G. Sutcliffe and A. Voronkov, editors, *LPAR-12*, volume 3835 of *Lecture Notes in Computer Science*. Springer-Verlag.
- [6] A. Church. A formulation of the simple theory of types. *Journal of Symbolic Logic*, 5:56–68, 1940.
- [7] G. Gentzen. Untersuchungen über das logische Schließen. *Math. Zeitschrift*, 1935.
- [8] F. Haftmann. *Code generation from Isabelle theories*. <http://isabelle.in.tum.de/doc/codegen.pdf>.
- [9] F. Haftmann. *Haskell-style type classes with Isabelle/Isar*. <http://isabelle.in.tum.de/doc/classes.pdf>.
- [10] A. Krauss. *Defining Recursive Functions in Isabelle/HOL*. <http://isabelle.in.tum.de/doc/functions.pdf>.
- [11] A. Krauss. *Automating Recursive Definitions and Termination Proofs in Higher-Order Logic*. PhD thesis, Institut für Informatik, Technische Universität München, Germany, 2009.

- [12] X. Leroy et al. *The Objective Caml system – Documentation and user’s manual*. <http://caml.inria.fr/pub/docs/manual-ocaml/>.
- [13] D. Miller. A logic programming language with lambda-abstraction, function variables, and simple unification. *Journal of Logic and Computation*, 1(4), 1991.
- [14] R. Milner, M. Tofte, and R. Harper. *The Definition of Standard ML*. MIT Press, 1990.
- [15] O. Müller, T. Nipkow, D. v. Oheimb, and O. Slotosch. HOLCF = HOL + LCF. *Journal of Functional Programming*, 9:191–223, 1999.
- [16] W. Naraschewski and M. Wenzel. Object-oriented verification based on record subtyping in higher-order logic. In J. Grundy and M. Newey, editors, *Theorem Proving in Higher Order Logics: TPHOLs ’98*, volume 1479 of *Lecture Notes in Computer Science*. Springer-Verlag, 1998.
- [17] T. Nipkow. Structured Proofs in Isar/HOL. In H. Geuvers and F. Wiedijk, editors, *Types for Proofs and Programs (TYPES 2002)*, volume 2646 of *Lecture Notes in Computer Science*, pages 259–278. Springer-Verlag, 2003.
- [18] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle’s Logics: HOL*. <http://isabelle.in.tum.de/doc/logics-HOL.pdf>.
- [19] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*. Springer, 2002. LNCS 2283.
- [20] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL: A Proof Assistant for Higher-Order Logic*, volume 2283 of *Lecture Notes in Computer Science*. Springer-Verlag, 2002.
- [21] T. Nipkow and C. Prehofer. Type reconstruction for type classes. *Journal of Functional Programming*, 5(2):201–224, 1995.
- [22] M. Odersky and al. An overview of the scala programming language. Technical Report IC/2004/64, EPFL Lausanne, Switzerland, 2004.
- [23] L. C. Paulson. *Isabelle’s Logics: FOL and ZF*. <http://isabelle.in.tum.de/doc/logics-ZF.pdf>.
- [24] L. C. Paulson. *The Old Isabelle Reference Manual*. <http://isabelle.in.tum.de/doc/ref.pdf>.
- [25] L. C. Paulson. Natural deduction as higher-order resolution. *Journal of Logic Programming*, 3:237–258, 1986.
- [26] L. C. Paulson. The foundation of a generic theorem prover. *Journal of Automated Reasoning*, 5(3):363–397, 1989.

- [27] L. C. Paulson. Isabelle: The next 700 theorem provers. In P. Odifreddi, editor, *Logic and Computer Science*, pages 361–386. Academic Press, 1990.
- [28] L. C. Paulson. A fixedpoint approach to implementing (co)inductive definitions. In A. Bundy, editor, *Automated Deduction — CADE-12 International Conference*, LNAI 814, pages 148–161. Springer, 1994.
- [29] L. C. Paulson. *ML for the Working Programmer*. Cambridge University Press, 2nd edition, 1996.
- [30] S. Peyton Jones et al. The Haskell 98 language and libraries: The revised report. *Journal of Functional Programming*, 13(1):0–255, Jan 2003.
<http://www.haskell.org/definition/>.
- [31] P. Rudnicki. An overview of the MIZAR project. In *1992 Workshop on Types for Proofs and Programs*. Chalmers University of Technology, Bastad, 1992.
- [32] P. Schroeder-Heister. A natural extension of natural deduction. *Journal of Symbolic Logic*, 49(4), 1984.
- [33] A. Trybulec. Some features of the Mizar language. Presented at a workshop in Turin, Italy, 1993.
- [34] M. Wenzel. *The Isabelle/Isar Implementation*.
<http://isabelle.in.tum.de/doc/implementation.pdf>.
- [35] M. Wenzel. Type classes and overloading in higher-order logic. In E. L. Gunter and A. Felty, editors, *Theorem Proving in Higher Order Logics: TPHOLs '97*, volume 1275 of *Lecture Notes in Computer Science*. Springer-Verlag, 1997.
- [36] M. Wenzel. Isar — a generic interpretative approach to readable formal proof documents. In Y. Bertot, G. Dowek, A. Hirschowitz, C. Paulin, and L. Thery, editors, *Theorem Proving in Higher Order Logics: TPHOLs '99*, volume 1690 of *Lecture Notes in Computer Science*. Springer-Verlag, 1999.
- [37] M. Wenzel. *Isabelle/Isar — a versatile environment for human-readable formal proof documents*. PhD thesis, Institut für Informatik, Technische Universität München, 2002.
<http://tumb1.biblio.tu-muenchen.de/publ/diss/in/2002/wenzel.html>.
- [38] M. Wenzel. Isabelle/Isar — a generic framework for human-readable proof documents. In R. Matuszewski and A. Zalewska, editors, *From Insight to Proof — Festschrift in Honour of Andrzej Trybulec*, volume 10(23) of *Studies in Logic, Grammar, and Rhetoric*. University of Białystok, 2007.
<http://www.in.tum.de/~wenzelm/papers/isar-framework.pdf>.

- [39] M. Wenzel and S. Berghofer. *The Isabelle System Manual*.
<http://isabelle.in.tum.de/doc/system.pdf>.
- [40] M. Wenzel and L. C. Paulson. Isabelle/Isar. In F. Wiedijk, editor, *The Seventeen Provers of the World*, LNAI 3600. 2006.
- [41] F. Wiedijk. Mizar: An impression. Unpublished paper, 1999.
<http://www.cs.kun.nl/~freek/mizar/mizarintro.ps.gz>.
- [42] F. Wiedijk and M. Wenzel. A comparison of the mathematical proof languages Mizar and Isar. *Journal of Automated Reasoning*, 29(3-4), 2002.

Index

- (method), **89**
- . (command), **86**
- .. (command), **86**
- ?thesis (variable), 85
- _ (fact), 81
- { (command), **76**
- } (command), **76**

- abbrev (antiquotation), **41**
- abbreviation (command), **51**
- also (command), **95**
- altstring (syntax), **29**, 29, 36
- and (keyword), 34, 78
- any (inner syntax), **115**, 117
- apply (command), 81, 82, **92**
- apply_end (command), **92**
- aprop (inner syntax), **116**, 117
- args (syntax), **35**
- arith (HOL attribute), **164**
- arith (HOL method), **164**
- arith_split (HOL attribute), **164**
- arities (command), **69**
- arity (syntax), **31**
- assms (fact), 82
- assume (command), **77**
- assumes (element), **56**
- assumption (inference), **9**
- assumption (method), **89**
- atom (syntax), **35**
- atomize (attribute), **145**
- atomize (method), **145**
- attribute_setup (command), **66**
- attributes (syntax), **35**
- auto (method), **142**

- ax_specification (HOL command),
178
- axiomatization (command), **51**
- axioms (command), **72**
- axmdecl (syntax), **36**

- back (command), **92**
- best (method), **141**
- bestsimp (method), **142**
- blast (method), **141**
- break (antiquotation option), 46
- by (command), **86**

- calculation (fact), 95
- case (command), **97**
- case_conclusion (attribute), **97**
- case_names (attribute), **97**
- case_tac (HOL method), **168**
- case_tac (ZF method), **185**
- cases (attribute), **104**
- cases (method), 85, 99, **100**
- cd (command), **127**
- chapter (command), **39**
- clamod (syntax), **141**
- clarify (method), **141**
- clarsimp (method), **142**
- clasimpmod (syntax), **142**
- class (antiquotation), **41**
- class (command), **62**
- class_deps (command), **62**
- classdecl (syntax), **31**
- classes (command), **68**
- classrel (command), **68**
- codatatype (ZF command), **183**
- code (HOL attribute), **169**, **176**

- code_abort (HOL command), **169**
- code_class (HOL command), **169**
- code_const (HOL command), **169**
- code_datatype (HOL command), **169**
- code_deps (HOL command), **169**
- code_include (HOL command), **169**
- code_inline (HOL attribute), **169**
- code_instance (HOL command), **169**
- code_library (HOL command), **176**
- code_module (HOL command), **176**
- code_modulename (HOL command), **169**
- code_monad (HOL command), **169**
- code_post (HOL attribute), **169**
- code_reflect (HOL command), **169**
- code_reserved (HOL command), **169**
- code_thms (HOL command), **169**
- code_type (HOL command), **169**
- coherent (HOL method), **165**
- coinduct (attribute), **104**
- coinduct (method), **100**
- coinductive (HOL command), **162**
- coinductive (ZF command), **183**
- coinductive_set (HOL command), **162**
- comment (syntax), **31**
- COMP (attribute), **131**
- cong (attribute), **138**
- const (antiquotation), **41**
- constrains (element), **56**
- consts (command), **71**
- consts (HOLCF command), **180**
- consts_code (HOL command), **176**
- consumes (attribute), **97**
- context (command), **50**
- contextelem (syntax), **56**
- contradiction (method), **140**
- corollary (command), **82**
- cut_tac (method), **134**
- datatype (HOL command), **153**
- datatype (ZF command), **183**
- declaration (command), **54**
- declare (command), **54**
- def (command), **77**
- default_sort (command), **68**
- defer (command), **92**
- defines (element), **56**
- definition (command), **51**
- defn (attribute), **51**
- defs (command), **71**
- dest (attribute), **144**
- dest (Pure attribute), **89**
- discharge (inference), **12**
- display (antiquotation option), **45**
- display_drafts (command), **47**
- domain (HOLCF command), **180**
- done (command), **92**
- drule (method), **130**
- drule_tac (method), **134**
- elim (attribute), **144**
- elim (method), **140**
- elim (Pure attribute), **89**
- elim_format (Pure attribute), **131**
- elim_resolution (inference), **9**
- end (global command), **49**
- end (local command), **50, 63**
- enriched_type (HOL command), **154**
- erule (method), **130**
- erule_tac (method), **134**
- eta_contract (antiquotation option), **45**
- eta_contract (ML), **108**
- expansion (inference), **12**
- export_code (HOL command), **169**
- fact (method), **36, 89**
- fail (method), **130**
- fast (method), **141**
- fastsimp (method), **142**
- file (antiquotation), **41**

- finally (command), **95**
- find_consts (command), **124**
- find_theorems (command), **124**
- finish (inference), **8**
- fix (command), **77**
- fixes (element), **56**
- float (syntax), **28**
- float_const (inner syntax), **120**
- float_token (inner syntax), **119**
- fold (method), **130**
- folded (attribute), **131**
- force (method), **142**
- from (command), **80**
- frule (method), **130**
- frule_tac (method), **134**
- full_prf (antiquotation), **41**
- full_prf (command), **106**
- fun (HOL command), **155**
- function (HOL command), **155**
- Goal_Display.goals_limit (ML), **108**
- Goal_Display.show_main_goal (ML), **108**
- goals (antiquotation), **41**
- goals_limit (antiquotation option), **46**
- goalspec (syntax), **86**
- guess (command), **93**
- have (command), **82**
- header (command), **39**
- hence (command), **82**
- hide_class (command), **73**
- hide_const (command), **73**
- hide_fact (command), **73**
- hide_type (command), **73**
- hypsubst (method), **132**
- id (inner syntax), **119**
- ident (syntax), **28**, **119**
- idt (inner syntax), **116**, **117**, **118**
- idts (inner syntax), **116**, **118**
- iff (attribute), **144**
- in (keyword), **51**
- ind_cases (HOL method), **168**
- ind_cases (ZF method), **185**
- indent (antiquotation option), **46**
- induct (attribute), **104**
- induct (method), **82**, **99**, **100**
- induct_simp (attribute), **103**
- induct_tac (HOL method), **168**
- induct_tac (ZF method), **185**
- inductive (HOL command), **162**
- inductive (ZF command), **183**
- inductive_cases (HOL command), **168**
- inductive_cases (ZF command), **185**
- inductive_set (HOL command), **162**
- infix (syntax), **110**
- init (inference), **8**
- insert (method), **130**
- inst (syntax), **33**
- instance (command), **62**, **69**
- instantiation (command), **62**, **70**
- insts (syntax), **33**
- int (syntax), **31**
- interp (syntax), **59**
- interpret (command), **59**
- interpretation (command), **59**
- intro (attribute), **144**
- intro (method), **140**
- intro (Pure attribute), **89**
- intro_classes (method), **62**
- intro_locales (method), **56**
- iprover (HOL method), **165**
- is (keyword), **79**
- judgment (command), **145**
- kill (command), **77**, **127**
- lemma (antiquotation), **41**
- lemma (command), **82**
- lemmas (command), **72**
- let (command), **79**

- lexicographic_order (HOL method), **157**
- linear_undo (command), **127**
- local_setup (command), **66**
- locale (command), **56**
- localeexpr (syntax), **55**
- logic (inner syntax), **116**, 117
- long_names (antiquotation option), **45**
- long_names (ML), **108**
- longid (inner syntax), **119**
- longident (syntax), **28**, 119
- margin (antiquotation option), **46**
- method (syntax), **85**
- method_setup (command), **93**
- mixfix (syntax), **110**
- ML (antiquotation), **41**
- ML (command), **66**
- ML_command (command), **66**
- ML_prf (command), **66**
- ML_struct (antiquotation), **41**
- ML_type (antiquotation), **41**
- ML_val (command), **66**
- mode (antiquotation option), **46**
- mono (HOL attribute), **162**
- moreover (command), **95**
- name (syntax), **30**
- nameref (syntax), **30**
- nat (syntax), **28**, 28, 119
- next (command), **76**
- no_notation (command), **113**
- no_syntax (command), **120**
- no_translations (command), **120**
- no_type_notation (command), **113**
- no_vars (attribute), **131**
- nonterminal (command), **120**
- notation (command), **113**
- note (command), **80**
- notepad (command), **75**
- notes (element), **56**
- nothing (fact), 81
- num (inner syntax), **119**
- num_const (inner syntax), **120**
- obtain (command), **93**
- obtains (element), **82**, 85
- OF (attribute), **89**
- of (attribute), **89**
- oops (command), **76**
- oracle (command), **73**
- output (keyword), 121
- overloading (command), **65**
- params (attribute), **97**
- parname (syntax), **30**
- parse_ast_translation (command), **121**
- parse_translation (command), **121**
- partial_function (HOL command), **159**
- partial_function_mono (HOL attribute), **159**
- pat_completeness (HOL method), **157**
- pr (command), **106**
- prefer (command), **92**
- prems (fact), 79
- presume (command), **77**
- Pretty.margin_default (ML), **110**
- prf (antiquotation), **41**
- prf (command), **106**
- primrec (HOL command), **155**
- primrec (ZF command), **184**
- print_abbrevs (command), **51**
- print_ast_translation (command), **121**
- print_attributes (command), **124**
- print_binds (command), **124**
- print_cases (command), **97**
- print_claset (command), **144**

- print_classes (command), **62**
- print_codeproc (HOL command), **169**
- print_codesetup (HOL command), **169**
- print_commands (command), **124**
- print_configs (command), **129**
- print_dependencies (command), **59**
- print_depth (ML), **110**
- print_drafts (command), **47**
- print_facts (command), **124**
- print_induct_rules (command), **104**
- print_interps (command), **59**
- print_locale (command), **56**
- print_locales (command), **56**
- print_methods (command), **124**
- print_simpset (command), **138**
- print_statement (command), **82**
- print_syntax (command), **122**
- print_tcset (ZF command), **182**
- print_theorems (command), **124**
- print_theory (command), **124**
- print_trans_rules (command), **95**
- print_translation (command), **121**
- proof
 - default, 88
 - fake, 88
 - terminal, 88
 - trivial, 88
- proof (command), 81, 82, **86**, 87, 90
- prop (antiquotation), **41**
- prop (command), **106**
- prop (inner syntax), **115**, 117
- prop (syntax), **32**
- proppat (syntax), **34**
- props (syntax), **34**
- pttrn (inner syntax), **116**, 118
- pttrns (inner syntax), **116**, 118
- pwd (command), **127**
- qed (command), **86**, 87
- quickcheck (HOL command), **166**
- quickcheck_params (HOL command), **166**
- quotes (antiquotation option), **46**
- raw_tactic (method), **134**
- real (syntax), **31**
- recdef (HOL command), **160**
- recdef_cong (HOL attribute), **161**
- recdef_simp (HOL attribute), **161**
- recdef_tc (HOL command), **160**
- recdef_wf (HOL attribute), **161**
- record (HOL command), **151**
- relation (HOL method), **157**
- rename_tac (method), **134**
- rep_datatype (HOL command), **153**
- resolution (inference), **9**
- rotate_tac (method), **134**
- rotated (attribute), **131**
- rule (attribute), **89**, **144**
- rule (method), 81, 87, **89**, 90, **140**
- rule_format (attribute), **145**
- rule_tac (method), **134**
- rulify (attribute), **145**
- safe (method), **141**
- schematic_corollary (command), **82**
- schematic_lemma (command), **82**
- schematic_theorem (command), **82**
- sect (command), **39**
- section (command), **39**
- selection (syntax), **36**
- setup (command), **66**
- short_names (antiquotation option), **45**
- short_names (ML), **108**
- show (command), 78, **82**, 87
- show_abbrevs (antiquotation option), **45**
- show_abbrevs (ML), **108**
- show_brackets (ML), **108**

- show_consts (ML), **108**
- show_hyps (ML), **108**
- show_question_marks (ML), **108**
- show_sorts (antiquotation option), **45**
- show_sorts (ML), **108**
- show_structs (antiquotation option), **45**
- show_tags (ML), **108**
- show_types (antiquotation option), **45**
- show_types (ML), **108**
- shows (element), **82**
- simp (attribute), **138**
- simp (method), **136**
- simp_all (method), **136**
- simplified (attribute), **139**
- simpmod (syntax), **136**
- simproc_setup (command), **138**
- size_change (HOL method), **157**
- slow (method), **141**
- slowsimp (method), **142**
- sorry (command), **76, 86**
- sort (inner syntax), **117, 118**
- sort (syntax), **31**
- source (antiquotation option), **46**
- specification (HOL command), **178**
- split (attribute), **138**
- split (method), **132**
- standard (attribute), **131**
- string (syntax), **29, 29**
- structmixfix (syntax), **110**
- subclass (command), **62, 69**
- subgoal_tac (method), **134**
- subgoals (antiquotation), **41**
- sublocale (command), **59**
- subset (command), **39**
- subsection (command), **39**
- subst (method), **132**
- subsubset (command), **39**
- subsubsection (command), **39**
- succeed (method), **130**
- swapped (attribute), **145**
- symident (syntax), **28**
- syntax (command), **120**
- syntax_declaration (command), **54**
- tactic (method), **134**
- tagged (attribute), **131**
- tags (syntax), **46**
- target (syntax), **50**
- TC (ZF attribute), **182**
- term (antiquotation), **41**
- term (command), **106**
- term (syntax), **32**
- term abbreviations, **80**
- term_type (antiquotation), **41**
- termination (HOL command), **155**
- termpat (syntax), **34**
- text (antiquotation), **41**
- text (command), **39**
- text (syntax), **31**
- text_raw (command), **39**
- THEN (attribute), **131**
- then (command), **80, 82**
- theorem (command), **82**
- theorems (command), **72**
- theory (antiquotation), **41**
- theory (command), **49**
- thesis (variable), **80**
- thin_tac (method), **134**
- this (fact), **75, 80**
- this (method), **89**
- this (variable), **80**
- thm (antiquotation), **41**
- thm (command), **106**
- thm_deps (command), **124**
- thmdecl (syntax), **36**
- thmdef (syntax), **36**
- thmref (syntax), **36**
- thmrefs (syntax), **36**
- thus (command), **82**
- tid (inner syntax), **119**

- translations (command), **120**
- tvar (inner syntax), **119**
- txt (command), **39**
- txt_raw (command), **39**
- typ (antiquotation), **41**
- typ (command), **106**
- type (antiquotation), **41**
- type (inner syntax), **117**, 118
- type (syntax), **32**
- type_notation (command), **113**
- type_synonym (command), **69**
- typecheck (ZF method), **182**
- typed_print_translation (command),
121
- typeddecl (command), **69**
- typedef (HOL command), **148**
- typefree (syntax), **29**, 119
- typeof (antiquotation), **41**
- types_code (HOL command), **176**
- typespec (syntax), **33**
- typespecsorts (syntax), **33**
- typevar (syntax), **29**, 29, 119
- ultimately (command), **95**
- undo (command), **127**
- unfold (method), **130**
- unfold_locales (method), **56**
- unfolded (attribute), **131**
- unfolding (command), **80**
- unique_names (antiquotation option), **45**
- unique_names (ML), **108**
- untagged (attribute), **131**
- unused_thms (command), **124**
- use (command), 50, **66**
- use_thy (command), **127**
- uses (keyword), **50**, 67
- using (command), **80**
- value (HOL command), **166**
- var (inner syntax), **119**
- var (syntax), **28**, 29, 119
- vars (syntax), **34**
- verbatim (syntax), **29**, 29
- weak-discharge (inference), **12**
- where (attribute), **89**
- with (command), **80**
- write (command), **113**
- xnum (inner syntax), **119**
- xstr (inner syntax), **119**