



The Isabelle/Isar Reference Manual

Makarius Wenzel

With Contributions by Clemens Ballarin, Stefan Berghofer,
Jasmin Blanchette, Timothy Bourke, Lukas Bulwahn,
Amine Chaieb, Lucas Dixon, Florian Haftmann,
Brian Huffman, Lars Hupel, Gerwin Klein,
Alexander Krauss, Ondřej Kunčar, Andreas Lochbihler,
Tobias Nipkow, Lars Noschinski, David von Oheimb,
Larry Paulson, Sebastian Skalberg, Christian Sternagel

August 27, 2014

Preface

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.

In its pioneering times, the Isabelle/Isar version of the *Proof General* user interface [2, 3] has contributed to the success of for interactive theory and proof development in this advanced theorem proving environment, even though it was somewhat biased towards old-style proof scripts. The more recent Isabelle/jEdit Prover IDE [56] emphasizes the document-oriented approach of Isabelle/Isar again more explicitly.

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 [53, 54]. *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 Isabelle/HOL are described in part III. Although the main language elements are already provided by the Isabelle/Pure framework, examples given in the generic parts will usually refer to Isabelle/HOL.

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.

Contents

I	Basic Concepts	1
1	Synopsis	2
1.1	Notepad	2
1.1.1	Types and terms	2
1.1.2	Facts	2
1.1.3	Block structure	5
1.2	Calculational reasoning	6
1.2.1	Special names in Isar proofs	6
1.2.2	Transitive chains	7
1.2.3	Degenerate calculations and bigstep reasoning	8
1.3	Induction	9
1.3.1	Induction as Natural Deduction	9
1.3.2	Induction with local parameters and premises	11
1.3.3	Implicit induction context	12
1.3.4	Advanced induction with term definitions	13
1.4	Natural Deduction	13
1.4.1	Rule statements	13
1.4.2	Isar context elements	15
1.4.3	Pure rule composition	16
1.4.4	Structured backward reasoning	16
1.4.5	Structured rule application	17
1.4.6	Example: predicate logic	18
1.5	Generalized elimination and cases	22
1.5.1	General elimination rules	22
1.5.2	Rules with cases	23
1.5.3	Obtaining local contexts	24
2	The Isabelle/Isar Framework	25

2.1	The Pure framework	27
2.1.1	Primitive inferences	28
2.1.2	Reasoning with rules	29
2.2	The Isar proof language	31
2.2.1	Context elements	32
2.2.2	Structured statements	34
2.2.3	Structured proof refinement	35
2.2.4	Calculational reasoning	37
2.3	Example: First-Order Logic	38
2.3.1	Equational reasoning	39
2.3.2	Basic group theory	40
2.3.3	Propositional logic	41
2.3.4	Classical logic	43
2.3.5	Quantifiers	44
2.3.6	Canonical reasoning patterns	45
II	General Language Elements	48
3	Outer syntax — the theory language	49
3.1	Commands	50
3.2	Lexical matters	50
3.3	Common syntax entities	53
3.3.1	Names	53
3.3.2	Numbers	53
3.3.3	Comments	54
3.3.4	Type classes, sorts and arities	54
3.3.5	Types and terms	55
3.3.6	Term patterns and declarations	57
3.3.7	Attributes and theorems	58
4	Document preparation	62
4.1	Markup commands	62
4.2	Document Antiquotations	65
4.2.1	Styled antiquotations	69
4.2.2	General options	70

4.3	Markup via command tags	71
4.4	Railroad diagrams	72
4.5	Draft presentation	76
5	Specifications	77
5.1	Defining theories	77
5.2	Local theory targets	79
5.3	Bundled declarations	81
5.4	Term definitions	82
5.5	Axiomatizations	84
5.6	Generic declarations	85
5.7	Locales	86
5.7.1	Locale expressions	87
5.7.2	Locale declarations	88
5.7.3	Locale interpretation	91
5.8	Classes	95
5.8.1	The class target	97
5.8.2	Co-regularity of type classes and arities	98
5.9	Unrestricted overloading	98
5.10	Incorporating ML code	99
5.11	Primitive specification elements	102
5.11.1	Sorts	102
5.11.2	Types	103
5.11.3	Constants and definitions	103
5.12	Naming existing theorems	105
5.13	Oracles	106
5.14	Name spaces	106
6	Proofs	108
6.1	Proof structure	108
6.1.1	Formal notepad	108
6.1.2	Blocks	109
6.1.3	Omitting proofs	110
6.2	Statements	110
6.2.1	Context elements	110

6.2.2	Term abbreviations	112
6.2.3	Facts and forward chaining	113
6.2.4	Goals	115
6.3	Refinement steps	118
6.3.1	Proof method expressions	118
6.3.2	Initial and terminal proof steps	120
6.3.3	Fundamental methods and attributes	122
6.3.4	Emulating tactic scripts	126
6.3.5	Defining proof methods	127
6.4	Generalized elimination	128
6.5	Calculational reasoning	130
6.6	Proof by cases and induction	132
6.6.1	Rule contexts	132
6.6.2	Proof methods	135
6.6.3	Declaring rules	140
7	Inner syntax — the term language	142
7.1	Printing logical entities	142
7.1.1	Diagnostic commands	142
7.1.2	Details of printed content	145
7.1.3	Alternative print modes	147
7.2	Mixfix annotations	148
7.2.1	The general mixfix form	149
7.2.2	Infixes	151
7.2.3	Binders	151
7.3	Explicit notation	152
7.4	The Pure syntax	154
7.4.1	Lexical matters	154
7.4.2	Priority grammars	155
7.4.3	The Pure grammar	156
7.4.4	Inspecting the syntax	160
7.4.5	Ambiguity of parsed expressions	161
7.5	Syntax transformations	161
7.5.1	Abstract syntax trees	162
7.5.2	Raw syntax and translations	164

7.5.3	Syntax translation functions	169
7.5.4	Built-in syntax transformations	172
8	Other commands	175
8.1	Inspecting the context	175
9	Generic tools and packages	179
9.1	Configuration options	179
9.2	Basic proof tools	180
9.2.1	Miscellaneous methods and attributes	180
9.2.2	Low-level equational reasoning	183
9.2.3	Further tactic emulations	184
9.3	The Simplifier	188
9.3.1	Simplification methods	188
9.3.2	Declaring rules	193
9.3.3	Ordered rewriting with permutative rules	196
9.3.4	Simplifier tracing and debugging	198
9.3.5	Simplification procedures	200
9.3.6	Configurable Simplifier strategies	202
9.3.7	Forward simplification	206
9.4	The Classical Reasoner	207
9.4.1	Basic concepts	207
9.4.2	Rule declarations	211
9.4.3	Structured methods	213
9.4.4	Fully automated methods	214
9.4.5	Partially automated methods	218
9.4.6	Single-step tactics	219
9.4.7	Modifying the search step	220
9.5	Object-logic setup	221
9.6	Tracing higher-order unification	223
III	Isabelle/HOL	224
10	Higher-Order Logic	225

11	Derived specification elements	227
11.1	Inductive and coinductive definitions	227
11.1.1	Derived rules	229
11.1.2	Monotonicity theorems	230
11.2	Recursive functions	231
11.2.1	Proof methods related to recursive definitions	236
11.2.2	Functions with explicit partiality	238
11.2.3	Old-style recursive function definitions (TFL)	239
11.3	Datatypes	241
11.4	Records	243
11.4.1	Basic concepts	243
11.4.2	Record specifications	244
11.4.3	Record operations	245
11.4.4	Derived rules and proof tools	246
11.5	Typedef axiomatization	247
11.6	Functorial structure of types	250
11.7	Quotient types	251
11.8	Definition by specification	254
11.9	Adhoc overloading of constants	255
12	Proof tools	257
12.1	Adhoc tuples	257
12.2	Transfer package	257
12.3	Lifting package	259
12.4	Coercive subtyping	264
12.5	Arithmetic proof support	265
12.6	Intuitionistic proof search	266
12.7	Model Elimination and Resolution	266
12.8	Algebraic reasoning via Gröbner bases	267
12.9	Coherent Logic	268
12.10	Proving propositions	269
12.11	Checking and refuting propositions	271
12.12	Unstructured case analysis and induction	275
13	Executable code	278

IV	Appendix	290
A	Isabelle/Isar quick reference	291
A.1	Proof commands	291
A.1.1	Primitives and basic syntax	291
A.1.2	Abbreviations and synonyms	292
A.1.3	Derived elements	292
A.1.4	Diagnostic commands	292
A.2	Proof methods	293
A.3	Attributes	294
A.4	Rule declarations and methods	294
A.5	Emulating tactic scripts	295
A.5.1	Commands	295
A.5.2	Methods	295
B	Predefined Isabelle symbols	296
C	ML tactic expressions	302
C.1	Resolution tactics	302
C.2	Simplifier tactics	303
C.3	Classical Reasoner tactics	303
C.4	Miscellaneous tactics	304
C.5	Tacticals	304
	Bibliography	305
	Index	310

Part I

Basic Concepts

Synopsis

1.1 Notepad

An Isar proof body serves as mathematical notepad to compose logical content, consisting of types, terms, facts.

1.1.1 Types and terms

notepad
begin

Locally fixed entities:

fix x — local constant, without any type information yet
fix $x :: 'a$ — variant with explicit type-constraint for subsequent use

fix $a b$
assume $a = b$ — type assignment at first occurrence in concrete term

Definitions (non-polymorphic):

def $x \equiv t :: 'a$

Abbreviations (polymorphic):

let $?f = \lambda x. x$
term $?f ?f$

Notation:

write $x (***)$
end

1.1.2 Facts

A fact is a simultaneous list of theorems.

Producing facts

notepad
begin

Via assumption (“lambda”):

assume $a: A$

Via proof (“let”):

have $b: B$ **sorry**

Via abbreviation (“let”):

note $c = a\ b$

end

Referencing facts

notepad
begin

Via explicit name:

assume $a: A$

note a

Via implicit name:

assume A

note $this$

Via literal proposition (unification with results from the proof text):

assume A

note ‘ A ’

assume $\bigwedge x. B\ x$

note ‘ $B\ a$ ’

note ‘ $B\ b$ ’

end

Manipulating facts

notepad
begin

Instantiation:

```

assume a:  $\bigwedge x. B\ x$ 
note a
note a [of b]
note a [where  $x = b$ ]

```

Backchaining:

```

assume 1: A
assume 2:  $A \implies C$ 
note 2 [OF 1]
note 1 [THEN 2]

```

Symmetric results:

```

assume  $x = y$ 
note this [symmetric]

```

```

assume  $x \neq y$ 
note this [symmetric]

```

Adhoc-simplification (take care!):

```

assume P ( $[] @ xs$ )
note this [simplified]
end

```

Projections

Isar facts consist of multiple theorems. There is notation to project interval ranges.

```

notepad
begin
  assume stuff: A B C D
  note stuff(1)
  note stuff(2–3)
  note stuff(2–)
end

```

Naming conventions

- Lower-case identifiers are usually preferred.
- Facts can be named after the main term within the proposition.
- Facts should *not* be named after the command that introduced them (**assume**, **have**). This is misleading and hard to maintain.

- Natural numbers can be used as “meaningless” names (more appropriate than a_1 , a_2 etc.)
- Symbolic identifiers are supported (e.g. $*$, $**$, $***$).

1.1.3 Block structure

The formal notepad is block structured. The fact produced by the last entry of a block is exported into the outer context.

```
notepad
begin
{
  have a: A sorry
  have b: B sorry
  note a b
}
note this
note 'A'
note 'B'
end
```

Explicit blocks as well as implicit blocks of nested goal statements (e.g. **have**) automatically introduce one extra pair of parentheses in reserve. The **next** command allows to “jump” between these sub-blocks.

```
notepad
begin
{
  have a: A sorry
next
  have b: B
  proof –
    show B sorry
next
  have c: C sorry
next
  have d: D sorry
qed
}
```

Alternative version with explicit parentheses everywhere:


```

{
  {
    have a: A sorry
  }
  {
    have b: B
    proof -
      {
        show B sorry
      }
      {
        have c: C sorry
      }
      {
        have d: D sorry
      }
    qed
  }
}
end

```

1.2 Calculational reasoning

For example, see `~/src/HOL/Isar_Examples/Group.thy`.

1.2.1 Special names in Isar proofs

- term *?thesis* — the main conclusion of the innermost pending claim
- term `...` — the argument of the last explicitly stated result (for infix application this is the right-hand side)
- fact *this* — the last result produced in the text

```

notepad
begin
  have x = y
  proof -
    term ?thesis
    show ?thesis sorry
    term ?thesis — static!
  qed
end

```

```

qed
term ...
thm this
end

```

Calculational reasoning maintains the special fact called “*calculation*” in the background. Certain language elements combine primary *this* with secondary *calculation*.

1.2.2 Transitive chains

The Idea is to combine *this* and *calculation* via typical *trans* rules (see also `print_trans_rules`):

```

thm trans
thm less_trans
thm less_le_trans

```

```

notepad
begin

```

Plain bottom-up calculation:

```

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

```

Variant using the ... abbreviation:

```

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

```

Top-down version with explicit claim at the head:

```

have a = d
proof -
  have a = b sorry

```

```

also
have ... =  $c$  sorry
also
have ... =  $d$  sorry
finally
show ?thesis .
qed
next

```

Mixed inequalities (require suitable base type):

```

fix  $a\ b\ c\ d :: nat$ 

have  $a < b$  sorry
also
have  $b \leq c$  sorry
also
have  $c = d$  sorry
finally
have  $a < d$  .
end

```

Notes

- The notion of *trans* rule is very general due to the flexibility of Isabelle/Pure rule composition.
- User applications may declare their own rules, with some care about the operational details of higher-order unification.

1.2.3 Degenerate calculations and bigstep reasoning

The Idea is to append *this* to *calculation*, without rule composition.

```

notepad
begin

```

A vacuous proof:

```

have  $A$  sorry
moreover
have  $B$  sorry
moreover
have  $C$  sorry

```

```

ultimately
  have A and B and C .
next

```

Slightly more content (trivial bigstep reasoning):

```

have A sorry
moreover
  have B sorry
moreover
  have C sorry
ultimately
  have A  $\wedge$  B  $\wedge$  C by blast
next

```

More ambitious bigstep reasoning involving structured results:

```

have A  $\vee$  B  $\vee$  C sorry
moreover
  { assume A have R sorry }
moreover
  { assume B have R sorry }
moreover
  { assume C have R sorry }
ultimately
  have R by blast — “big-bang integration” of proof blocks (occasionally fragile)
end

```

1.3 Induction

1.3.1 Induction as Natural Deduction

In principle, induction is just a special case of Natural Deduction (see also §1.4). For example:

```

thm nat.induct
print _statement nat.induct

notepad
begin
  fix n :: nat
  have P n
  proof (rule nat.induct) — fragile rule application!
    show P 0 sorry
  next

```

```

fix  $n :: nat$ 
assume  $P n$ 
show  $P (Suc n)$  sorry
qed
end

```

In practice, much more proof infrastructure is required.

The proof method *induct* provides:

- implicit rule selection and robust instantiation
- context elements via symbolic case names
- support for rule-structured induction statements, with local parameters, premises, etc.

```

notepad
begin
  fix  $n :: nat$ 
  have  $P n$ 
  proof (induct  $n$ )
    case 0
    show  $?case$  sorry
  next
    case ( $Suc n$ )
    from  $Suc.hyps$  show  $?case$  sorry
  qed
end

```

Example

The subsequent example combines the following proof patterns:

- outermost induction (over the datatype structure of natural numbers), to decompose the proof problem in top-down manner
- calculational reasoning (§1.2) to compose the result in each case
- solving local claims within the calculation by simplification

```

lemma
  fixes  $n :: nat$ 
  shows  $(\sum_{i=0..n} i) = n * (n + 1) \text{ div } 2$ 
proof (induct  $n$ )

```

```

case 0
have ( $\sum i=0..0. i$ ) = (0::nat) by simp
also have ... = 0 * (0 + 1) div 2 by simp
finally show ?case .
next
case (Suc n)
have ( $\sum i=0..Suc\ n. i$ ) = ( $\sum i=0..n. i$ ) + (n + 1) by simp
also have ... = n * (n + 1) div 2 + (n + 1) by (simp add: Suc.hyps)
also have ... = (n * (n + 1) + 2 * (n + 1)) div 2 by simp
also have ... = (Suc n * (Suc n + 1)) div 2 by simp
finally show ?case .
qed

```

This demonstrates how induction proofs can be done without having to consider the raw Natural Deduction structure.

1.3.2 Induction with local parameters and premises

Idea: Pure rule statements are passed through the induction rule. This achieves convenient proof patterns, thanks to some internal trickery in the *induct* method.

Important: Using compact HOL formulae with $\forall / \longrightarrow$ is a well-known anti-pattern! It would produce useless formal noise.

```

notepad
begin
  fix n :: nat
  fix P :: nat  $\Rightarrow$  bool
  fix Q :: 'a  $\Rightarrow$  nat  $\Rightarrow$  bool

  have P n
  proof (induct n)
    case 0
    show P 0 sorry
  next
    case (Suc n)
    from 'P n' show P (Suc n) sorry
  qed

  have A n  $\Longrightarrow$  P n
  proof (induct n)
    case 0
    from 'A 0' show P 0 sorry

```

```

next
  case (Suc n)
  from 'A n  $\implies$  P n'
    and 'A (Suc n)' show P (Suc n) sorry
qed

have  $\bigwedge x. Q x n$ 
proof (induct n)
  case 0
  show Q x 0 sorry
next
  case (Suc n)
  from ' $\bigwedge x. Q x n$ ' show Q x (Suc n) sorry

```

Local quantification admits arbitrary instances:

```

  note 'Q a n' and 'Q b n'
qed
end

```

1.3.3 Implicit induction context

The *induct* method can isolate local parameters and premises directly from the given statement. This is convenient in practical applications, but requires some understanding of what is going on internally (as explained above).

```

notepad
begin
  fix n :: nat
  fix Q :: 'a  $\implies$  nat  $\implies$  bool'

  fix x :: 'a'
  assume A x n
  then have Q x n
  proof (induct n arbitrary: x)
    case 0
    from 'A x 0' show Q x 0 sorry
  next
    case (Suc n)
    from ' $\bigwedge x. A x n \implies Q x n$ ' — arbitrary instances can be produced here
      and 'A x (Suc n)' show Q x (Suc n) sorry
  qed
end

```

1.3.4 Advanced induction with term definitions

Induction over subexpressions of a certain shape are delicate to formalize. The Isar *induct* method provides infrastructure for this.

Idea: sub-expressions of the problem are turned into a defined induction variable; often accompanied with fixing of auxiliary parameters in the original expression.

```

notepad
begin
  fix  $a :: 'a \Rightarrow nat$ 
  fix  $A :: nat \Rightarrow bool$ 

  assume  $A (a x)$ 
  then have  $P (a x)$ 
  proof (induct a x arbitrary: x)
    case 0
      note  $prem = 'A (a x)'$ 
      and  $defn = '0 = a x'$ 
      show  $P (a x)$  sorry
    next
      case (Suc n)
      note  $hyp = '\wedge x. n = a x \Longrightarrow A (a x) \Longrightarrow P (a x)'$ 
      and  $prem = 'A (a x)'$ 
      and  $defn = 'Suc n = a x'$ 
      show  $P (a x)$  sorry
  qed
end

```

1.4 Natural Deduction

1.4.1 Rule statements

Isabelle/Pure “theorems” are always natural deduction rules, which sometimes happen to consist of a conclusion only.

The framework connectives \wedge and \Longrightarrow indicate the rule structure declaratively. For example:

```

thm conjI
thm impI
thm nat.induct

```

The object-logic is embedded into the Pure framework via an implicit derivability judgment $Trueprop :: bool \Rightarrow prop$.

Thus any HOL formulae appears atomic to the Pure framework, while the rule structure outlines the corresponding proof pattern.

This can be made explicit as follows:

```
notepad
begin
  write Trueprop (Tr)

  thm conjI
  thm impI
  thm nat.induct
end
```

Isar provides first-class notation for rule statements as follows.

```
print_statement conjI
print_statement impI
print_statement nat.induct
```

Examples

Introductions and eliminations of some standard connectives of the object-logic can be written as rule statements as follows. (The proof “**by blast**” serves as sanity check.)

```
lemma (P ==> False) ==> ~ P by blast
```

```
lemma ~ P ==> P ==> Q by blast
```

```
lemma P ==> Q ==> P & Q by blast
```

```
lemma P & Q ==> (P ==> Q ==> R) ==> R by blast
```

```
lemma P ==> P v Q by blast
```

```
lemma Q ==> P v Q by blast
```

```
lemma P v Q ==> (P ==> R) ==> (Q ==> R) ==> R by blast
```

```
lemma (∧ x. P x) ==> (∀ x. P x) by blast
```

```
lemma (∀ x. P x) ==> P x by blast
```

```
lemma P x ==> (∃ x. P x) by blast
```

```
lemma (∃ x. P x) ==> (∧ x. P x ==> R) ==> R by blast
```

```
lemma x ∈ A ==> x ∈ B ==> x ∈ A ∩ B by blast
```

```
lemma x ∈ A ∩ B ==> (x ∈ A ==> x ∈ B ==> R) ==> R by blast
```

```
lemma x ∈ A ==> x ∈ A ∪ B by blast
```

lemma $x \in B \implies x \in A \cup B$ **by** *blast*

lemma $x \in A \cup B \implies (x \in A \implies R) \implies (x \in B \implies R) \implies R$ **by** *blast*

1.4.2 Isar context elements

We derive some results out of the blue, using Isar context elements and some explicit blocks. This illustrates their meaning wrt. Pure connectives, without goal states getting in the way.

notepad

begin

```
{
  fix x
  have B x sorry
}
```

have $\bigwedge x. B x$ **by** *fact*

next

```
{
  assume A
  have B sorry
}
```

have $A \implies B$ **by** *fact*

next

```
{
  def x  $\equiv$  t
  have B x sorry
}
```

have $B t$ **by** *fact*

next

```
{
  obtain x :: 'a where B x sorry
  have C sorry
}
```

have C **by** *fact*

end

1.4.3 Pure rule composition

The Pure framework provides means for:

- backward-chaining of rules by *resolution*
- closing of branches by *assumption*

Both principles involve higher-order unification of λ -terms modulo $\alpha\beta\eta$ -equivalence (cf. Huet and Miller).

notepad

begin

assume $a: A$ **and** $b: B$

thm *conjI*

thm *conjI* [*of* $A B$] — instantiation

thm *conjI* [*of* $A B$, *OF* $a b$] — instantiation and composition

thm *conjI* [*OF* $a b$] — composition via unification (trivial)

thm *conjI* [*OF* ' A ' ' B ']

thm *conjI* [*OF* *disjI1*]

end

Note: Low-level rule composition is tedious and leads to unreadable / unmaintainable expressions in the text.

1.4.4 Structured backward reasoning

Idea: Canonical proof decomposition via **fix** / **assume** / **show**, where the body produces a natural deduction rule to refine some goal.

notepad

begin

fix $A B :: 'a \Rightarrow bool$

have $\bigwedge x. A x \Longrightarrow B x$

proof —

fix x

assume $A x$

show $B x$ **sorry**

qed

have $\bigwedge x. A x \Longrightarrow B x$

proof —

```

{
  fix x
  assume A x
  show B x sorry
} — implicit block structure made explicit
note ‘ $\wedge x. A x \implies B x$ ’
  — side exit for the resulting rule
qed
end

```

1.4.5 Structured rule application

Idea: Previous facts and new claims are composed with a rule from the context (or background library).

```

notepad
begin
  assume r1:  $A \implies B \implies C$  — simple rule (Horn clause)

  have A sorry — prefix of facts via outer sub-proof
  then have C
  proof (rule r1)
    show B sorry — remaining rule premises via inner sub-proof
  qed

  have C
  proof (rule r1)
    show A sorry
    show B sorry
  qed

  have A and B sorry
  then have C
  proof (rule r1)
  qed

  have A and B sorry
  then have C by (rule r1)

next

  assume r2:  $A \implies (\wedge x. B1 x \implies B2 x) \implies C$  — nested rule

  have A sorry

```

```

then have  $C$ 
proof (rule r2)
  fix  $x$ 
  assume  $B1\ x$ 
  show  $B2\ x$  sorry
qed

```

The compound rule premise $\bigwedge x. B1\ x \implies B2\ x$ is better addressed via **fix** / **assume** / **show** in the nested proof body.

```

end

```

1.4.6 Example: predicate logic

Using the above principles, standard introduction and elimination proofs of predicate logic connectives of HOL work as follows.

```

notepad
begin
  have  $A \implies B$  and  $A$  sorry
  then have  $B$  ..

  have  $A$  sorry
  then have  $A \vee B$  ..

  have  $B$  sorry
  then have  $A \vee B$  ..

  have  $A \vee B$  sorry
  then have  $C$ 
  proof
    assume  $A$ 
    then show  $C$  sorry
  next
    assume  $B$ 
    then show  $C$  sorry
  qed

  have  $A$  and  $B$  sorry
  then have  $A \wedge B$  ..

  have  $A \wedge B$  sorry
  then have  $A$  ..

  have  $A \wedge B$  sorry

```

then have B ..

have $False$ **sorry**
then have A ..

have $True$..

have $\neg A$
proof
assume A
then show $False$ **sorry**
qed

have $\neg A$ **and** A **sorry**
then have B ..

have $\forall x. P x$
proof
fix x
show $P x$ **sorry**
qed

have $\forall x. P x$ **sorry**
then have $P a$..

have $\exists x. P x$
proof
show $P a$ **sorry**
qed

have $\exists x. P x$ **sorry**
then have C
proof
fix a
assume $P a$
show C **sorry**
qed

Less awkward version using **obtain**:

have $\exists x. P x$ **sorry**
then obtain a **where** $P a$..
end

Further variations to illustrate Isar sub-proofs involving **show**:

```

notepad
begin
  have  $A \wedge B$ 
  proof — two strictly isolated subproofs
    show  $A$  sorry
  next
    show  $B$  sorry
  qed

  have  $A \wedge B$ 
  proof — one simultaneous sub-proof
    show  $A$  and  $B$  sorry
  qed

  have  $A \wedge B$ 
  proof — two subproofs in the same context
    show  $A$  sorry
    show  $B$  sorry
  qed

  have  $A \wedge B$ 
  proof — swapped order
    show  $B$  sorry
    show  $A$  sorry
  qed

  have  $A \wedge B$ 
  proof — sequential subproofs
    show  $A$  sorry
    show  $B$  using ' $A$ ' sorry
  qed
end

```

Example: set-theoretic operators

There is nothing special about logical connectives (\wedge , \vee , \forall , \exists etc.). Operators from set-theory or lattice-theory work analogously. It is only a matter of rule declarations in the library; rules can be also specified explicitly.

```

notepad
begin
  have  $x \in A$  and  $x \in B$  sorry
  then have  $x \in A \cap B$  ..

```

```
have  $x \in A$  sorry
then have  $x \in A \cup B$  ..

have  $x \in B$  sorry
then have  $x \in A \cup B$  ..

have  $x \in A \cup B$  sorry
then have  $C$ 
proof
  assume  $x \in A$ 
  then show  $C$  sorry
next
  assume  $x \in B$ 
  then show  $C$  sorry
qed

next
have  $x \in \bigcap A$ 
proof
  fix  $a$ 
  assume  $a \in A$ 
  show  $x \in a$  sorry
qed

have  $x \in \bigcap A$  sorry
then have  $x \in a$ 
proof
  show  $a \in A$  sorry
qed

have  $a \in A$  and  $x \in a$  sorry
then have  $x \in \bigcup A$  ..

have  $x \in \bigcup A$  sorry
then obtain  $a$  where  $a \in A$  and  $x \in a$  ..
end
```


1.5 Generalized elimination and cases

1.5.1 General elimination rules

The general format of elimination rules is illustrated by the following typical representatives:

thm *exE* — local parameter

thm *conjE* — local premises

thm *disjE* — split into cases

Combining these characteristics leads to the following general scheme for elimination rules with cases:

- prefix of assumptions (or “major premises”)
- one or more cases that enable to establish the main conclusion in an augmented context

notepad

begin

assume *r*:

$A1 \implies A2 \implies$ (* assumptions *)

$(\bigwedge x y. B1 x y \implies C1 x y \implies R) \implies$ (* case 1 *)

$(\bigwedge x y. B2 x y \implies C2 x y \implies R) \implies$ (* case 2 *)

R (* main conclusion *)

have *A1* and *A2* sorry

then have *R*

proof (*rule r*)

fix *x y*

assume *B1 x y* and *C1 x y*

show *?thesis* sorry

next

fix *x y*

assume *B2 x y* and *C2 x y*

show *?thesis* sorry

qed

end

Here *?thesis* is used to refer to the unchanged goal statement.

1.5.2 Rules with cases

Applying an elimination rule to some goal, leaves that unchanged but allows to augment the context in the sub-proof of each case.

Isar provides some infrastructure to support this:

- native language elements to state eliminations
- symbolic case names
- method *cases* to recover this structure in a sub-proof

```
print_statement exE
print_statement conjE
print_statement disjE
```

```
lemma
  assumes A1 and A2 — assumptions
  obtains
    (case1) x y where B1 x y and C1 x y
  | (case2) x y where B2 x y and C2 x y
  sorry
```

Example

```
lemma tertium_non_datur:
  obtains
    (T) A
  | (F)  $\neg A$ 
  by blast
```

```
notepad
begin
  fix x y :: 'a
  have C
  proof (cases x = y rule: tertium_non_datur)
    case T
    from 'x = y' show ?thesis sorry
  next
    case F
    from 'x  $\neq$  y' show ?thesis sorry
  qed
end
```

Example

Isabelle/HOL specification mechanisms (datatype, inductive, etc.) provide suitable derived cases rules.

```
datatype foo = Foo | Bar foo
```

```
notepad
begin
  fix x :: foo
  have C
  proof (cases x)
    case Foo
    from 'x = Foo' show ?thesis sorry
  next
    case (Bar a)
    from 'x = Bar a' show ?thesis sorry
  qed
end
```

1.5.3 Obtaining local contexts

A single “case” branch may be inlined into Isar proof text via **obtain**. This proves $(\bigwedge x. B x \implies thesis) \implies thesis$ on the spot, and augments the context afterwards.

```
notepad
begin
  fix B :: 'a  $\Rightarrow$  bool

  obtain x where B x sorry
  note 'B x'
```

Conclusions from this context may not mention x again!

```
{
  obtain x where B x sorry
  from 'B x' have C sorry
}
note 'C'
end
```

The Isabelle/Isar Framework

Isabelle/Isar [53, 54, 30, 58, 55] 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 [49, 45, 59], using logical symbols for certain reasoning schemes where Mizar would prefer English words; see [60] 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 [39, 40], notably composition of higher-order natural deduction rules, which is a generalization of Gentzen’s original calculus [16]. 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 [33] (simply-typed set-theory) is being used most of the time; Isabelle/ZF [37] 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 $x \in B$	$\frac{x \in A \quad x \in B}{x \in A \cap B}$
then have $x \in A \cap B$..	

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	$[A][A \in \mathcal{A}]$
fix A	\vdots
assume $A \in \mathcal{A}$	$\frac{x \in A}{x \in \bigcap \mathcal{A}}$
show $x \in A$ <i><proof></i>	
qed	

This Isar reasoning pattern again refers to the primitive rule depicted above. The system determines it in the “**proof**” step, which could have been spelled 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 [39, 40] is an intuitionistic fragment of higher-order logic [14]. 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 [6], 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 [47]. Here object-logic inferences are internalized as formulae over \bigwedge and \Longrightarrow . Combining such rule statements may involve higher-order unification [38].

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:

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

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 $\forall :: (i \Rightarrow o) \Rightarrow o$, and axioms for object-level rules such as *conjI*: $A \Longrightarrow$

$B \implies A \wedge B$ or *allI*: $(\bigwedge x. B x) \implies \forall x. B x$. 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* [26] 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 \cap -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:

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

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:

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

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 $\#$:

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

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 \langle *proof* \rangle ” 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** *«inference»* element provides a general interface to hypotheses: “**assume** *«inference»* 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 *«inference»*, 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** *«weak-discharge»* A
def $x \equiv a \equiv$ **fix** x **assume** *«expansion»* $x \equiv a$

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

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

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

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** [54, §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** *«elimination case»* $\bar{A} \ \bar{x}$

$$\frac{\begin{array}{l} case: \Gamma \vdash \bigwedge thesis. (\bigwedge \bar{x}. \bar{A} \ \bar{x} \Longrightarrow thesis) \Longrightarrow thesis \\ result: \Gamma \cup \bar{A} \ \bar{y} \vdash B \end{array}}{\Gamma \vdash B} \text{ (elimination)}$$

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 ‘B’ </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:

```

statement  $\equiv$  name: props and ...
           | context* conclusion
context    $\equiv$  fixes vars and ...
           | assumes name: props and ...
conclusion  $\equiv$  shows name: props and ...
           | obtains vars and ... where name: props and ...
           | ...

```

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  $\bar{A} \bar{x} \mid \dots \equiv$ 
  fixes thesis
  assumes [intro]:  $\bigwedge \bar{x}. \bar{A} \bar{x} \implies \textit{thesis}$  and  $\dots$ 
  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:

<pre> theorem fixes x and y assumes $A x$ and $B y$ shows $C x y$ proof – from ‘$A x$’ and ‘$B y$’ show $C x y$ ⟨<i>proof</i>⟩ qed </pre>	<pre> theorem obtains x and y where $A x$ and $B y$ proof – have $A a$ and $B b$ ⟨<i>proof</i>⟩ then show <i>thesis</i> .. qed </pre>
--	---

Here local facts ‘ $A x$ ’ and ‘ $B y$ ’ 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 $\bigwedge x y. A x \implies B y \implies \textit{thesis}$ 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).

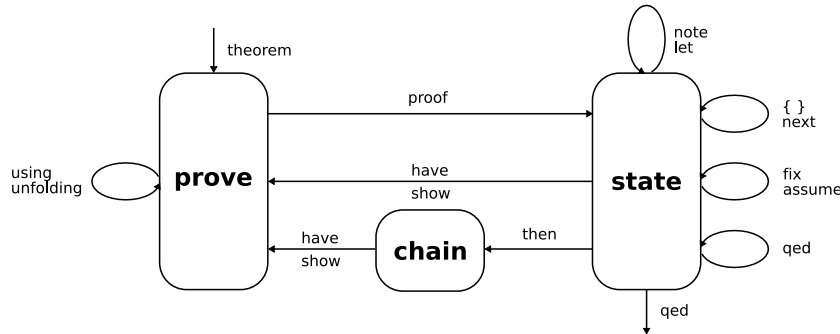


Figure 2.1: Isar/VM modes

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

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

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 $\bigwedge x y. A x \implies B y \implies C x y$ proof – fix x and y assume $A x$ and $B y$ show $C x y$ $\langle proof \rangle$ qed </pre>	<pre> have $\bigwedge x y. A x \implies B y \implies C x y$ proof – fix x assume $A x$ fix y assume $B y$ show $C x y$ $\langle proof \rangle$ qed </pre>
<pre> have $\bigwedge x y. A x \implies B y \implies C x y$ proof – fix y assume $B y$ fix x assume $A x$ show $C x y$ sorry qed </pre>	<pre> have $\bigwedge x y. A x \implies B y \implies C x y$ 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 [54, §6] and [5].

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 Base
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 judgment 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*]: $x = y \Longrightarrow B x \Longrightarrow B y$

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

theorem *sym* [*sym*]:

assumes $x = y$

shows $y = x$

proof –

have $x = x$..

with ‘ $x = y$ ’ **show** $y = x$..

qed

theorem *forw_subst* [*trans*]:

assumes $y = x$ **and** $B x$

shows $B y$

proof –

from ‘ $y = x$ ’ **have** $x = y$..

from *this* **and** ‘ $B x$ ’ **show** $B y$..

qed

theorem *back_subst* [*trans*]:

assumes $B x$ **and** $x = y$

```

  shows  $B y$ 
proof -
  from ' $x = y$ ' and ' $B x$ '
  show  $B y$  ..
qed

```

```

theorem trans [trans]:
  assumes  $x = y$  and  $y = z$ 
  shows  $x = z$ 
proof -
  from ' $y = z$ ' and ' $x = y$ '
  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 [16].

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

$disj :: o \Rightarrow o \Rightarrow o$ (**infix** \vee 30) **where**
 $disjI_1$ [*intro*]: $A \Longrightarrow A \vee B$ **and**
 $disjI_2$ [*intro*]: $B \Longrightarrow A \vee B$ **and**
 $disjE$ [*elim*]: $A \vee B \Longrightarrow (A \Longrightarrow C) \Longrightarrow (B \Longrightarrow C) \Longrightarrow C$

axiomatization

$conj :: o \Rightarrow o \Rightarrow o$ (**infix** \wedge 35) **where**
 $conjI$ [*intro*]: $A \Longrightarrow B \Longrightarrow A \wedge B$ **and**
 $conjD_1$: $A \wedge B \Longrightarrow A$ **and**
 $conjD_2$: $A \wedge B \Longrightarrow B$

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 ' $A \wedge B$ ' **show** A **by** (*rule* $conjD_1$)

from ' $A \wedge B$ ' **show** B **by** (*rule* $conjD_2$)

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 ' $A \Longrightarrow \perp$ ')
qed

```

```

theorem notE [elim]:
  assumes  $\neg A$  and  $A$ 
  shows  $B$ 
proof -
  from ' $\neg A$ ' have  $A \longrightarrow \perp$  unfolding not_def .
  from ' $A \longrightarrow \perp$ ' and ' $A$ ' have  $\perp$  ..
  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 ' $\neg \neg C$ ' 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 ' $\neg (C \vee \neg C)$ ' **show** \perp ..

qed

then have $C \vee \neg C$..

with ' $\neg (C \vee \neg C)$ ' **show** \perp ..

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 $(\neg C \implies C) \implies C$ 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 statement “**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 [14], 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*]: $(\bigwedge x. B x) \implies \forall x. B x$ **and**

allD [*dest*]: $(\forall x. B x) \implies B a$

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** ‘ $\exists x. \forall y. R\ x\ y$ ’ .. — \exists elimination

fix y **have** $R\ x\ y$ **using** ‘ $\forall y. R\ x\ y$ ’ .. — \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** $A \Longrightarrow B$ **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** $A \Longrightarrow \perp$ **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:


```

have  $A \longrightarrow B$ 
proof
  assume  $A$ 
  show  $B$   $\langle proof \rangle$ 
qed

```

```

have  $A \longrightarrow B$  and  $A$   $\langle proof \rangle$ 
then have  $B$  ..

```

```

have  $A$   $\langle proof \rangle$ 
then have  $A \vee B$  ..

```

```

have  $A \vee B$   $\langle proof \rangle$ 
then have  $C$ 
proof

```

```

have  $B$   $\langle proof \rangle$ 
then have  $A \vee B$  ..

```

```

  assume  $A$ 
  then show  $C$   $\langle proof \rangle$ 
next
  assume  $B$ 
  then show  $C$   $\langle proof \rangle$ 
qed

```

```

have  $A$  and  $B$   $\langle proof \rangle$ 
then have  $A \wedge B$  ..

```

```

have  $A \wedge B$   $\langle proof \rangle$ 
then obtain  $A$  and  $B$  ..

```

```

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 theory language

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 [57]). 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 Isabelle/jEdit [56] and Proof General [2] do not require explicit semicolons: command spans are determined by inspecting the content of the editor 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 Commands

`print _ commands*` : *any* →
`help*` : *any* →



`print _ commands` prints all outer syntax keywords and commands.

`help pats` retrieves outer syntax commands according to the specified name patterns.

Examples

Some common diagnostic commands are retrieved like this (according to usual naming conventions):

```
help print
help find
```

3.2 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.

```

    ident  = letter (subscript? quasiletter)*
    longident = ident(.ident)+
    symident = sym+ | \<ident>
    nat     = digit+
    float   = nat.nat | -nat.nat
    var     = ?ident | ?ident.nat
    typefree = 'ident
    typevar = ?typefree | ?typefree.nat
    string  = " ... "
    altstring = ' ... '
    cartouche = \<open> ... \<close>
    verbatim = {* ... *}

    letter = latin | \<latin> | \<latin latin> | greek |
    subscript = \<^sub>
    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.

A *cartouche* consists of arbitrary text, with properly balanced blocks of “\<open> ... \<close>”. Note that the rendering of cartouche delimiters is usually like this: “⟨ ... ⟩”.

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 §3.3.3).

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 L^AT_EX, Proof General, or Isabelle/jEdit. 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.


3.3 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.


3.3.1 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



3.3.2 Numbers



The outer lexical syntax (§3.2) 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*.




3.3.3 Comments

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

text*comment*


3.3.4 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*

3.3.5 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

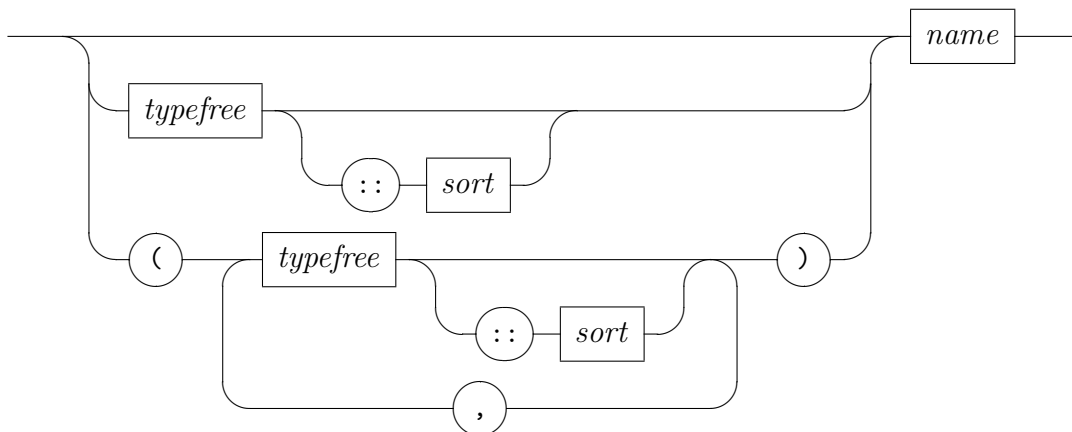


Type declarations and definitions usually refer to *typespec* on the left-hand side. This models basic type constructor application at the outer syntax level. Note that only plain postfix notation is available here, but no infixes.

typespec




typespec_sorts




3.3.6 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*.



term_pat



prop_pat



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.


3.3.7 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 token syntax *altstring* ‘ φ ’ or *cartouche* $\langle\varphi\rangle$ (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 “ $[[\textit{simproc } a]]$ ”) 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 PDF- \LaTeX , with support for hyperlinks and bookmarks within that format. This allows to produce papers, books, theses etc. from Isabelle theory sources.

\LaTeX output is generated while processing a *session* in batch mode, as explained in the *The Isabelle System Manual* [57]. The main Isabelle tools to get started with document preparation are `isabelle mkroot` and `isabelle build`.

The classic Isabelle/HOL tutorial [32] also explains some aspects of theory presentation.

4.1 Markup commands

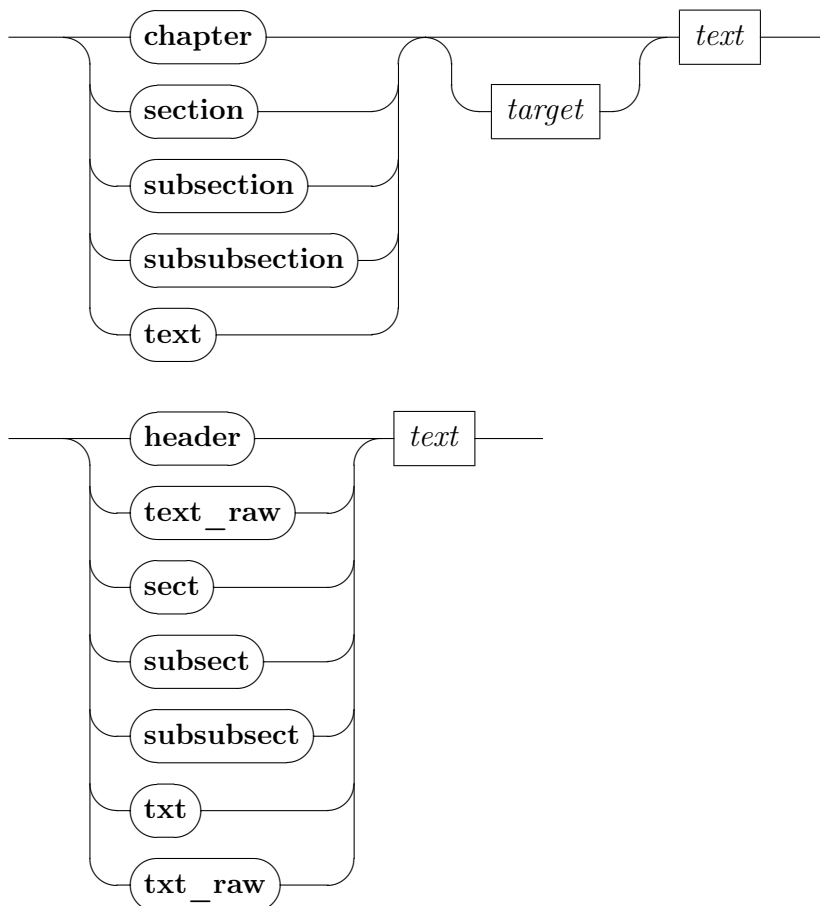
```

header   : toplevel  $\rightarrow$  toplevel
chapter : local_theory  $\rightarrow$  local_theory
section  : local_theory  $\rightarrow$  local_theory
subsection : local_theory  $\rightarrow$  local_theory
subsubsection : local_theory  $\rightarrow$  local_theory
text     : local_theory  $\rightarrow$  local_theory
text_raw : local_theory  $\rightarrow$  local_theory
sect     : proof  $\rightarrow$  proof
subject  : proof  $\rightarrow$  proof
subsubject : proof  $\rightarrow$  proof
txt      : proof  $\rightarrow$  proof
txt_raw  : proof  $\rightarrow$  proof

```

Markup commands provide a structured way to insert text into the document generated from a theory. Each markup command takes a single *text* argument, which is passed as argument to a corresponding \LaTeX macro. The default macros provided by `~/lib/texinputs/isabelle.sty` can be redefined according to the needs of the underlying document and \LaTeX styles.

Note that formal comments (§3.3.3) 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, **subsect**, 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_op : antiquotation
  ML_type : antiquotation
ML_structure : antiquotation
  ML_functor : antiquotation
  file : antiquotation
  url : antiquotation

```

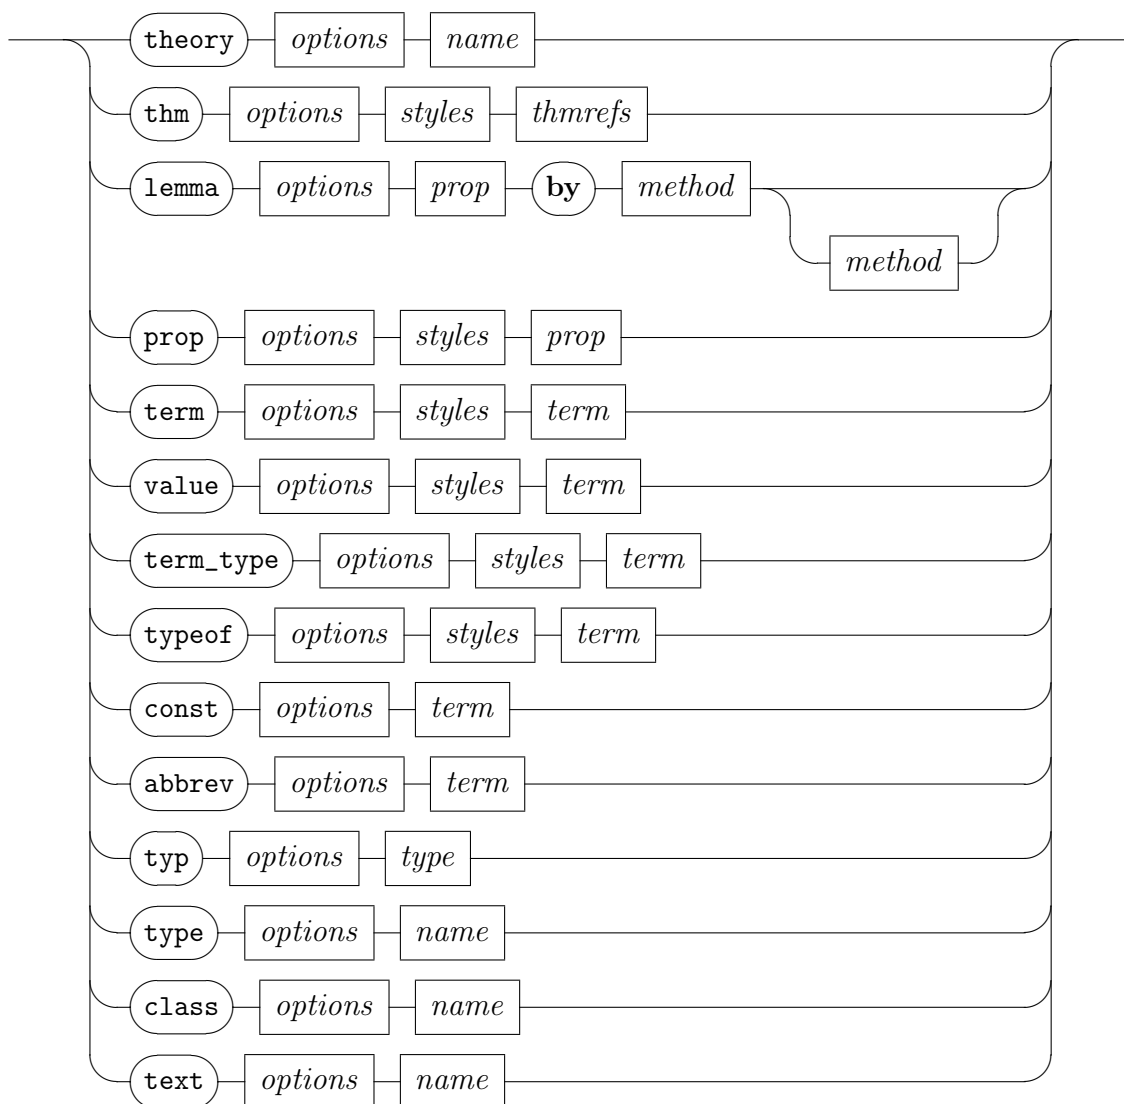
The overall content of an Isabelle/Isar theory may alternate between formal and informal text. The main body consists of formal specification and proof commands, interspersed with markup commands (§4.1) or document comments (§3.3.3). The argument of markup commands quotes informal text to be printed 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.

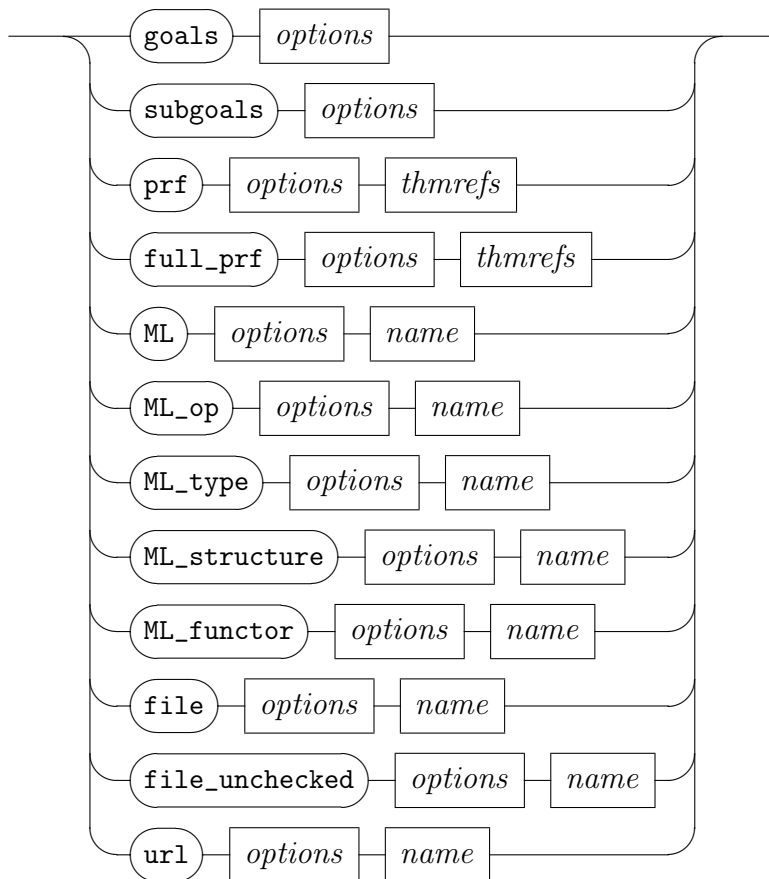
Antiquotations usually spare the author tedious typing of logical entities in full detail. Even more importantly, some degree of consistency-checking between the main body of formal text and its informal explanation is achieved, since terms and types appearing in antiquotations are checked within the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




antiquotation




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 (§3.3.7).
- `@{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 .
- `@{value t }` evaluates a term t and prints its result, see also **value**.
- `@{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 .
- `@{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.
- `@{typ τ }` prints a well-formed type τ .
- `@{type κ }` prints a (logical or syntactic) type constructor κ .
- `@{class c }` prints a class c .

$\@{\textit{text } s}$ 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.

$\@{\textit{goals}}$ 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.

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

$\@{\textit{prf } a_1 \dots a_n}$ 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.

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

$\@{\textit{ML } s}$, $\@{\textit{ML_op } s}$, $\@{\textit{ML_type } s}$, $\@{\textit{ML_structure } s}$, and $\@{\textit{ML_functor } s}$ check text s as ML value, infix operator, type, structure, and functor respectively. The source is printed verbatim.

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

$\@{\textit{file_unchecked } path}$ is like $\@{\textit{file } path}$, but does not check the existence of the $path$ within the file-system.

$\@{\textit{url } name}$ produces markup for the given URL, which results in an active hyperlink within the text.

4.2.1 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 $A_1 \implies \dots A_n \implies C$

4.2.2 General options

The following options are available to tune the printed output of antiquotations. Note that many of these coincide with system and configuration options 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.

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

names_short = 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.

names_unique = 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 subgoals 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 specifically for a logic session in the corresponding ROOT file.


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:

```

theory  theory begin/end
proof   all proof commands
ML     all commands involving ML code

```

The Isabelle document preparation system [57] 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).

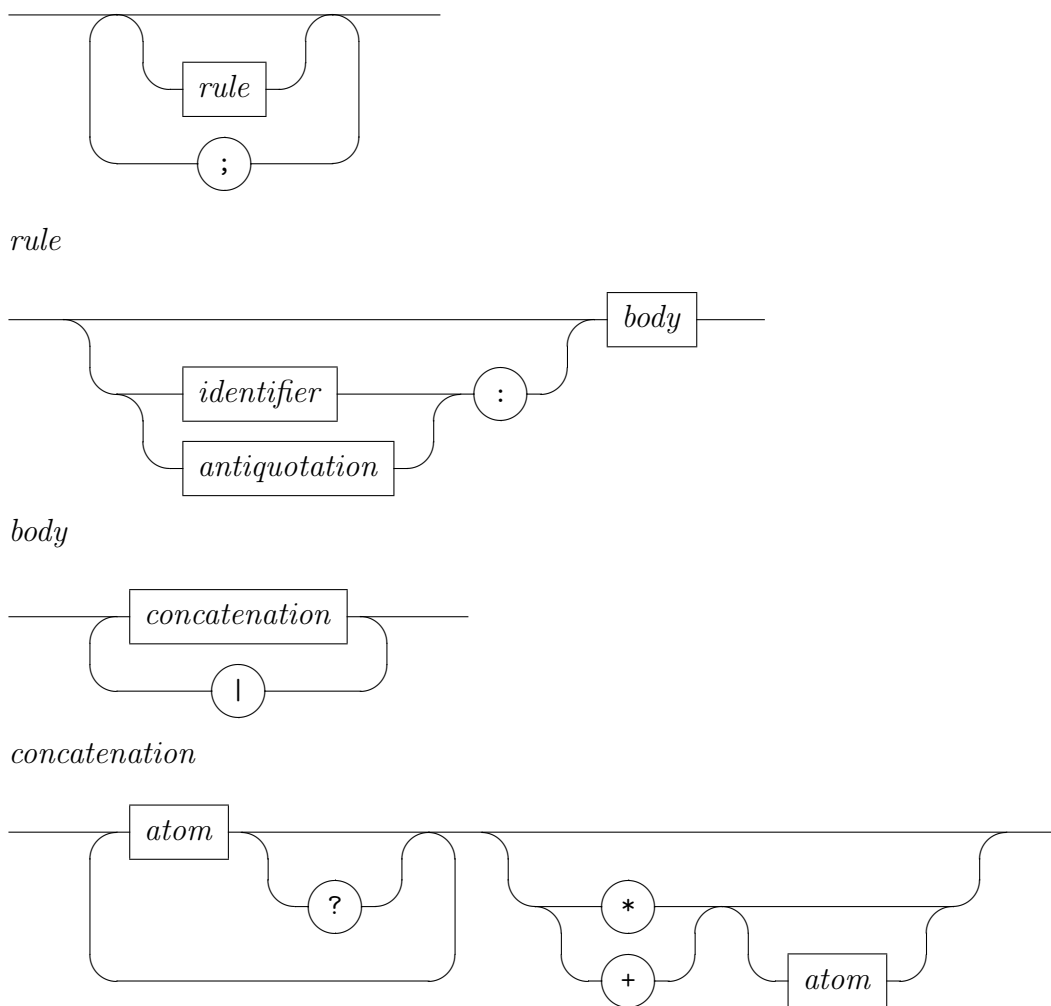
Command tags merely produce certain markup environments for typesetting. The meaning of these is determined by L^AT_EX macros, as defined in `~/lib/texinputs/isabelle.sty` or by the document author. The Isabelle document preparation tools also provide some high-level options to specify the meaning of 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 [57] for further details.

4.4 Railroad diagrams


rail : antiquotation



The *rail* antiquotation allows to include syntax diagrams into Isabelle documents. \LaTeX requires the style file `~/lib/texinputs/pdfsetup.sty`, which can be used via `\usepackage{pdfsetup}` in `root.tex`, for example. The rail specification language is quoted here as Isabelle *string* or text *cartouche*; it has its own grammar given below.



atom



The lexical syntax of *identifier* coincides with that of *ident* in regular Isabelle syntax, but *string* uses single quotes instead of double quotes of the standard *string* category.

Each *rule* defines a formal language (with optional name), using a notation that is similar to EBNF or regular expressions with recursion. The meaning and visual appearance of these rail language elements is illustrated by the following representative examples.

- Empty ()

- Nonterminal A

— A —

- Nonterminal via Isabelle antiquotation @`{syntax method}`

— *method* —


- Terminal 'xyz'

— xyz —


- Terminal in keyword style @'xyz'




- Terminal via Isabelle antiquotation @@{method rule}




- Concatenation A B C




- Newline inside concatenation A B C \<newline> D E F




- Variants A | B | C




- Option A ?




- Repetition A *




- Repetition with separator $A * sep$



- Strict repetition $A +$




- Strict repetition with separator $A + sep$



4.5 Draft presentation

`display_drafts*` : *any* \rightarrow



`display_drafts paths` performs simple output of a given list of raw source files. Only those symbols that do not require additional \LaTeX packages are displayed properly, everything else is left verbatim.

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 : toplevel → theory  
end   : theory → toplevel
```


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.




imports



keywords



keyword_decls



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 takes care that sources contributing to the development graph are always up-to-date: changed files are automatically rechecked whenever a theory header specification is processed.

The optional **keywords** specification declares outer syntax (chapter 3) that is introduced in this theory later on (rare in end-user applications). Both minor keywords and major keywords of the Isar command language need to be specified, in order to make parsing of proof documents work properly. Command keywords need to be classified according to their structural role in the formal text. Examples may be seen in Isabelle/HOL sources itself, such as **keywords** "typedef" :: *thy_goal*

or **keywords** "datatype" :: *thy_decl* for theory-level declarations with and without proof, respectively. Additional *tags* provide defaults for document preparation (§4.3).

It is possible to specify an alternative completion via `== text`, while the default is the corresponding keyword name.

end concludes the current theory definition. Note that some other commands, e.g. local theory targets **locale** or **class** may involve a **begin** that needs to be matched by **end**, according to the usual rules for nested blocks.


5.2 Local theory targets

context : *theory* → *local_theory*
end : *local_theory* → *theory*


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. §5.7) or type classes (cf. §5.8); the name “_” signifies the global theory context.

Unnamed contexts may introduce additional parameters and assumptions, and results produced in the context are generalized accordingly. Such auxiliary contexts may be nested within other targets, like **locale**, **class**, **instantiation**, **overloading**.



target



context *c* **begin** opens a named context, by recommencing an existing locale or class *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.

context *bundles elements* **begin** opens an unnamed context, by extending the enclosing global or local theory target by the given declaration bundles (§5.3) and context elements (**fixes**, **assumes** etc.). This means any results stemming from definitions and proofs in the extended context will be exported into the enclosing target by lifting over extra parameters and premises.

end concludes the current local theory, according to the nesting of contexts. 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 $c.a \ ?x \equiv t[?x]$ at the theory level (for arbitrary $?x$), together with a local abbreviation $c \equiv c.a \ x$ 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 $c.a: A[?x] \implies B[?x]$, again for arbitrary $?x$.

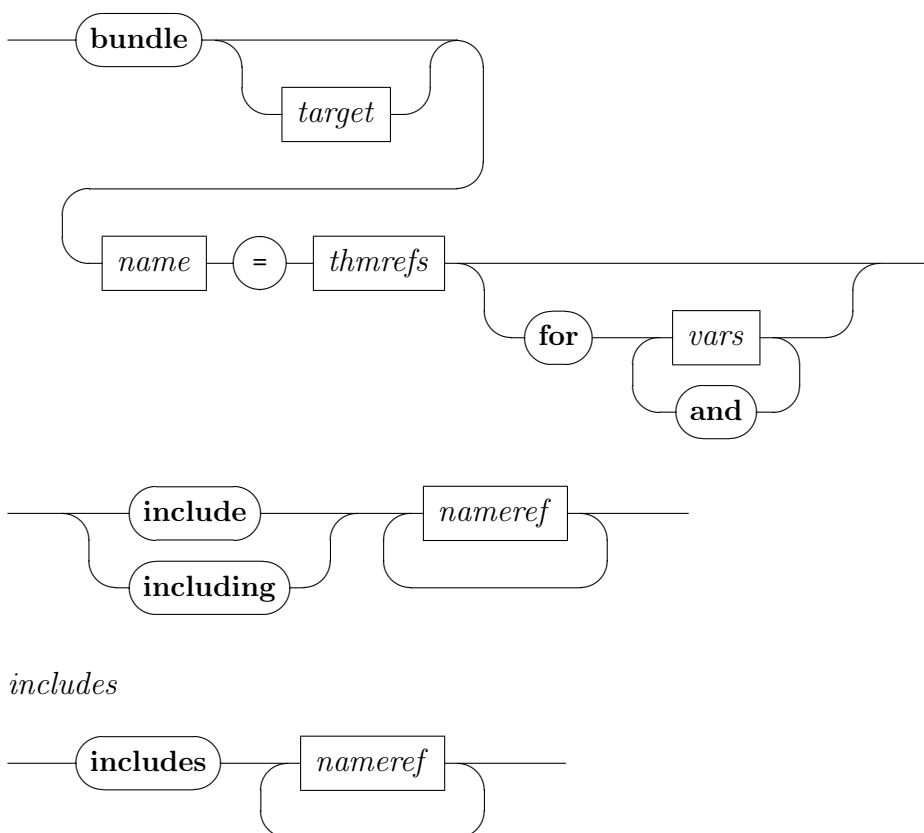
The Isabelle/HOL library contains numerous applications of locales and classes, e.g. see `~~/src/HOL/Algebra`. An example for an unnamed auxiliary contexts is given in `~~/src/HOL/Isar_Examples/Group_Context.thy`.

5.3 Bundled declarations

bundle : $local_theory \rightarrow local_theory$
print_bundles* : $context \rightarrow$
include : $proof(state) \rightarrow proof(state)$
including : $proof(prove) \rightarrow proof(prove)$
includes : $syntax$

The outer syntax of fact expressions (§3.3.7) involves theorems and attributes, which are evaluated in the context and applied to it. Attributes may declare theorems to the context, as in *this_rule [intro] that_rule [elim]* for example. Configuration options (§9.1) are special declaration attributes that operate on the context without a theorem, as in *[[show_types = false]]* for example.

Expressions of this form may be defined as *bundled declarations* in the context, and included in other situations later on. Including declaration bundles augments a local context casually without logical dependencies, which is in contrast to locales and locale interpretation (§5.7).



bundle $b = decls$ defines a bundle of declarations in the current context.

The RHS is similar to the one of the **declare** command. Bundles defined in local theory targets are subject to transformations via morphisms, when moved into different application contexts; this works analogously to any other local theory specification.

print_bundles prints the named bundles that are available in the current context.

include $b_1 \dots b_n$ includes the declarations from the given bundles into the current proof body context. This is analogous to **note** (§6.2.3) with the expanded bundles.

including is similar to **include**, but works in proof refinement (backward mode). This is analogous to **using** (§6.2.3) with the expanded bundles.

includes $b_1 \dots b_n$ is similar to **include**, but works in situations where a specification context is constructed, notably for **context** and long statements of **theorem** etc.

Here is an artificial example of bundling various configuration options:

```
bundle trace = [[simp_trace, linarith_trace, metis_trace, smt_trace]]
```

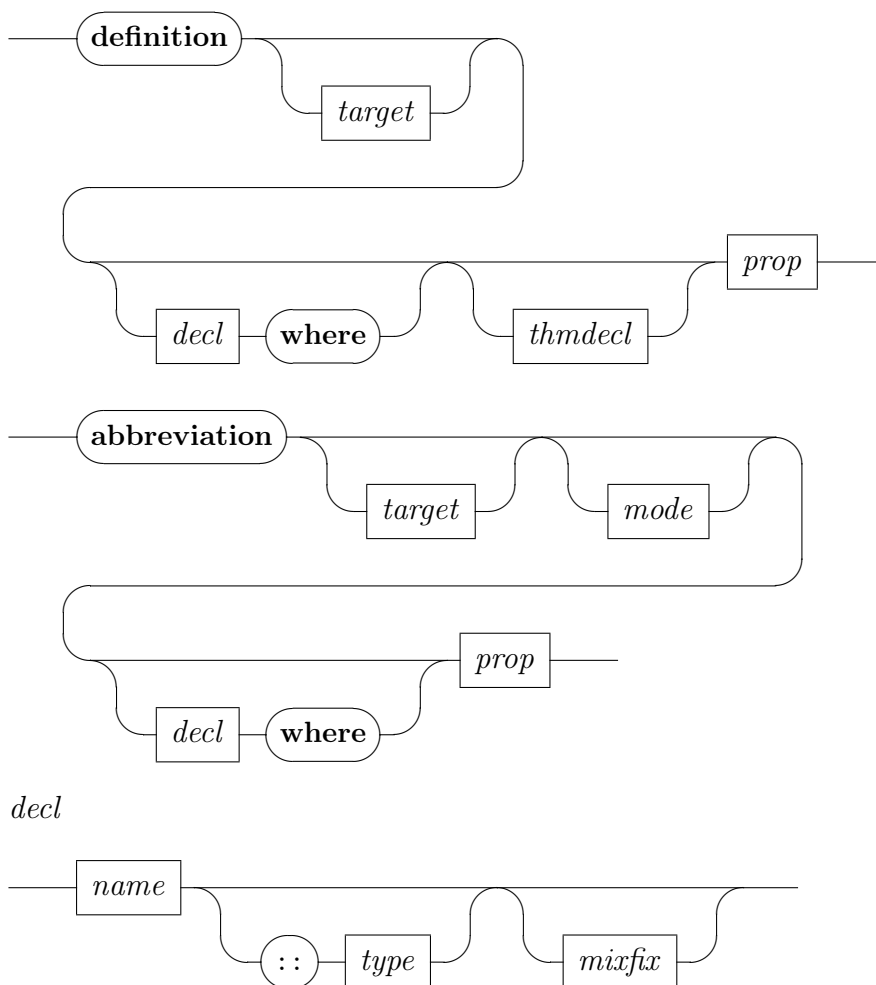
```
lemma x = x
  including trace by metis
```

5.4 Term definitions

```

definition : local_theory → local_theory
      defn   : attribute
print_defn_rules* : context →
      abbreviation : local_theory → local_theory
print_abbrevs* : context →
```

Term definitions may either happen within the logic (as equational axioms of a certain form, see also see §5.11.3), or outside of it as rewrite system on abstract syntax. The second form is called “abbreviation”.



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 $g \equiv \lambda x y. u$.

print_defn_rules prints the definitional rewrite rules declared via *defn* in the current context.

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

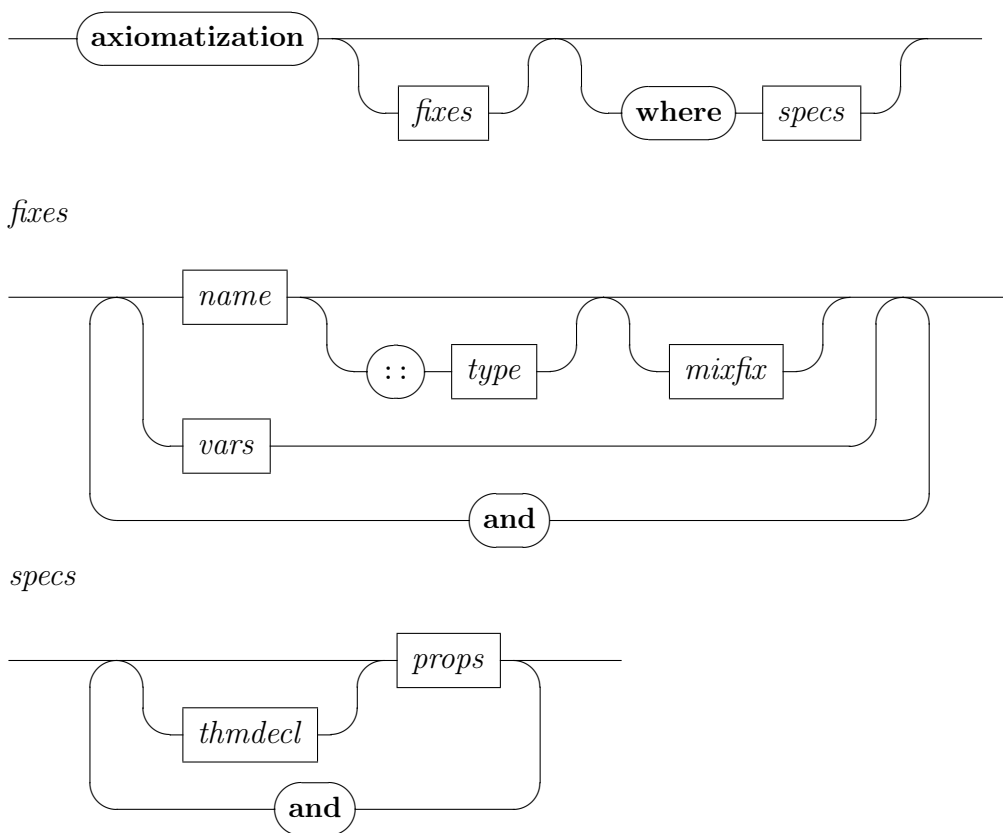
Abbreviations participate in the usual type-inference process, but are expanded before the logic ever sees them. Pretty printing of terms involves higher-order rewriting with rules stemming from reverted abbreviations. This needs some care to avoid overlapping or looping syntactic replacements!

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 §7.5.2.

print_abbrevs prints all constant abbreviations of the current context.

5.5 Axiomatizations

axiomatization : *theory* → *theory* (*axiomatic!*)



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 for the same constants later on, but it is always possible to emit axiomatizations without referring to particular constants. Note that lack of precise dependency tracking of axiomatizations may disrupt the well-formedness of an otherwise definitional theory.

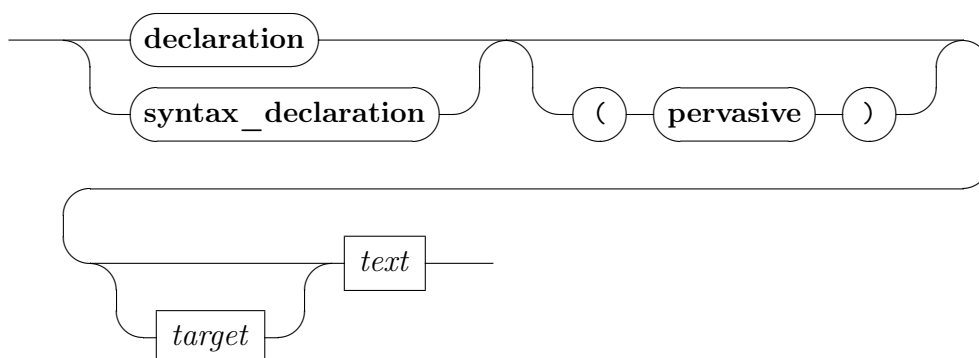
Axiomatization is restricted to a global theory context: support for local theory targets §5.2 would introduce an extra dimension of uncertainty what the written specifications really are, and make it infeasible to argue why they are correct.


Axiomatic specifications are required when declaring a new logical system within Isabelle/Pure, but in an application environment like Isabelle/HOL the user normally stays within definitional mechanisms provided by the logic and its libraries.

5.6 Generic declarations

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

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 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. §5.12), so **declare** only has the effect of applying attributes as included in the theorem specification.

5.7 Locales

A locale is a functor that maps parameters (including implicit type parameters) and a specification to a list of declarations. The syntax of locales is modeled after the Isar proof context commands (cf. §6.2.1).


Locale hierarchies are supported by maintaining a graph of dependencies between locale instances in the global theory. Dependencies may be introduced through `import` (where a locale is defined as sublocale of the imported instances) or by proving that an existing locale is a sublocale of one or several locale instances.

A locale may be opened with the purpose of appending to its list of declarations (cf. §5.2). When opening a locale declarations from all dependencies are collected and are presented as a local theory. In this process, which is called *roundup*, redundant locale instances are omitted. A locale instance is redundant if it is subsumed by an instance encountered earlier. A more detailed description of this process is available elsewhere [4].


5.7.1 Locale expressions

A *locale expression* denotes a context composed of instances of existing locales. The context consists of the declaration elements from the locale instances. Redundant locale instances are omitted according to roundup.


locale_expr




instance




qualifier



pos_insts



named_insts



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. This is useful for shadowing names bound in outer contexts, and for declaring syntax. 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”. Qualifiers play no role in determining whether one locale instance subsumes another.


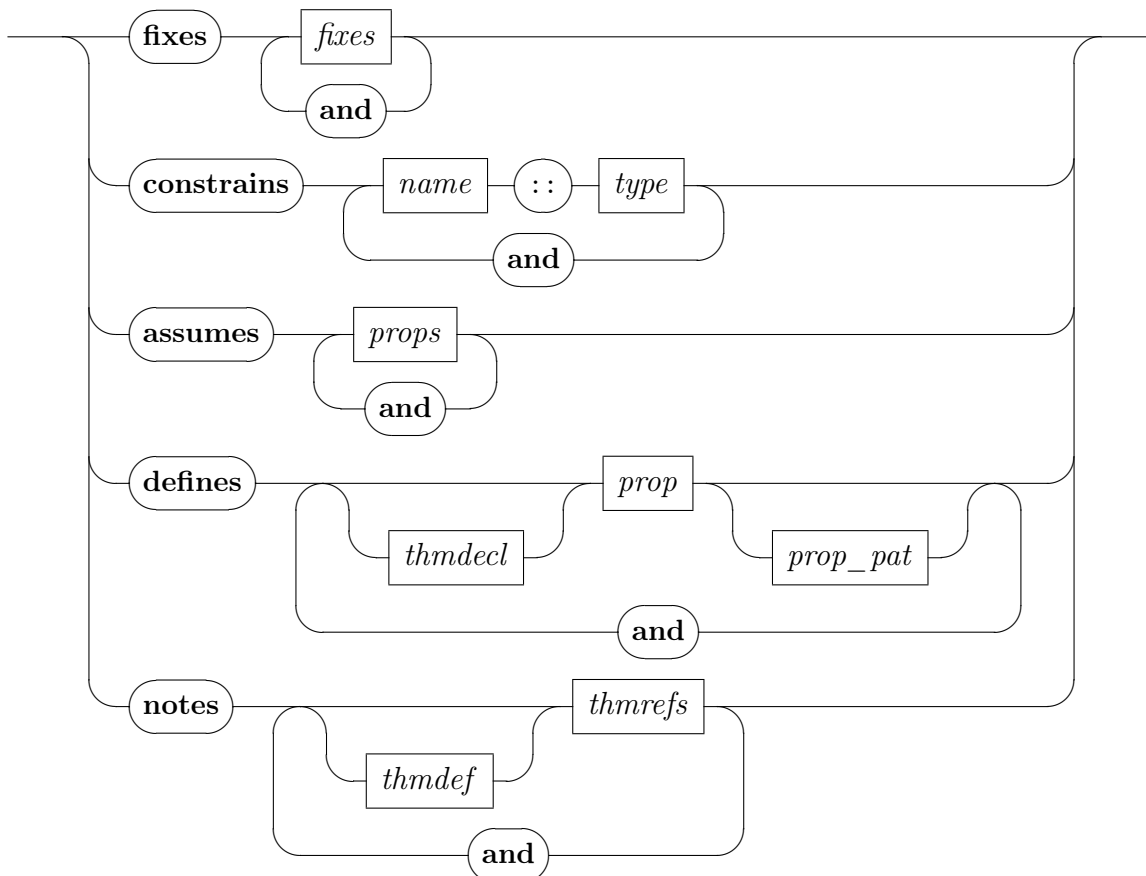
5.7.2 Locale declarations

```

locale      : theory → local_theory
print_locale* : context →
print_locales* : context →
locale_deps* : context →
  intro_locales : method
  unfold_locales : method

```



locale*context_elem*

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 locale expression; see §6.2.1 above. Its **for** clause defines the parameters of *import*. These are parameters of the defined locale. Locale parameters whose instantiation 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.

Both **assumes** and **defines** elements contribute to the locale specification. When defining an operation derived from the parameters, **definition** (§5.4) is usually more appropriate.

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 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.

locale_deps visualizes all locales and their relations as a Hasse diagram. This includes locales defined as type classes (§5.8). See also **print_dependencies** below.

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.

5.7.3 Locale interpretation

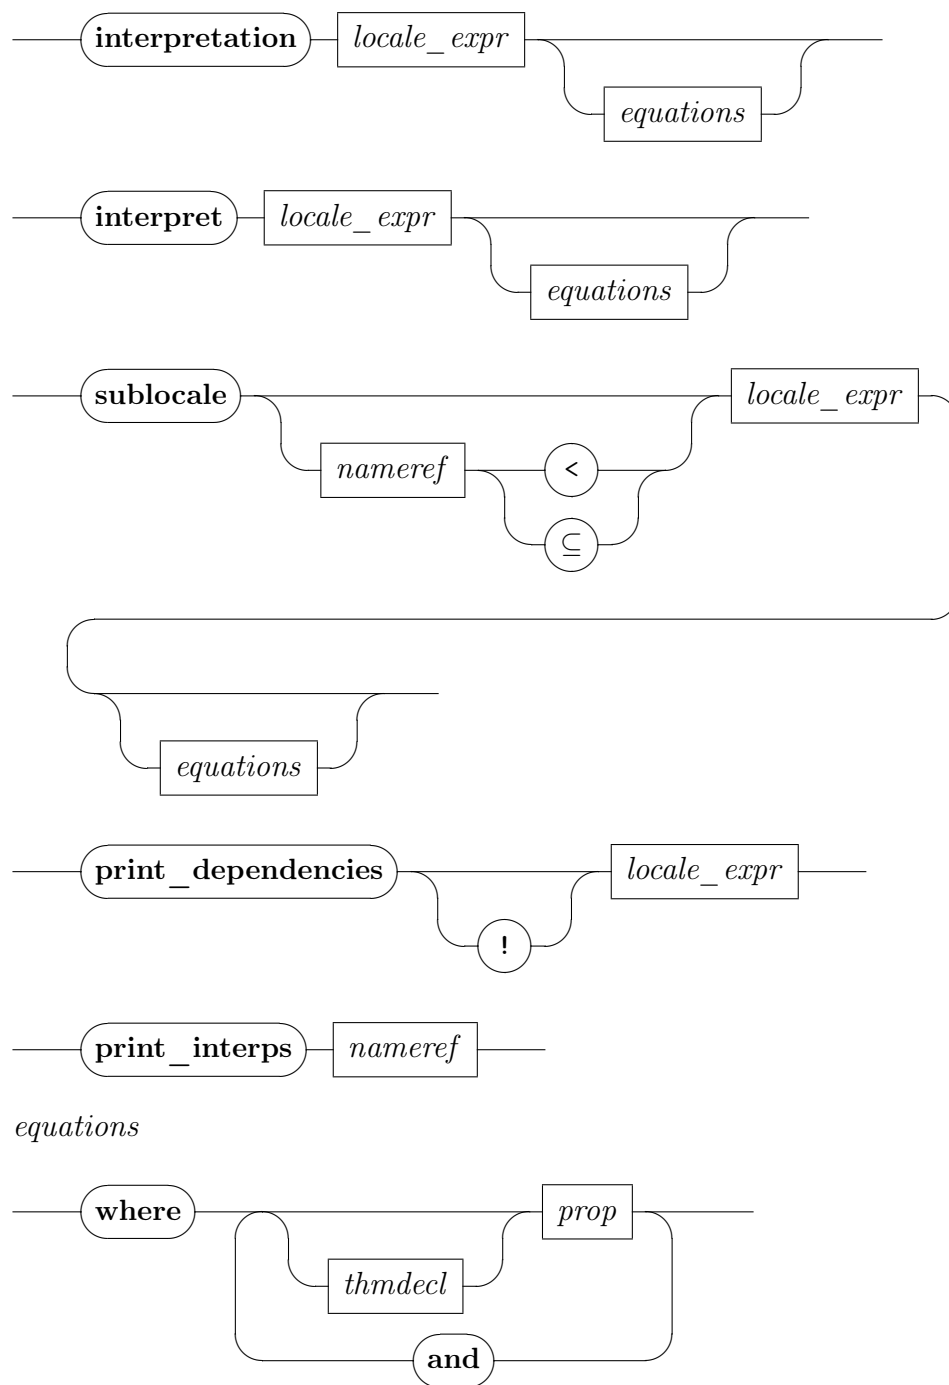
```

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

```

Locales may be instantiated, and the resulting instantiated declarations added to the current context. This requires proof (of the instantiated spec-

ification) and is called *locale interpretation*. Interpretation is possible in locales (**sublocale**), global and local theories (**interpretation**) and also within proof bodies (**interpret**).



interpret *expr* **where** *eqns* interprets *expr* in a global or local theory.

The command generates proof obligations for the instantiated specifications. Once these are discharged by the user, instantiated declarations (in particular, facts) are added to the theory in a post-processing phase.

The command is aware of interpretations that are already active. Post-processing is achieved through a variant of roundup that takes the interpretations of the current global or local theory into account. In order to simplify the proof obligations according to existing interpretations use methods *intro_locales* or *unfold_locales*.

When adding declarations to locales, interpreted versions of these declarations are added to the global theory for all interpretations in the global theory as well. That is, interpretations in global theories dynamically participate in any declarations added to locales.

In contrast, the lifetime of an interpretation in a local theory is limited to the current context block. At the closing **end** of the block the interpretation and its declarations disappear. This enables establishing facts based on interpretations without creating permanent links to the interpreted locale instances, as would be the case with **sublocale**. While **interpret** (**in** *c*) ... is technically possible, it is not useful since its result is discarded immediately.

Free variables in the interpreted expression are allowed. They are turned into schematic variables in the generated declarations. In order to use a free variable whose name is already bound in the context — for example, because a constant of that name exists — add it to the **for** clause.

The equations *eqns* yield *rewrite morphisms*, which are unfolded during post-processing and are useful for interpreting concepts introduced through definitions. The equations must be proved.

interpret *expr* **where** *eqns* interprets *expr* in the proof context and is otherwise similar to interpretation in local theories. Note that for **interpret** the *eqns* 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 locale hierarchy is changed as if *name* imported *expr* (hence the name **sublocale**). When the context of *name* is subsequently entered, traversing the locale hierarchy will involve the locale

instances of $expr$, and their declarations will be added to the context. This makes **sublocale** dynamic: extensions of a locale that is instantiated in $expr$ may take place after the **sublocale** declaration and still become available in the context. Circular **sublocale** declarations are allowed as long as they do not lead to infinite chains.

If interpretations of $name$ exist in the current global 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.

The equations $eqns$ amend the morphism through which $expr$ is interpreted. This enables mapping definitions from the interpreted locales to entities of $name$ and can help break infinite chains induced by circular **sublocale** declarations.

In a named context block the **sublocale** command may also be used, but the locale argument must be omitted. The command then refers to the locale (or class) target of the context block.

print_dependencies $expr$ is useful for understanding the effect of an interpretation of $expr$ in the current context. It lists all locale instances for which interpretations would be added to the current context. Variant **print_dependencies!** does not generalize parameters and assumes an empty context — that is, it prints all locale instances that would be considered for interpretation. The latter is useful for understanding the dependencies of a locale expression.

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

! If a global theory inherits declarations (body elements) for a locale from one
 • parent and an interpretation of that locale from another parent, the interpretation will not be applied to the 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 local 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.

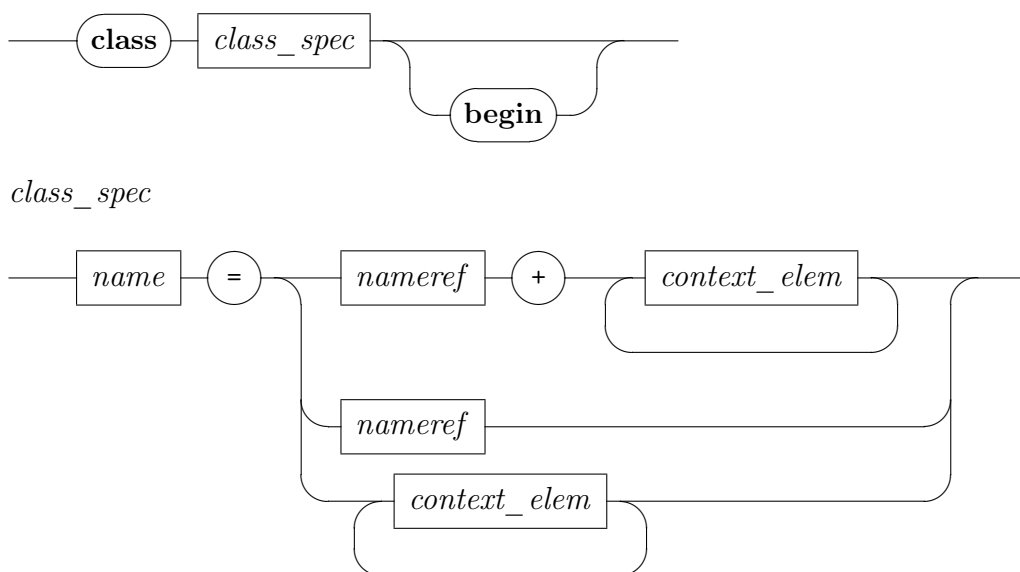
5.8 Classes

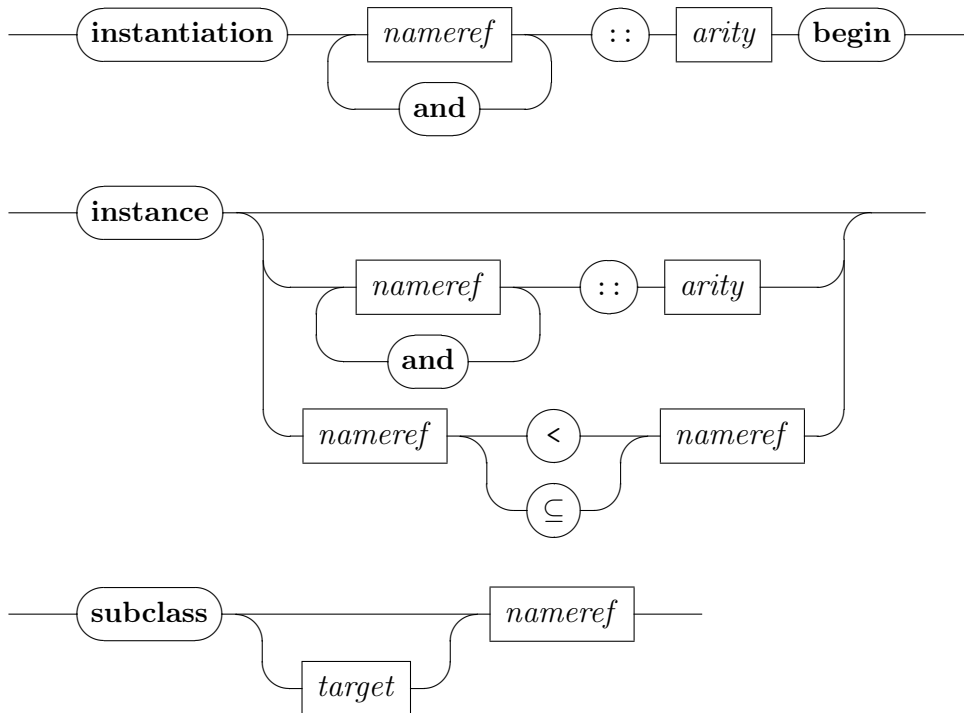
```

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

```

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 [52] 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 [19] for a short tutorial.





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 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 weakened 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.

5.8.1 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.

5.8.2 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 [34].

5.9 Unrestricted overloading


overloading : $theory \rightarrow local_theory$

Isabelle/Pure's definitional schemes support certain forms of overloading (see §5.11.3). Overloading means that a constant being declared as $c :: \alpha decl$ may be defined separately on type instances $c :: (\beta_1, \dots, \beta_n) t decl$ 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 §5.8). Sometimes low-level overloading is desirable. The **overloading** target provides a convenient view for end-users.



spec



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 §5.11.3 for a precise description). It is at the discretion of the user to avoid malformed theory specifications!

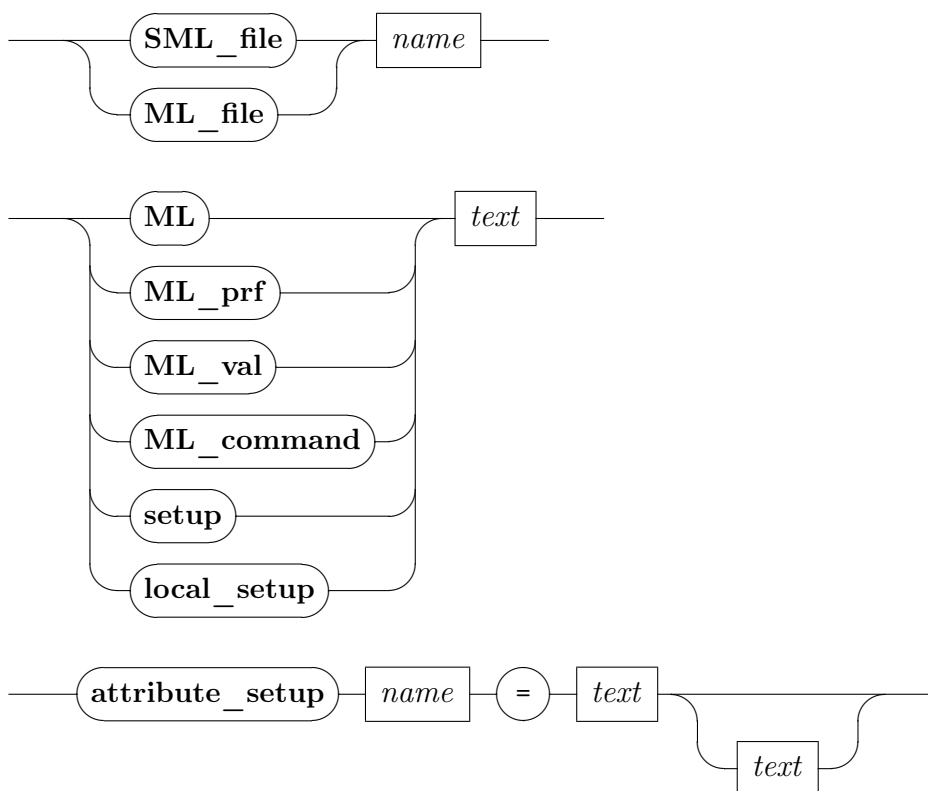
5.10 Incorporating ML code

```

SML_file  : theory → theory
ML_file   : local_theory → local_theory
ML        : local_theory → local_theory
ML_prf    : proof → proof
ML_val    : any →
ML_command : any →
  setup    : theory → theory
  local_setup : local_theory → local_theory
attribute_setup : theory → theory

```

ML_print_depth : *attribute* default 10
 ML_source_trace : *attribute* default *false*
 $ML_exception_trace$: *attribute* default *false*



SML_file *name* reads and evaluates the given Standard ML file. Top-level SML bindings are stored within the theory context; the initial environment is restricted to the Standard ML implementation of Poly/ML, without the many add-ons of Isabelle/ML. Multiple **SML_file** commands may be used to build larger Standard ML projects, independently of the regular Isabelle/ML environment.

ML_file *name* reads and evaluates the given ML file. The current theory context is passed down to the ML toplevel and may be modified, using `Context.>>` or derived ML commands. Top-level ML bindings are stored within the (global or local) theory context.

ML *text* is similar to **ML_file**, but evaluates directly 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)) *}
```

```
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)) *}
```

ML_print_depth controls the printing depth of the ML toplevel pretty printer; the precise effect depends on the ML compiler and run-time

system. Typically the limit should be less than 10. Bigger values such as 100–1000 are occasionally useful for debugging.

ML_source_trace indicates whether the source text that is given to the ML compiler should be output: it shows the raw Standard ML after expansion of Isabelle/ML antiquotations.


ML_exception_trace indicates whether the ML run-time system should print a detailed stack trace on exceptions. The result is dependent on the particular ML compiler version. Note that after Poly/ML 5.3 some optimizations in the run-time systems may hinder exception traces.

The boundary for the exception trace is the current Isar command transactions. It is occasionally better to insert the combinator `Runtime.exn_trace` into ML code for debugging [50], closer to the point where it actually happens.

5.11 Primitive specification elements

5.11.1 Sorts

default_sort : *local_theory* → *local_theory*



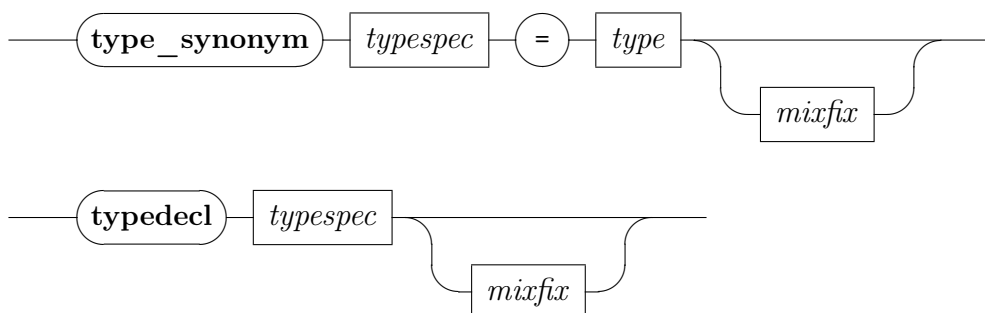
default_sort *s* makes sort *s* the new default sort for any type variable that is given explicitly in the text, but lacks a sort constraint (wrt. the current context). Type variables generated by type inference are not affected.

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.

5.11.2 Types

type_synonym : $local_theory \rightarrow local_theory$
typedecl : $local_theory \rightarrow local_theory$



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 the semantic type definitions in Isabelle/HOL, type synonyms are merely syntactic abbreviations without any logical significance. Internally, type synonyms are fully expanded.

typedecl $(\alpha_1, \dots, \alpha_n) t$ 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 type-class instance $t :: (s, \dots, s)s$.

! If you introduce a new type axiomatically, i.e. via **typedecl** and **axiomatization** (§5.5), the minimum requirement is that it has a non-empty model, to avoid immediate collapse of the logical environment. Moreover, one needs to demonstrate that the interpretation of such free-form axiomatizations can coexist with other axiomatization schemes for types, notably **typedef** in Isabelle/HOL (§11.5), or any other extension that people might have introduced elsewhere.

5.11.3 Constants and definitions

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

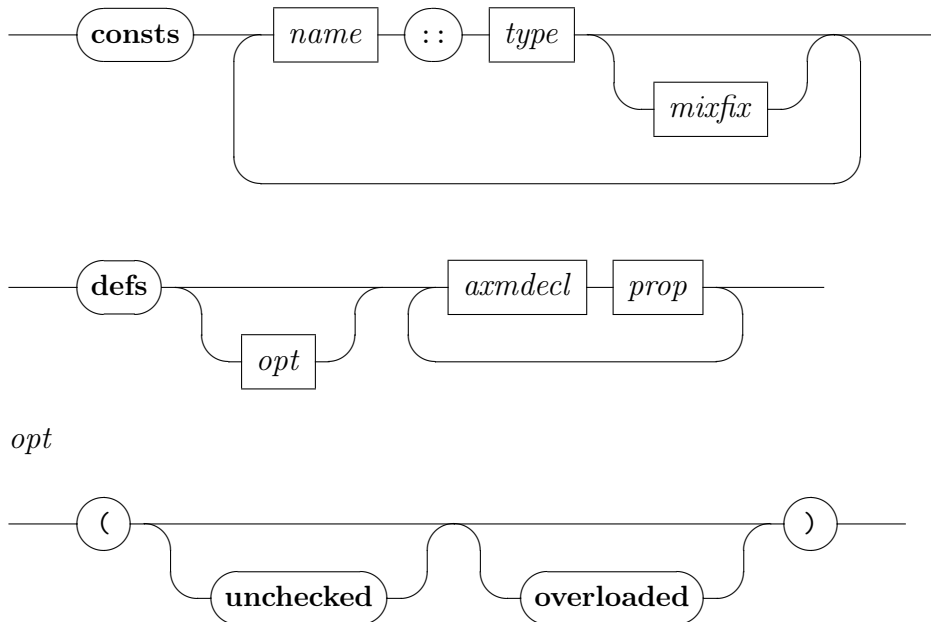
Definitions essentially express abbreviations within the logic. The simplest form of a definition is $c :: \sigma \equiv t$, where c is a newly declared constant. Isabelle also allows derived forms where the arguments of c appear on the

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 $0 :: nat \equiv length ([] :: \alpha list)$ 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 $c :: \sigma$ 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.

5.12 Naming existing theorems

lemmas : $local_theory \rightarrow local_theory$
theorems : $local_theory \rightarrow local_theory$



lemmas $a = b_1 \dots b_n$ **for** $x_1 \dots x_m$ evaluates given facts (with attributes) in the current context, which may be augmented by local variables. Results are standardized before being stored, i.e. schematic variables are renamed to enforce index 0 uniformly.

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


5.13 Oracles

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

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 $name = text$ 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.

5.14 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)*



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

```

next : proof(state) → proof(state)
  { : proof(state) → proof(state)
  } : proof(state) → proof(state)

```

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

oops : $proof \rightarrow local_theory \mid theory$

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 in any way.

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 L^AT_EX macros can be easily adapted to print something like “...” instead of the keyword “**oops**”.

6.2 Statements

6.2.1 Context elements


fix : $proof(state) \rightarrow proof(state)$
assume : $proof(state) \rightarrow proof(state)$
presume : $proof(state) \rightarrow proof(state)$
def : $proof(state) \rightarrow proof(state)$

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 $\vdash \varphi[t]$.



def



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 *term_pat* or *prop_pat* (see §3.3.6).

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

```



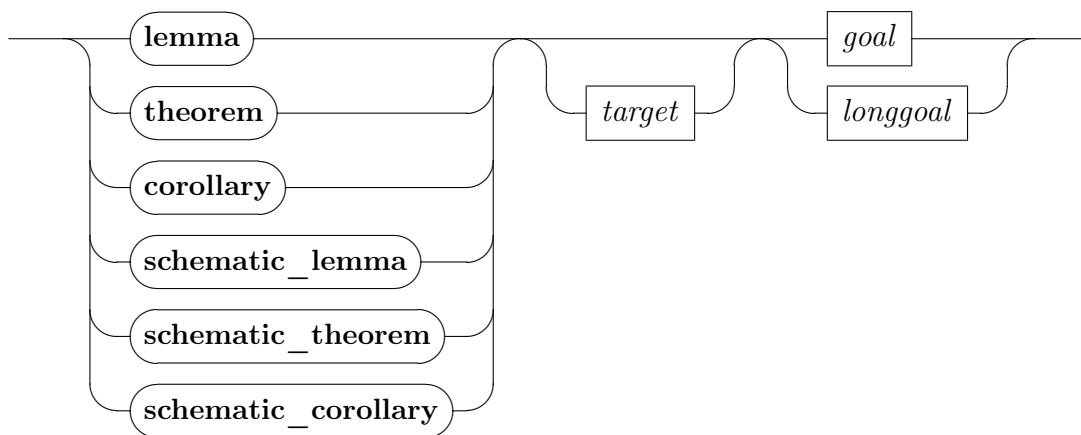
lemma : local_theory → proof(prove)
theorem : local_theory → proof(prove)
corollary : local_theory → proof(prove)
schematic_lemma : local_theory → proof(prove)
schematic_theorem : local_theory → proof(prove)
schematic_corollary : local_theory → proof(prove)
have : proof(state) | proof(chain) → proof(prove)
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 §5.7); 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 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 also *includes* in §5.3 and *context_elem* in §5.7.

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.

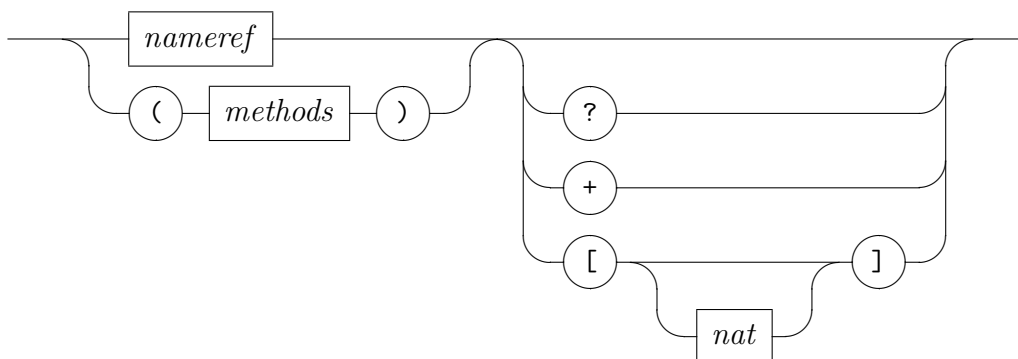

Any goal statement causes some term abbreviations (such as *?thesis*) to be bound automatically, see also §6.2.2.

The optional case names of **obtains** have a twofold meaning: (1) in the proof of this claim they refer to the local context introductions, (2) in the resulting rule they become annotations for symbolic case splits, e.g. for the *cases* method (§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*.

goal_spec



6.3.2 Initial and terminal proof steps

proof : $proof(prove) \rightarrow proof(state)$
qed : $proof(state) \rightarrow proof(state) \mid local_theory \mid theory$
by : $proof(prove) \rightarrow proof(state) \mid local_theory \mid theory$
.. : $proof(prove) \rightarrow proof(state) \mid local_theory \mid theory$
. : $proof(prove) \rightarrow proof(state) \mid local_theory \mid theory$
sorry : $proof(prove) \rightarrow proof(state) \mid local_theory \mid theory$

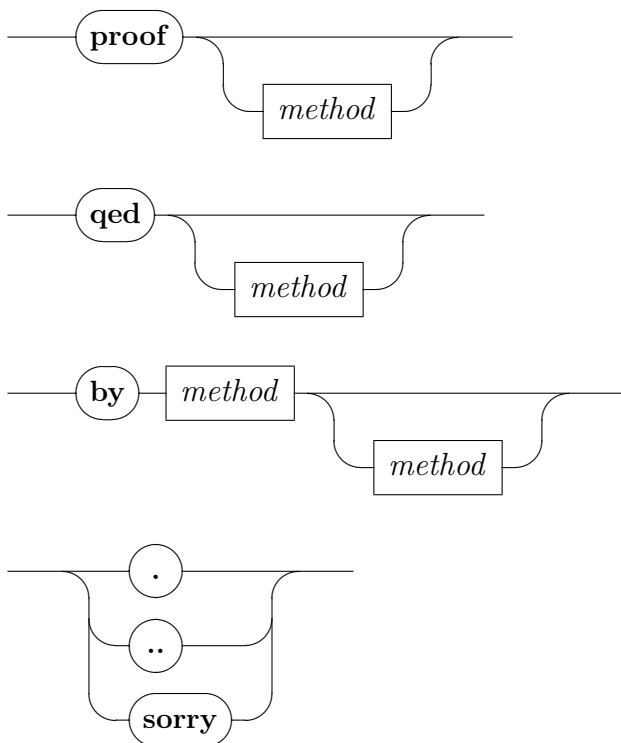
Arbitrary goal refinement via tactics is considered harmful. Structured proof composition in Isar admits proof methods to be invoked in two places only.

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 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* (or its classical variant *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**.

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

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* is enabled. Facts emerging from fake proofs are not the real thing. Internally, the derivation object is tainted by an oracle invocation, which may be inspected via the theorem status [50].

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

```







print_rules* : context →
    - : method
    fact : method
    assumption : method
    this : method
    rule : method
    intro : attribute
    elim : attribute
    dest : attribute
    rule : attribute
    OF : attribute
    of : attribute
    where : attribute

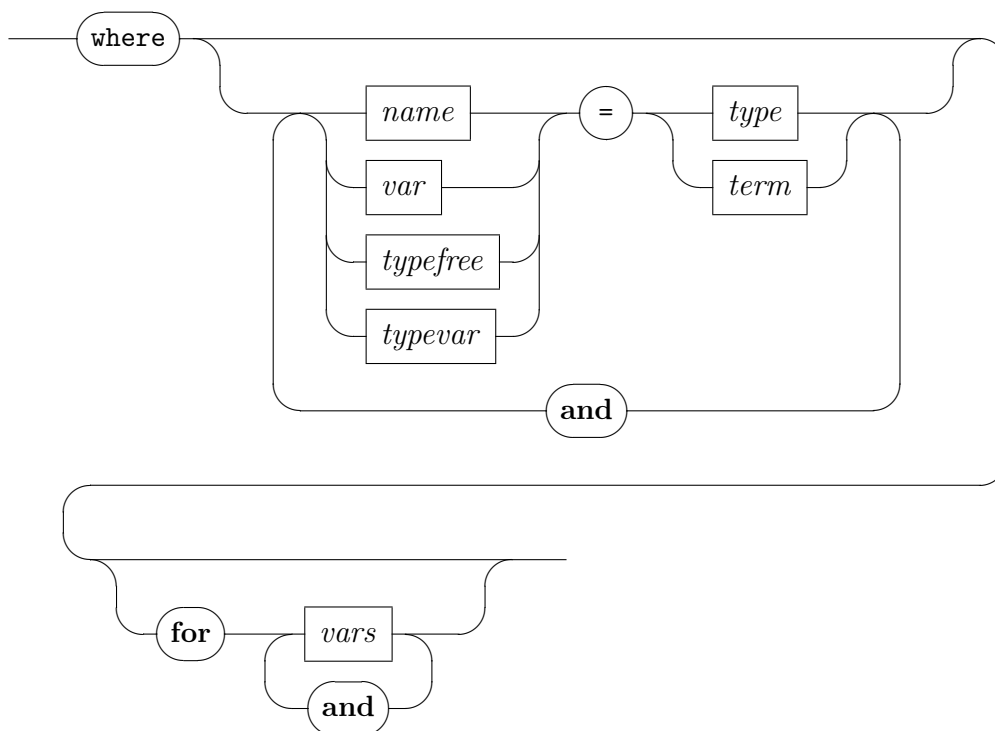
```





rulemod





print_rules prints rules declared via attributes *intro*, *elim*, *dest* of Isabelle/Pure.

See also the analogous **print_claset** command for similar rule declarations of the classical reasoner (§9.4).

“*–*” (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. §3.3.7): “**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 canonical right-to-left order, which means that premises stemming from the a_i emerge in parallel in the result, without interfering with each other. In many practical situations, the a_i do not have premises themselves, so *rule* [*OF* $a_1 \dots a_n$] can be actually read as functional application (modulo unification).

Argument positions may be effectively skipped by using “_” (underscore), which refers to the propositional identity rule in the Pure theory.

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.

An optional context of local variables **for** $x_1 \dots x_m$ may be specified: the instantiated theorem is exported, and these variables become schematic (usually with some shifting of indices).


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.

An optional context of local variables **for** $x_1 \dots x_m$ may be specified as for *of* above.

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* : $proof(prove) \rightarrow proof(prove)$
apply_end* : $proof(state) \rightarrow proof(state)$
done* : $proof(prove) \rightarrow proof(state) \mid local_theory \mid theory$
defer* : $proof \rightarrow proof$
prefer* : $proof \rightarrow proof$
back* : $proof \rightarrow proof$



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. Any proof command may return multiple results, and this command explores the possibilities step-by-step. It is mainly useful for experimentation and interactive exploration, and should be avoided in finished proofs.

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 -> 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 : $proof(state) \mid proof(chain) \rightarrow proof(prove)$
guess* : $proof(state) \mid proof(chain) \rightarrow proof(prove)$

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 see §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).

```

<using  $b_1 \dots b_k$ > obtain  $x_1 \dots x_m$  where  $a: \varphi_1 \dots \varphi_n$  <proof>  $\equiv$ 
  have  $\bigwedge thesis. (\bigwedge x_1 \dots x_m. \varphi_1 \implies \dots \varphi_n \implies thesis) \implies thesis$ 
  proof succeed
  fix thesis
  assume that [Pure.intro?]:  $\bigwedge x_1 \dots x_m. \varphi_1 \implies \dots \varphi_n \implies thesis$ 
  then show thesis
  apply –
  using  $b_1 \dots b_k$  <proof>
qed
fix  $x_1 \dots x_m$  assume*  $a: \varphi_1 \dots \varphi_n$ 

```

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) → proof(state)
finally  : proof(state) → proof(chain)
moreover : proof(state) → proof(state)
ultimately : proof(state) → proof(chain)
print_trans_rules* : context →
    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 “...” 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

also₀ ≡ **note** *calculation* = *this*
also_{n+1} ≡ **note** *calculation* = *trans* [*OF calculation this*]
finally ≡ **also from** *calculation*
moreover ≡ **note** *calculation* = *calculation this*
ultimately ≡ **moreover from** *calculation*



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.

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 patterns) of the current context.

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

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.

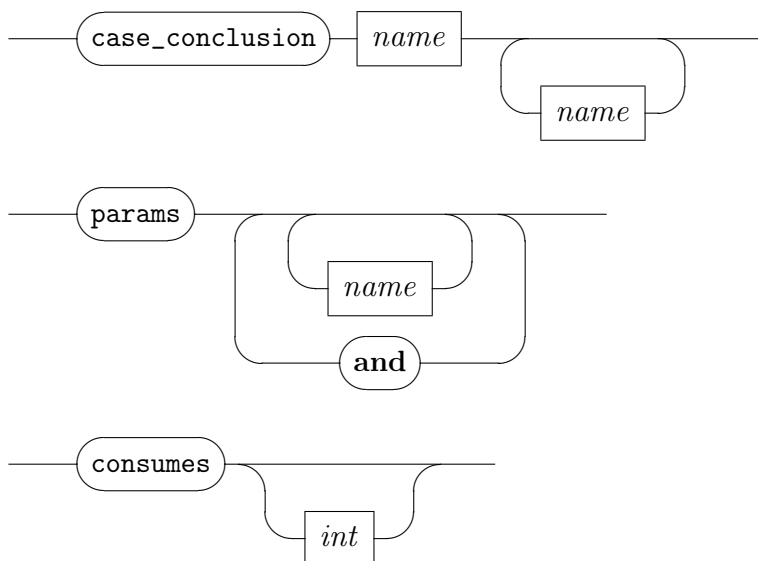
Under normal circumstances, the text of cases emerge from standard elimination or induction rules, which in turn are derived from previous theory specifications in a canonical way (say from **inductive** definitions).

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 *prefix* of the list of premises. Each of the cases c_i can be of the form $c[h_1 \dots h_n]$ where the $h_1 \dots h_n$ are the names of the hypotheses in case c_i from left to right.

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. §11.1). Rules without any *consumes* declaration given are treated as if *consumes 0* had been specified.

A negative n is interpreted relatively to the total number of premises of the rule in the target context. Thus its absolute value specifies the remaining number of premises, after subtracting the prefix of major premises as indicated above. This form of declaration has the technical advantage of being stable under more morphisms, notably those that export the result from a nested **context** with additional assumptions.

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.

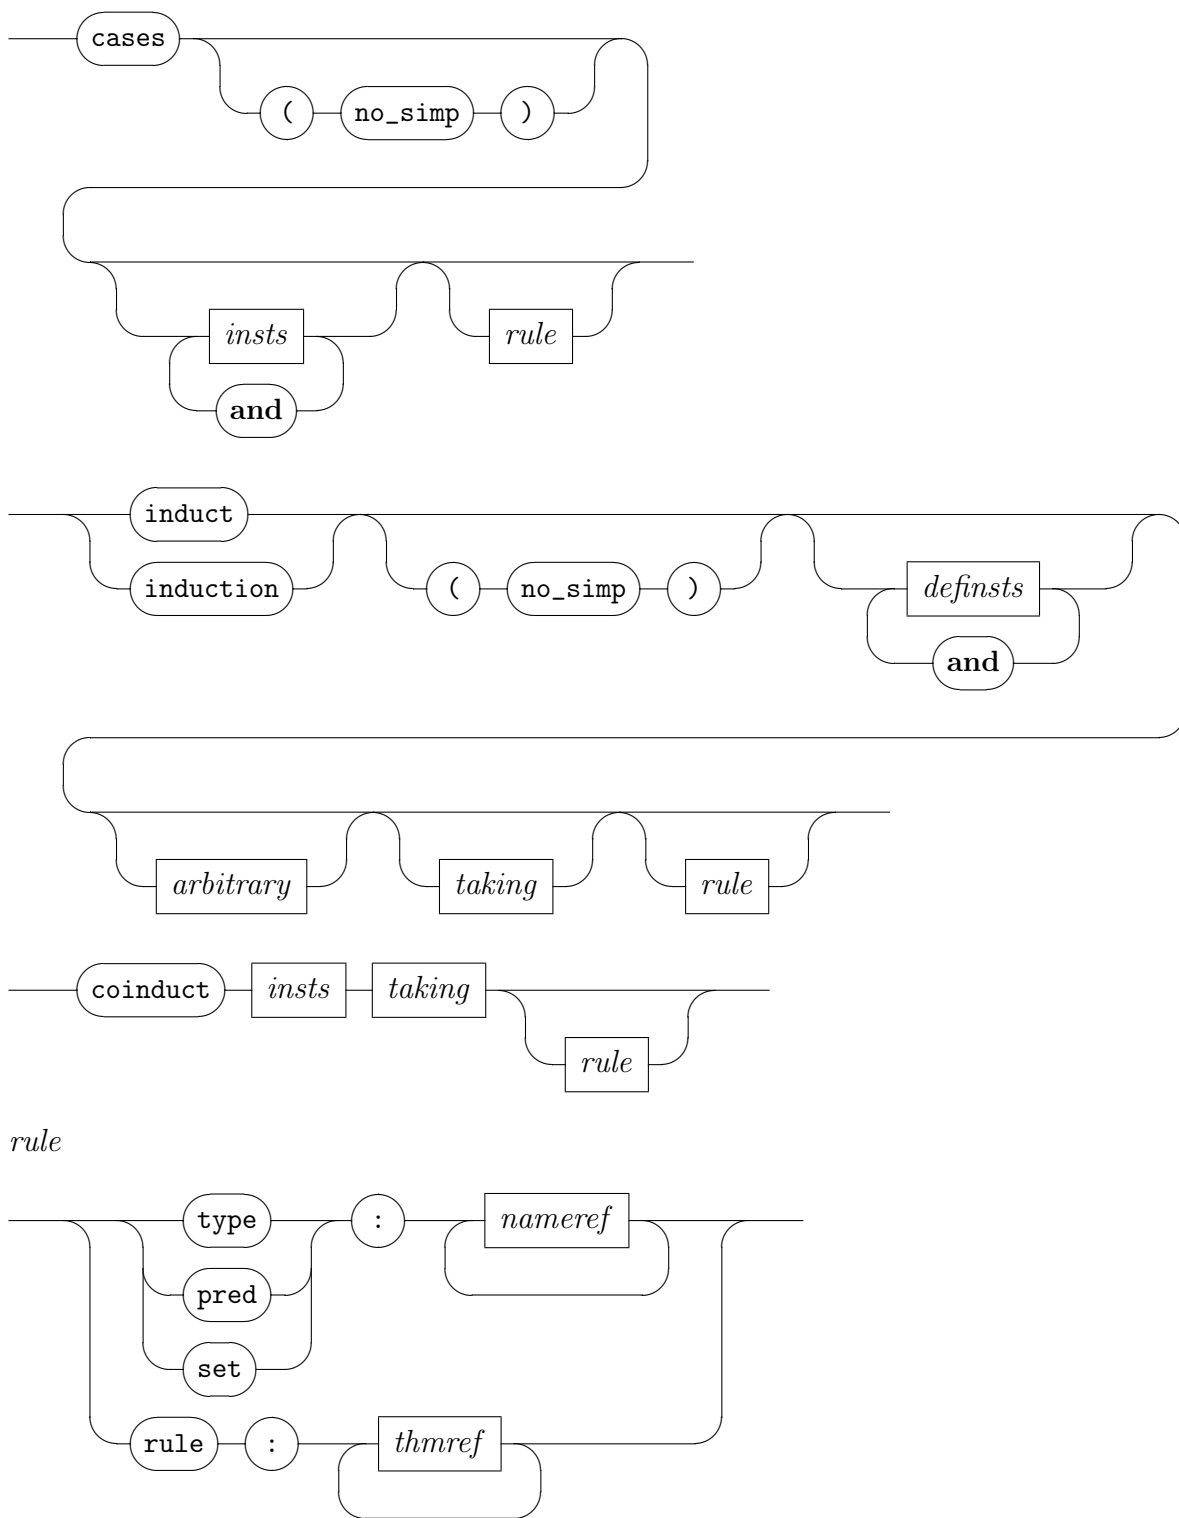
6.6.2 Proof methods

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


The *cases*, *induct*, *induction*, 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.


Method *induction* differs from *induct* only in the names of the facts in the local context invoked by the **case** command.




definst



definsts



arbitrary



taking



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 t</i>	datatype exhaustion (type of <i>t</i>)
$\vdash A$	<i>cases ...</i>	inductive predicate/set elimination (of <i>A</i>)
...	<i>cases ... rule: R</i>	explicit rule <i>R</i>

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 and *induction insts R* are analogous to the *cases* method, but refer 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. One can separate variables by “*and*” to generalize them in other goals than the first. 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*: $t_1 \dots t_n$ ” 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> ...		predicate/set coinduction (of A)
... <i>coinduct</i> ... <i>rule</i> : 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*: $t_1 \dots t_n$ ” 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 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 *induct* 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.

In *induction* proofs, local assumptions introduced by cases are split into three different kinds: *IH*, the induction hypotheses, *hyps*, the remaining hypotheses stemming from the rule, and *prems*, the assumptions from the goal statement. The names are *c.IH*, *c.hyps* and *c.prems*, as above.


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

The inner syntax of Isabelle provides concrete notation for the main entities of the logical framework, notably λ -terms with types and type classes. Applications may either extend existing syntactic categories by additional notation, or define new sub-languages that are linked to the standard term language via some explicit markers. For example `F00 foo` could embed the syntax corresponding for some user-defined nonterminal *foo* — within the bounds of the given lexical syntax of Isabelle/Pure.

The most basic way to specify concrete syntax for logical entities works via mixfix annotations (§7.2), which may be usually given as part of the original declaration or via explicit notation commands later on (§7.3). This already covers many needs of concrete syntax without having to understand the full complexity of inner syntax layers.

Further details of the syntax engine involves the classical distinction of lexical language versus context-free grammar (see §7.4), and various mechanisms for *syntax transformations* (see §7.5).

7.1 Printing logical entities

7.1.1 Diagnostic commands

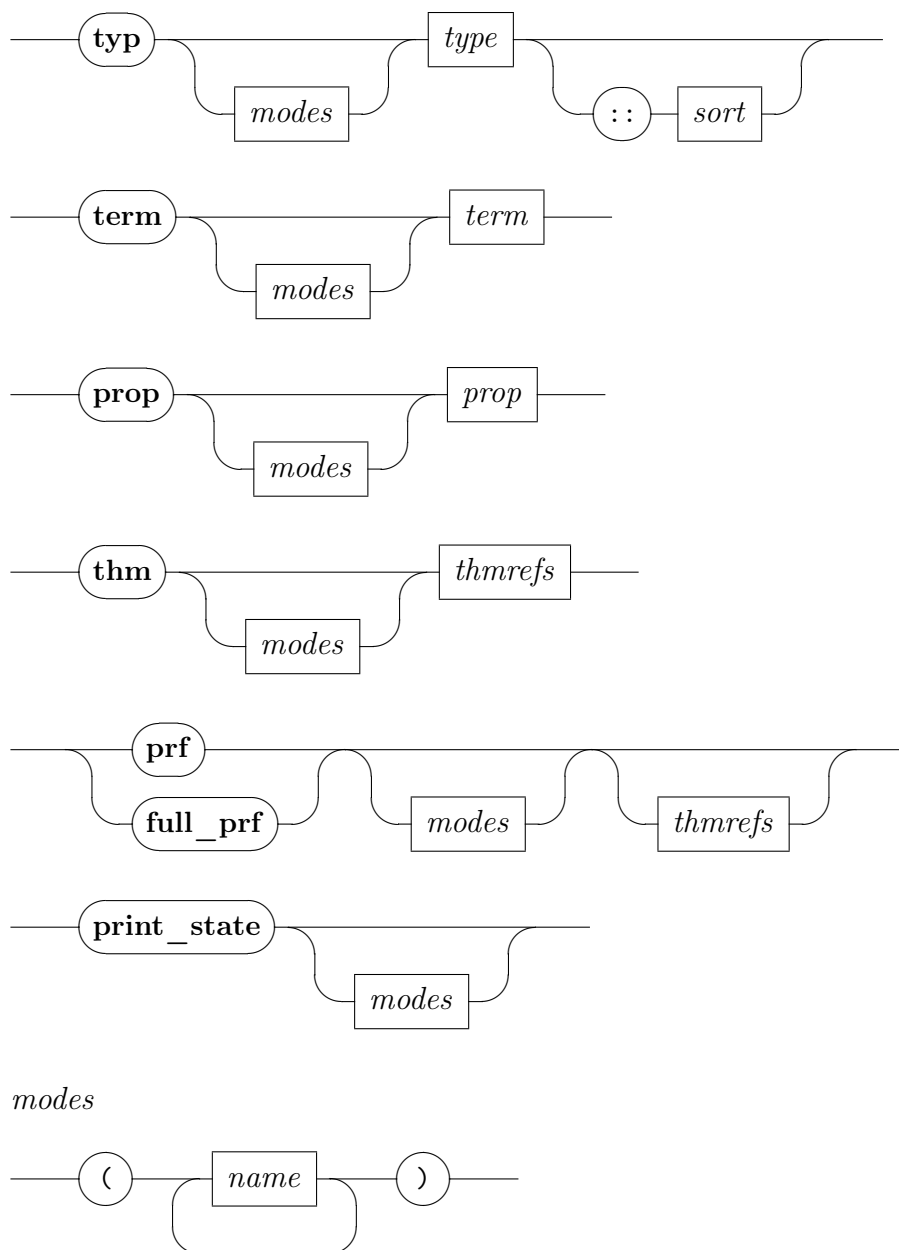
```

typ*   : context →
term*  : context →
prop*  : context →
thm*   : context →
prf*   : context →
full_prf* : context →
print_state* : any →

```

These diagnostic commands assist interactive development by printing inter-

nal logical entities in a human-readable fashion.



typ τ reads and prints a type expression according to the current context.

typ $\tau :: s$ uses type-inference to determine the most general way to make τ conform to sort s . For concrete τ this checks if the type belongs to that sort. Dummy type parameters “_” (underscore) are assigned to fresh

type variables with most general sorts, according to the principles of type-inference.

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.

print_state prints the current proof state (if present), including current facts and goals.

All of the diagnostic commands above admit a list of *modes* to be specified, which is appended to the current print mode; see also §7.1.3. Thus the output behavior may be modified according to particular print mode features. For example, **print_state** (*latex xsymbols*) prints the current proof state with mathematical symbols and special characters represented in L^AT_EX source, according to the Isabelle style [57].

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

<i>show_markup</i>	:	<i>attribute</i>	
<i>show_types</i>	:	<i>attribute</i>	default <i>false</i>
<i>show_sorts</i>	:	<i>attribute</i>	default <i>false</i>
<i>show_consts</i>	:	<i>attribute</i>	default <i>false</i>
<i>show_abbrevs</i>	:	<i>attribute</i>	default <i>true</i>
<i>show_brackets</i>	:	<i>attribute</i>	default <i>false</i>
<i>names_long</i>	:	<i>attribute</i>	default <i>false</i>
<i>names_short</i>	:	<i>attribute</i>	default <i>false</i>
<i>names_unique</i>	:	<i>attribute</i>	default <i>true</i>
<i>eta_contract</i>	:	<i>attribute</i>	default <i>true</i>
<i>goals_limit</i>	:	<i>attribute</i>	default 10
<i>show_main_goal</i>	:	<i>attribute</i>	default <i>false</i>
<i>show_hyps</i>	:	<i>attribute</i>	default <i>false</i>
<i>show_tags</i>	:	<i>attribute</i>	default <i>false</i>
<i>show_question_marks</i>	:	<i>attribute</i>	default <i>true</i>

These configuration options control the detail of information that is displayed for types, terms, theorems, goals etc. See also §9.1.

show_markup controls direct inlining of markup into the printed representation of formal entities — notably type and sort constraints. This enables Prover IDE users to retrieve that information via tooltips or popups while hovering with the mouse over the output window, for example. Consequently, this option is enabled by default for Isabelle/jEdit, but disabled for TTY and Proof General /Emacs where document markup would not work.

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 enabled, types are always shown as well. In Isabelle/jEdit, manual setting of these options is normally not required thanks to *show_markup* above.

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.

show_brackets controls bracketing in pretty printed output. If enabled, 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.

names_long, *names_short*, and *names_unique* 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.

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 $\lambda h. F (\lambda x. h x)$.

Enabling *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.

goals_limit controls the maximum number of subgoals to be printed.

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 enabling *show_hyps*, 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 Alternative print modes

```
print_mode_value: unit -> string list
Print_Mode.with_modes: string list -> ('a -> 'b) -> 'a -> 'b
```

The *print mode* facility allows to modify various operations for printing. Commands like **typ**, **term**, **thm** (see §7.1.1) take additional print modes as optional argument. The underlying ML operations are as follows.

`print_mode_value ()` yields the list of currently active print mode names.

This should be understood as symbolic representation of certain individual features for printing (with precedence from left to right).

`Print_Mode.with_modes modes f x` evaluates $f x$ in an execution context where the print mode is prepended by the given *modes*. This provides a thread-safe way to augment print modes. It is also monotonic in the set of mode names: it retains the default print mode that certain user-interfaces might have installed for their proper functioning!

! The old global reference `print_mode` should never be used directly in applications. Its main reason for being publicly accessible is to support historic versions of Proof General.

The pretty printer for inner syntax maintains alternative mixfix productions for any print mode name invented by the user, say in commands like **notation** or **abbreviation**. Mode names can be arbitrary, but the following ones have a specific meaning by convention:

- "" (the empty string): default mode; implicitly active as last element in the list of modes.
- **input**: dummy print mode that is never active; may be used to specify notation that is only available for input.

- `internal` dummy print mode that is never active; used internally in Isabelle/Pure.
- `xsymbols`: enable proper mathematical symbols instead of ASCII art.¹
- `HTML`: additional mode that is active in HTML presentation of Isabelle theory sources; allows to provide alternative output notation.
- `latex`: additional mode that is active in L^AT_EX document preparation of Isabelle theory sources; allows to provide alternative output notation.

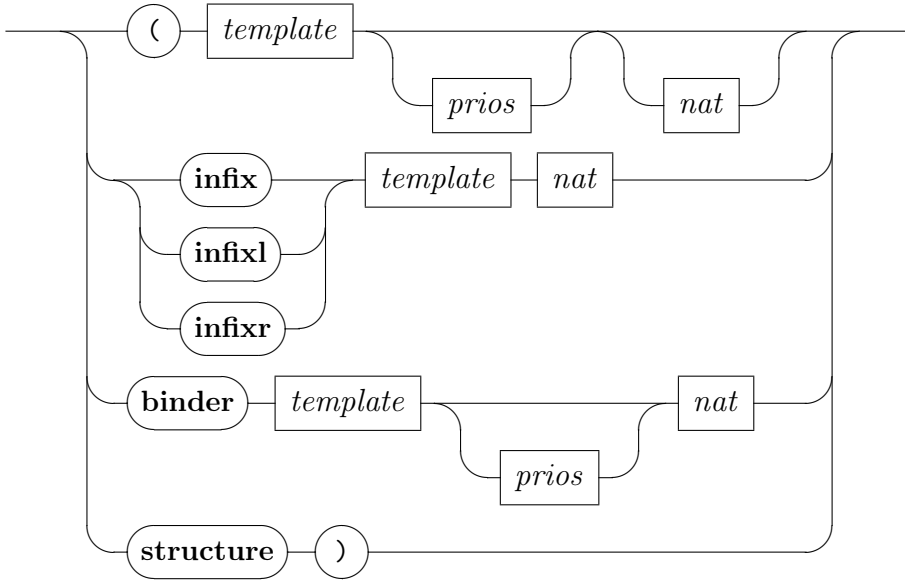
7.2 Mixfix annotations

Mixfix annotations specify concrete *inner syntax* of Isabelle types and terms. Locally fixed parameters in toplevel theorem statements, locale and class specifications also admit mixfix annotations in a fairly uniform manner. A mixfix annotation describes the concrete syntax, the translation to abstract syntax, and the pretty printing. Special case annotations provide a simple means of specifying infix operators and binders.


Isabelle mixfix syntax is inspired by OBJ [15]. It allows to specify any context-free priority grammar, which is more general than the fixity declarations of ML and Prolog.

¹This traditional mode name stems from the “X-Symbol” package for old versions Proof General with XEmacs, although that package has been superseded by Unicode in recent years.


mixfix



template



prios



The string given as *template* 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 §5.7). Only locally fixed variables may be declared as **structure**.

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.

7.2.1 The general mixfix form

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.²

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.

²Omitting priorities is prone to syntactic ambiguities unless the delimiter tokens determine fully bracketed notation, as in *if _ then _ else _ fi*.

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.

The general idea of pretty printing with blocks and breaks is also described in [41]; it goes back to [36].

7.2.2 Infixes

Infix operators are specified by convenient short forms that abbreviate general mixfix annotations as follows:

$$\begin{aligned} (\mathbf{infix} \text{ "sy" } p) &\mapsto \text{"(_ sy/ _)" [p + 1, p + 1] p} \\ (\mathbf{infixl} \text{ "sy" } p) &\mapsto \text{"(_ sy/ _)" [p, p + 1] p} \\ (\mathbf{infixr} \text{ "sy" } p) &\mapsto \text{"(_ sy/ _)" [p + 1, p] p} \end{aligned}$$

The mixfix template "(_ sy/ _)" specifies 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 alternative notation $\mathbf{op} \text{ sy}$ is introduced in addition. Thus any infix operator may be written in prefix form (as in ML), independently of the number of arguments in the term.

7.2.3 Binders

A *binder* is a variable-binding construct such as a quantifier. The idea to formalize $\forall x. b$ as $All (\lambda x. b)$ for $All :: ('a \Rightarrow bool) \Rightarrow bool$ already goes back to [14]. Isabelle declarations of certain higher-order operators may be annotated with **binder** annotations as follows:

$$c :: \text{"}(\tau_1 \Rightarrow \tau_2) \Rightarrow \tau_3\text{" } (\mathbf{binder} \text{ "sy" } [p] q)$$

This introduces concrete binder syntax $sy\ x.\ b$, where x is a bound variable of type τ_1 , the body b has type τ_2 and the whole term has type τ_3 . The optional integer p specifies the syntactic priority of the body; the default is q , which is also the priority of the whole construct.

Internally, the binder syntax is expanded to something like this:

$$c_binder :: "idts \Rightarrow \tau_2 \Rightarrow \tau_3" \quad ("(3sy_./ _)" [0, p] q)$$

Here $idts$ is the nonterminal symbol for a list of identifiers with optional type constraints (see also §7.4.3). The mixfix template $"(3sy_./ _)"$ defines argument positions for the bound identifiers and the body, separated by a dot with optional line break; the entire phrase is a pretty printing block of indentation level 3. Note that there is no extra space after sy , so it needs to be included user specification if the binder syntax ends with a token that may be continued by an identifier token at the start of $idts$.

Furthermore, a syntax translation to transforms $c_binder\ x_1 \dots x_n\ b$ into iterated application $c\ (\lambda x_1. \dots c\ (\lambda x_n. b) \dots)$. This works in both directions, for parsing and printing.

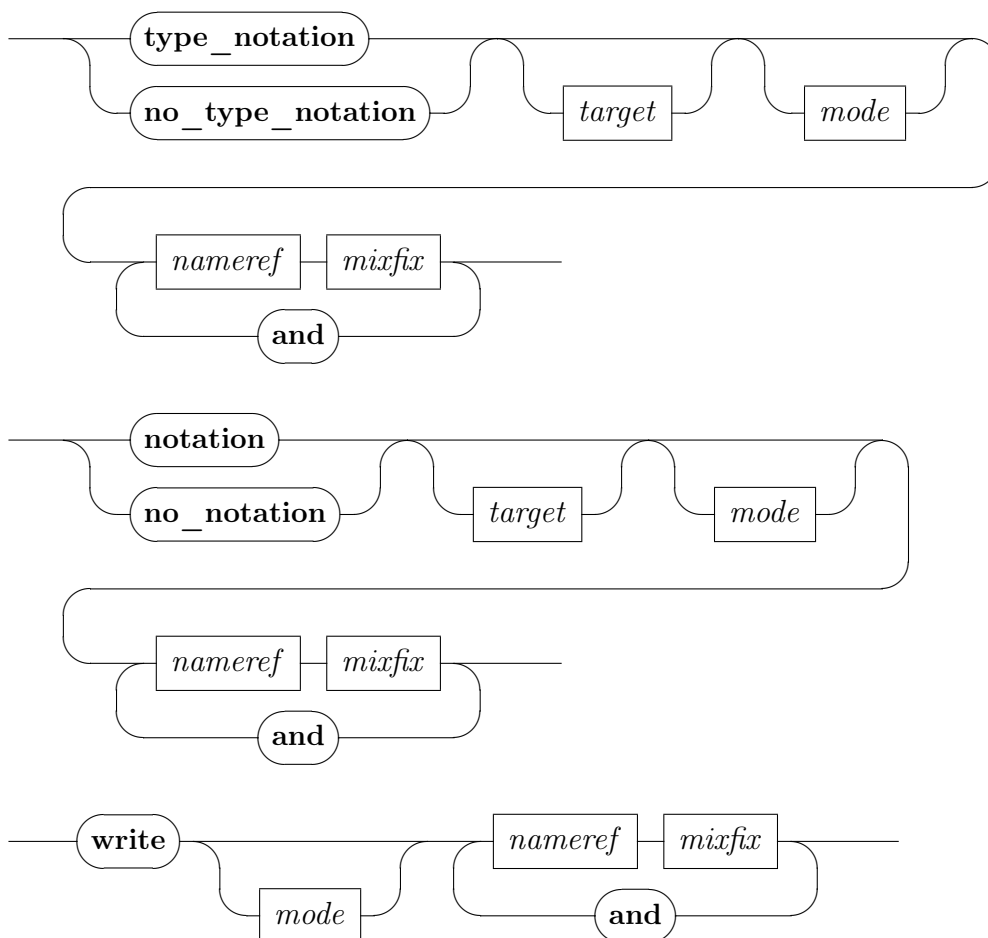
7.3 Explicit notation

```

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)

```

Commands that introduce new logical entities (terms or types) usually allow to provide mixfix annotations on the spot, which is convenient for default notation. Nonetheless, the syntax may be modified later on by declarations for explicit notation. This allows to add or delete mixfix annotations for of existing logical entities within the current context.



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.

7.4 The Pure syntax

7.4.1 Lexical matters

The inner lexical syntax vaguely resembles the outer one (§3.2), 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. §7.4.2);
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 (§3.2).

<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_token</i>	=	<i>nat</i> <i>-nat</i>
<i>float_token</i>	=	<i>nat.nat</i> <i>-nat.nat</i>
<i>xnum_token</i>	=	<i>#nat</i> <i>#-nat</i>
<i>str_token</i>	=	' ' ... ' '
<i>string_token</i>	=	" ... "
<i>cartouche</i>	=	\<open> ... \<close>

The token categories *num_token*, *float_token*, *xnum_token*, *str_token*, *string_token*, and *cartouche* are not used in Pure. Object-logics may implement numerals and string literals by adding appropriate syntax declarations, together with some translation functions (e.g. see `~/src/HOL/Tools/string_syntax.ML`).

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

7.4.2 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:

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

7.4.3 The Pure grammar

The priority grammar of the *Pure* theory is defined approximately like this:

$$\begin{array}{lcl}
 any & = & prop \mid logic \\
 \\
 prop & = & (prop) \\
 & | & prop^{(4)} :: type \quad (3) \\
 & | & any^{(3)} == any^{(3)} \quad (2) \\
 & | & any^{(3)} \equiv any^{(3)} \quad (2) \\
 & | & prop^{(3)} \&\&\& prop^{(2)} \quad (2) \\
 & | & prop^{(2)} ==> prop^{(1)} \quad (1) \\
 & | & prop^{(2)} \implies prop^{(1)} \quad (1) \\
 & | & [\mid prop ; \dots ; prop \mid] ==> prop^{(1)} \quad (1) \\
 & | & [[prop ; \dots ; prop]] \implies prop^{(1)} \quad (1) \\
 & | & !! idts . prop \quad (0) \\
 & | & \wedge idts . prop \quad (0) \\
 & | & OFCLASS (type , logic) \\
 & | & SORT_CONSTRAINT (type) \\
 & | & TERM logic \\
 & | & PROP apropos \\
 \\
 apropos & = & (apropos) \\
 & | & id \mid longid \mid var \mid _ \mid \dots \\
 & | & CONST id \mid CONST longid \\
 & | & XCONST id \mid XCONST longid \\
 & | & logic^{(1000)} any^{(1000)} \dots any^{(1000)} \quad (999) \\
 \\
 logic & = & (logic)
 \end{array}$$

	$logic^{(4)} :: type$	(3)
	$id \mid longid \mid var \mid _ \mid \dots$	
	$CONST id \mid CONST longid$	
	$XCONST id \mid XCONST longid$	
	$logic^{(1000)} any^{(1000)} \dots any^{(1000)}$	(999)
	$\diamond index^{(1000)}$	
	$\% pptrns . any^{(3)}$	(3)
	$\lambda pptrns . any^{(3)}$	(3)
	$op == \mid op \equiv \mid op \&\&\&$	
	$op ==> \mid op \implies$	
	$TYPE (type)$	
	$idt = (idt) \mid id \mid _$	
	$id :: type$	(0)
	$_ :: type$	(0)
	$index = \langle \text{b} \text{sub} \rangle logic^{(0)} \langle \text{e} \text{sub} \rangle \mid \mid 1$	
	$idts = idt \mid idt^{(1)} idts$	(0)
	$pttrn = idt$	
	$pttrns = pttrn \mid pttrn^{(1)} pttrns$	(0)
	$type = (type)$	
	$tid \mid tvar \mid _$	
	$tid :: sort \mid tvar :: sort \mid _ :: sort$	
	$type_name \mid type^{(1000)} type_name$	
	$(type , \dots , type) type_name$	
	$type^{(1)} \Rightarrow type$	(0)
	$type^{(1)} \Rightarrow type$	(0)
	$[type , \dots , type] \Rightarrow type$	(0)
	$[type , \dots , type] \Rightarrow type$	(0)
	$type_name = id \mid longid$	
	$sort = class_name \mid \{ \}$	
	$\{ class_name , \dots , class_name \}$	

$$\textit{class_name} = \textit{id} \mid \textit{longid}$$

Here literal terminals are printed *verbatim*; see also §7.4.1 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*.

index denotes an optional index term for indexed syntax. If omitted, it refers to the first **structure** variable in the context. The special dummy “i” serves as pattern variable in mixfix annotations that introduce indexed notation.

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 .

pttrn and *pttrns* 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 “ $_ :: \text{sort}$ ” 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 §3.3.6) 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).

- `CONST` ensures that the given identifier is treated as constant term, and passed through the parse tree in fully internalized form. This is particularly relevant for translation rules (§7.5.2), notably on the RHS.
- `XCONST` is similar to `CONST`, but retains the constant name as given. This is only relevant to translation rules (§7.5.2), notably on the LHS.

7.4.4 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 §7.4.1.

prods lists the productions of the underlying priority grammar; see §7.4.2.

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 §7.1.3.

parse_rules and *print_rules* relate to syntax translations (macros); see §7.5.2.

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 §7.5.3.

parse_translation and *print_translation* list the sets of constants that invoke regular translation functions; see §7.5.3.

7.4.5 Ambiguity of parsed expressions

```
syntax_ambiguity_warning : attribute default true
syntax_ambiguity_limit  : attribute default 10
```

Depending on the grammar and the given input, parsing may be ambiguous. Isabelle lets the Earley parser enumerate all possible parse trees, and then tries to make the best out of the situation. Terms that cannot be type-checked are filtered out, which often leads to a unique result in the end. Unlike regular type reconstruction, which is applied to the whole collection of input terms simultaneously, the filtering stage only treats each given term in isolation. Filtering is also not attempted for individual types or raw ASTs (as required for **translations**).

Certain warning or error messages are printed, depending on the situation and the given configuration options. Parsing ultimately fails, if multiple results remain after the filtering phase.

syntax_ambiguity_warning controls output of explicit warning messages about syntax ambiguity.

syntax_ambiguity_limit determines the number of resulting parse trees that are shown as part of the printed message in case of an ambiguity.

7.5 Syntax transformations

The inner syntax engine of Isabelle provides separate mechanisms to transform parse trees either via rewrite systems on first-order ASTs (§7.5.2), or ML functions on ASTs or syntactic λ -terms (§7.5.3). This works both for parsing and printing, as outlined in figure 7.1.

These intermediate syntax tree formats eventually lead to a pre-term with all names and binding scopes resolved, but most type information still missing. Explicit type constraints might be given by the user, or implicit position information by the system — both need to be passed-through carefully by syntax transformations.

Pre-terms are further processed by the so-called *check* and *uncheck* phases that are intertwined with type-inference (see also [50]). The latter allows

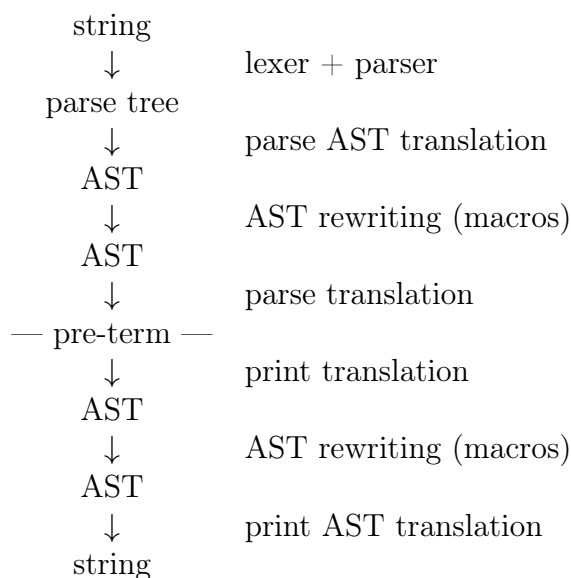


Figure 7.1: Parsing and printing with translations

to operate on higher-order abstract syntax with proper binding and type information already available.

As a rule of thumb, anything that manipulates bindings of variables or constants needs to be implemented as syntax transformation (see below). Anything else is better done via *check/uncheck*: a prominent example application is the **abbreviation** concept of Isabelle/Pure.

7.5.1 Abstract syntax trees

The ML datatype `Ast.ast` explicitly represents the intermediate AST format that is used for syntax rewriting (§7.5.2). It is defined in ML as follows:

```

datatype ast =
  Constant of string |
  Variable of string |
  Appl of ast list
  
```

An AST is either an atom (constant or variable) or a list of (at least two) subtrees. Occasional diagnostic output of ASTs uses notation that resembles S-expression of LISP. Constant atoms are shown as quoted strings, variable atoms as non-quoted strings and applications as a parenthesized list of subtrees. For example, the AST

```
Ast.Appl
```

```
[Ast.Constant "_abs", Ast.Variable "x", Ast.Variable "t"]
```

is pretty-printed as `("_abs" x t)`. Note that `()` and `(x)` are excluded as ASTs, because they have too few subtrees.

AST application is merely a pro-forma mechanism to indicate certain syntactic structures. Thus `(c a b)` could mean either term application or type application, depending on the syntactic context.

Nested application like `(("_abs" x t) u)` is also possible, but ASTs are definitely first-order: the syntax constant `"_abs"` does not bind the `x` in any way. Proper bindings are introduced in later stages of the term syntax, where `("_abs" x t)` becomes an `Abs` node and occurrences of `x` in `t` are replaced by bound variables (represented as de-Bruijn indices).

AST constants versus variables

Depending on the situation — input syntax, output syntax, translation patterns — the distinction of atomic ASTs as `Ast.Constant` versus `Ast.Variable` serves slightly different purposes.

Input syntax of a term such as $f a b = c$ does not yet indicate the scopes of atomic entities f , a , b , c : they could be global constants or local variables, even bound ones depending on the context of the term. `Ast.Variable` leaves this choice still open: later syntax layers (or translation functions) may capture such a variable to determine its role specifically, to make it a constant, bound variable, free variable etc. In contrast, syntax translations that introduce already known constants would rather do it via `Ast.Constant` to prevent accidental re-interpretation later on.

Output syntax turns term constants into `Ast.Constant` and variables (free or schematic) into `Ast.Variable`. This information is precise when printing fully formal λ -terms.

AST translation patterns (§7.5.2) that represent terms cannot distinguish constants and variables syntactically. Explicit indication of `CONST c` inside the term language is required, unless c is known as special *syntax constant* (see also **syntax**). It is also possible to use **syntax** declarations (without mixfix annotation) to enforce that certain unqualified names are always treated as constant within the syntax machinery.

The situation is simpler for ASTs that represent types or sorts, since the concrete syntax already distinguishes type variables from type constants (constructors). So `('a, 'b) foo` corresponds to an AST application of some constant

for *foo* and variable arguments for '*a*' and '*b*'. Note that the postfix application is merely a feature of the concrete syntax, while in the AST the constructor occurs in head position.

Authentic syntax names

Naming constant entities within ASTs is another delicate issue. Unqualified names are resolved in the name space tables in the last stage of parsing, after all translations have been applied. Since syntax transformations do not know about this later name resolution, there can be surprises in boundary cases.

Authentic syntax names for `Ast.Constant` avoid this problem: the fully-qualified constant name with a special prefix for its formal category (*class*, *type*, *const*, *fixed*) represents the information faithfully within the untyped AST format. Accidental overlap with free or bound variables is excluded as well. Authentic syntax names work implicitly in the following situations:

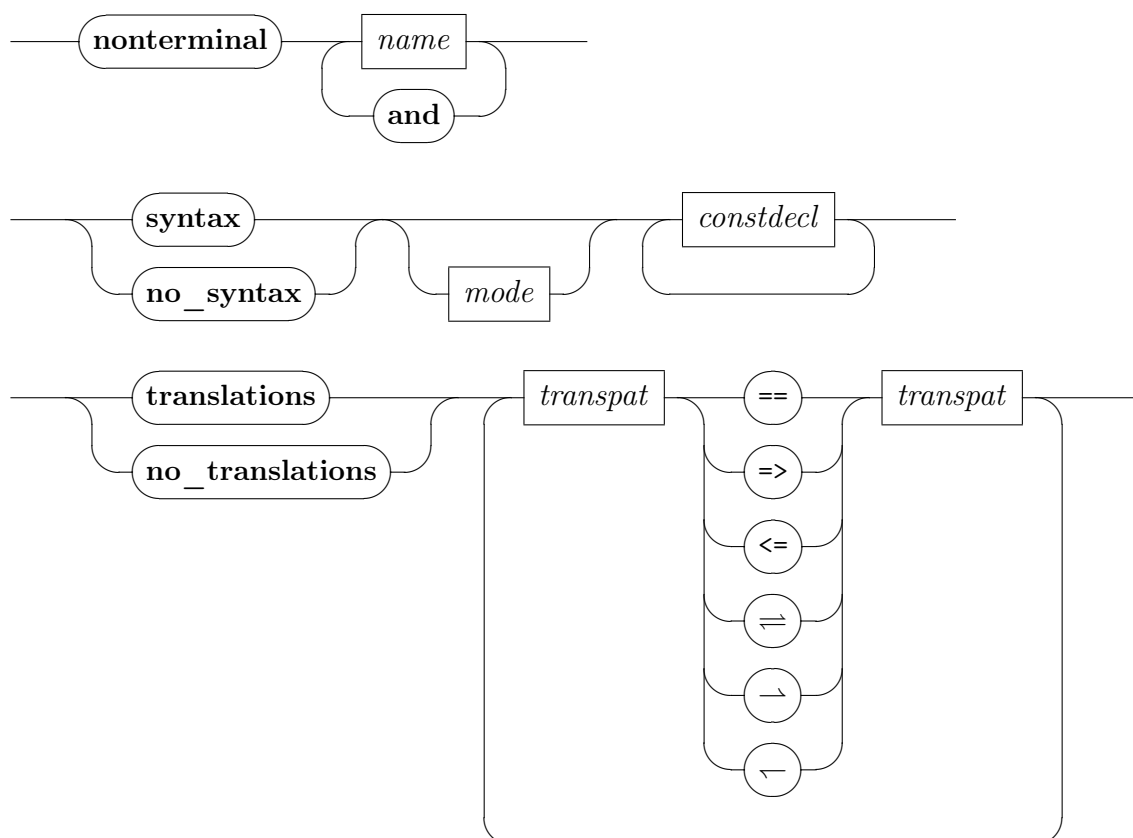
- Input of term constants (or fixed variables) that are introduced by concrete syntax via **notation**: the correspondence of a particular grammar production to some known term entity is preserved.
- Input of type constants (constructors) and type classes — thanks to explicit syntactic distinction independently on the context.
- Output of term constants, type constants, type classes — this information is already available from the internal term to be printed.

In other words, syntax transformations that operate on input terms written as prefix applications are difficult to make robust. Luckily, this case rarely occurs in practice, because syntax forms to be translated usually correspond to some concrete notation.


7.5.2 Raw syntax and translations

nonterminal	:	<i>theory</i>	→	<i>theory</i>	
syntax	:	<i>theory</i>	→	<i>theory</i>	
no_syntax	:	<i>theory</i>	→	<i>theory</i>	
translations	:	<i>theory</i>	→	<i>theory</i>	
no_translations	:	<i>theory</i>	→	<i>theory</i>	
<i>syntax_ast_trace</i>	:	<i>attribute</i>			default <i>false</i>
<i>syntax_ast_stats</i>	:	<i>attribute</i>			default <i>false</i>


Unlike mixfix notation for existing formal entities (§7.3), raw syntax declarations provide full access to the priority grammar of the inner syntax, without any sanity checks. This includes additional syntactic categories (via **nonterminal**) and free-form grammar productions (via **syntax**). Additional syntax translations (or macros, via **translations**) are required to turn resulting parse trees into proper representations of formal entities again.




constdecl



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*) $c :: \sigma (mx)$ augments the priority grammar and the pretty printer table for the given print mode (default ""). An optional keyword **output** means that only the pretty printer table is affected.

Following §7.2, the mixfix annotation $mx = \text{template } ps \ q$ together with type $\sigma = \tau_1 \Rightarrow \dots \tau_n \Rightarrow \tau$ and specify a grammar production. The *template* contains delimiter tokens that surround n argument positions ($_$). The latter correspond to nonterminal symbols A_i derived from the argument types τ_i as follows:

- *prop* if $\tau_i = \text{prop}$
- *logic* if $\tau_i = (\dots)\kappa$ for logical type constructor $\kappa \neq \text{prop}$
- *any* if $\tau_i = \alpha$ for type variables
- κ if $\tau_i = \kappa$ for nonterminal κ (syntactic type constructor)

Each A_i is decorated by priority p_i from the given list ps ; missing priorities default to 0.

The resulting nonterminal of the production is determined similarly from type τ , with priority q and default 1000.

Parsing via this production produces parse trees t_1, \dots, t_n for the argument slots. The resulting parse tree is composed as $c \ t_1 \ \dots \ t_n$, by using the syntax constant c of the syntax declaration.

Such syntactic constants are invented on the spot, without formal check wrt. existing declarations. It is conventional to use plain identifiers

prefixed by a single underscore (e.g. `_foobar`). Names should be chosen with care, to avoid clashes with other syntax declarations.

The special case of copy production is specified by `c = ""` (empty string). It means that the resulting parse tree t is copied directly, without any further decoration.

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) as first-order rewrite rules on ASTs (§7.5.1). The theory context maintains two independent lists translation rules: parse rules (\Rightarrow or \rightarrow) and print rules (\Leftarrow or \leftarrow). For convenience, both can be specified simultaneously as parse / print rules (\Rightarrow or \Leftarrow).

Translation patterns may be prefixed by the syntactic category to be used for parsing; the default is *logic* which means that regular term syntax is used. Both sides of the syntax translation rule undergo parsing and parse AST translations §7.5.3, in order to perform some fundamental normalization like $\lambda x y. b \rightsquigarrow \lambda x. \lambda y. b$, but other AST translation rules are *not* applied recursively here.

When processing AST patterns, the inner syntax lexer runs in a different mode that allows identifiers to start with underscore. This accommodates the usual naming convention for auxiliary syntax constants — those that do not have a logical counter part — by allowing to specify arbitrary AST applications within the term syntax, independently of the corresponding concrete syntax.

Atomic ASTs are distinguished as `Ast.Constant` versus `Ast.Variable` as follows: a qualified name or syntax constant declared via **syntax**, or parse tree head of concrete notation becomes `Ast.Constant`, anything else `Ast.Variable`. Note that `CONST` and `XCONST` within the term language (§7.4.3) allow to enforce treatment as constants.

AST rewrite rules (*lhs*, *rhs*) need to obey the following side-conditions:

- Rules must be left linear: *lhs* must not contain repeated variables.³
- Every variable in *rhs* must also occur in *lhs*.

³The deeper reason for this is that AST equality is not well-defined: different occurrences of the “same” AST could be decorated differently by accidental type-constraints or source position information, for example.

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

syntax_ast_trace and *syntax_ast_stats* control diagnostic output in the AST normalization process, when translation rules are applied to concrete input or output.

Raw syntax and translations provides a slightly more low-level access to the grammar and the form of resulting parse trees. It is often possible to avoid this untyped macro mechanism, and use type-safe **abbreviation** or **notation** instead. Some important situations where **syntax** and **translations** are really need are as follows:

- Iterated replacement via recursive **translations**. For example, consider list enumeration $[a, b, c, d]$ as defined in theory *List* in Isabelle/HOL.
- Change of binding status of variables: anything beyond the built-in **binder** mixfix annotation requires explicit syntax translations. For example, consider list filter comprehension $[x \leftarrow xs . P]$ as defined in theory *List* in Isabelle/HOL.

Applying translation rules

As a term is being parsed or printed, an AST is generated as an intermediate form according to figure 7.1. The AST is normalized by applying translation rules in the manner of a first-order term rewriting system. We first examine how a single rule is applied.

Let t be the abstract syntax tree to be normalized and (lhs, rhs) some translation rule. A subtree u of t is called *redex* if it is an instance of lhs ; in this case the pattern lhs is said to match the object u . A redex matched by lhs may be replaced by the corresponding instance of rhs , thus *rewriting* the AST t . Matching requires some notion of *place-holders* in rule patterns: **Ast.Variable** serves this purpose.

More precisely, the matching of the object u against the pattern lhs is performed as follows:

- Objects of the form **Ast.Variable** x or **Ast.Constant** x are matched by pattern **Ast.Constant** x . Thus all atomic ASTs in the object are treated as (potential) constants, and a successful match makes them actual constants even before name space resolution (see also §7.5.1).

- Object u is matched by pattern `Ast.Variable` x , binding x to u .
- Object `Ast.App1` us is matched by `Ast.App1` ts if us and ts have the same length and each corresponding subtree matches.
- In every other case, matching fails.

A successful match yields a substitution that is applied to rhs , generating the instance that replaces u .

Normalizing an AST involves repeatedly applying translation rules until none are applicable. This works yoyo-like: top-down, bottom-up, top-down, etc. At each subtree position, rules are chosen in order of appearance in the theory definitions.

The configuration options `syntax_ast_trace` and `syntax_ast_stats` might help to understand this process and diagnose problems.

! If syntax translation rules work incorrectly, the output of `print_syntax` with its `rules` sections reveals the actual internal forms of AST pattern, without potentially confusing concrete syntax. Recall that AST constants appear as quoted strings and variables without quotes.

! If `eta_contract` is set to `true`, terms will be η -contracted *before* the AST rewriter sees them. Thus some abstraction nodes needed for print rules to match may vanish. For example, `Ball A ($\lambda x. P x$)` would contract to `Ball A P` and the standard print rule would fail to apply. This problem can be avoided by hand-written ML translation functions (see also §7.5.3), which is in fact the same mechanism used in built-in `binder` declarations.

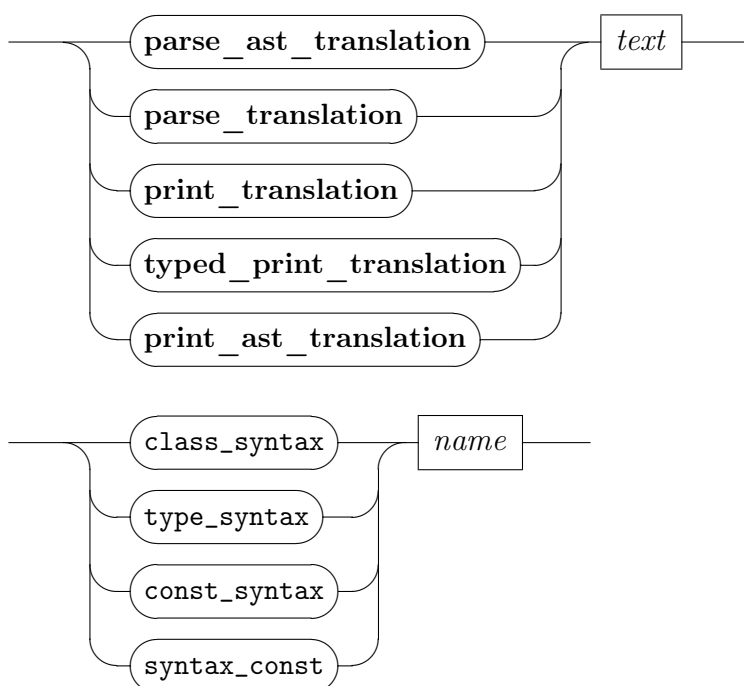
7.5.3 Syntax translation functions

```

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
  class_syntax       : ML antiquotation
  type_syntax        : ML antiquotation
  const_syntax       : ML antiquotation
  syntax_const       : ML antiquotation

```

Syntax translation functions written in ML admit almost arbitrary manipulations of inner syntax, at the expense of some complexity and obscurity in the implementation.



`parse_translation` etc. declare syntax translation functions to the theory. Any of these commands have a single *text* argument that refers to an ML expression of appropriate type as follows:

```

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

```

The argument list consists of (c, tr) pairs, where c is the syntax name of the formal entity involved, and tr a function that translates a syntax form c *args* into tr *ctxt args* (depending on the context). The Isabelle/ML naming convention for parse translations is c_tr and for print translations c_tr' .

The `print_syntax` command displays the sets of names associated with the translation functions of a theory under `parse_ast_translation` etc.

`@{class_syntax c}`, `@{type_syntax c}`, `@{const_syntax c}` inline the authentic syntax name of the given formal entities into the ML source. This is the fully-qualified logical name prefixed by a special marker to indicate its kind: thus different logical name spaces are properly distinguished within parse trees.

`@{const_syntax c}` inlines the name `c` of the given syntax constant, having checked that it has been declared via some `syntax` commands within the theory context. Note that the usual naming convention makes syntax constants start with underscore, to reduce the chance of accidental clashes with other names occurring in parse trees (unqualified constants etc.).

The translation strategy

The different kinds of translation functions are invoked during the transformations between parse trees, ASTs and syntactic terms (cf. figure 7.1). Whenever a combination of the form `c x1 ... xn` is encountered, and a translation function `f` of appropriate kind is declared for `c`, the result is produced by evaluation of `f [x1, ..., xn]` in ML.

For AST translations, the arguments `x1, ..., xn` are ASTs. A combination has the form `Ast.Constant c` or `Ast.Appl [Ast.Constant c, x1, ..., xn]`. For term translations, the arguments are terms and a combination has the form `Const (c, τ)` or `Const (c, τ) $ x1 $... $ xn`. Terms allow more sophisticated transformations than ASTs do, typically involving abstractions and bound variables. *Typed* print translations may even peek at the type `τ` of the constant they are invoked on, although some information might have been suppressed for term output already.

Regardless of whether they act on ASTs or terms, translation functions called during the parsing process differ from those for printing in their overall behaviour:

Parse translations are applied bottom-up. The arguments are already in translated form. The translations must not fail; exceptions trigger an error message. There may be at most one function associated with any syntactic name.

Print translations are applied top-down. They are supplied with arguments that are partly still in internal form. The result again undergoes translation; therefore a print translation should not introduce as head the very constant that invoked it. The function may raise exception `Match` to indicate failure; in this event it has no effect. Multiple functions associated with some syntactic name are tried in the order of declaration in the theory.

Only constant atoms — constructor `Ast.Constant` for ASTs and `Const` for terms — can invoke translation functions. This means that parse translations can only be associated with parse tree heads of concrete syntax, or syntactic constants introduced via other translations. For plain identifiers within the term language, the status of constant versus variable is not yet known during parsing. This is in contrast to print translations, where constants are explicitly known from the given term in its fully internal form.

7.5.4 Built-in syntax transformations

Here are some further details of the main syntax transformation phases of figure 7.1.

Transforming parse trees to ASTs

The parse tree is the raw output of the parser. It is transformed into an AST according to some basic scheme that may be augmented by AST translation functions as explained in §7.5.3.

The parse tree is constructed by nesting the right-hand sides of the productions used to recognize the input. Such parse trees are simply lists of tokens and constituent parse trees, the latter representing the nonterminals of the productions. Ignoring AST translation functions, parse trees are transformed to ASTs by stripping out delimiters and copy productions, while retaining some source position information from input tokens.

The Pure syntax provides predefined AST translations to make the basic λ -term structure more apparent within the (first-order) AST representation, and thus facilitate the use of **translations** (see also §7.5.2). This covers ordinary term application, type application, nested abstraction, iterated meta implications and function types. The effect is illustrated on some representative input strings as follows:

input source	AST
$f x y z$	<code>(f x y z)</code>
$'a ty$	<code>(ty 'a)</code>
$('a, 'b)ty$	<code>(ty 'a 'b)</code>
$\lambda x y z. t$	<code>("_abs" x ("_abs" y ("_abs" z t)))</code>
$\lambda x :: 'a. t$	<code>("_abs" ("_constrain" x 'a) t)</code>
$\llbracket P; Q; R \rrbracket \implies S$	<code>("==>" P ("==>" Q ("==>" R S)))</code>
$['a, 'b, 'c] \Rightarrow 'd$	<code>("fun" 'a ("fun" 'b ("fun" 'c 'd)))</code>

Note that type and sort constraints may occur in further places — translations need to be ready to cope with them. The built-in syntax transformation from parse trees to ASTs insert additional constraints that represent source positions.

Transforming ASTs to terms

After application of macros (§7.5.2), the AST is transformed into a term. This term still lacks proper type information, but it might contain some constraints consisting of applications with head `_constrain`, where the second argument is a type encoded as a pre-term within the syntax. Type inference later introduces correct types, or indicates type errors in the input.

Ignoring parse translations, ASTs are transformed to terms by mapping AST constants to term constants, AST variables to term variables or constants (according to the name space), and AST applications to iterated term applications.

The outcome is still a first-order term. Proper abstractions and bound variables are introduced by parse translations associated with certain syntax constants. Thus `(_abs x x)` eventually becomes a de-Bruijn term `Abs ("x", _, Bound 0)`.

Printing of terms

The output phase is essentially the inverse of the input phase. Terms are translated via abstract syntax trees into pretty-printed text.

Ignoring print translations, the transformation maps term constants, variables and applications to the corresponding constructs on ASTs. Abstractions are mapped to applications of the special constant `_abs` as seen before. Type constraints are represented via special `_constrain` forms, according to various policies of type annotation determined elsewhere. Sort constraints of type variables are handled in a similar fashion.

After application of macros (§7.5.2), the AST is finally pretty-printed. The built-in print AST translations reverse the corresponding parse AST translations.

For the actual printing process, the priority grammar (§7.4.2) plays a vital role: productions are used as templates for pretty printing, with argument slots stemming from nonterminals, and syntactic sugar stemming from literal tokens.

Each AST application with constant head c and arguments t_1, \dots, t_n (for $n = 0$ the AST is just the constant c itself) is printed according to the first grammar production of result name c . The required syntax priority of the argument slot is given by its nonterminal $A^{(p)}$. The argument t_i that corresponds to the position of $A^{(p)}$ is printed recursively, and then put in parentheses *if* its priority p requires this. The resulting output is concatenated with the syntactic sugar according to the grammar production.

If an AST application $(c x_1 \dots x_m)$ has more arguments than the corresponding production, it is first split into $((c x_1 \dots x_n) x_{n+1} \dots x_m)$ and then printed recursively as above.

Applications with too few arguments or with non-constant head or without a corresponding production are printed in prefix-form like $f t_1 \dots t_n$ for terms. Multiple productions associated with some name c are tried in order of appearance within the grammar. An occurrence of some AST variable x is printed as x outright.

White space is *not* inserted automatically. If blanks (or breaks) are required to separate tokens, they need to be specified in the mixfix declaration (§7.2).

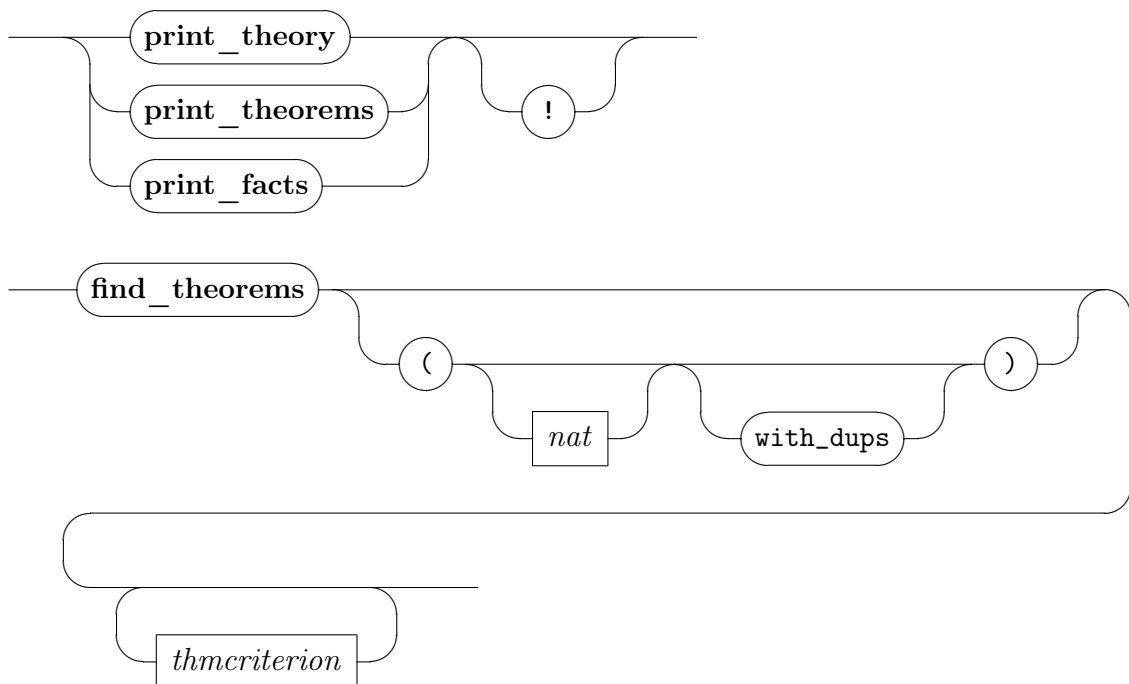
Other commands

8.1 Inspecting the context


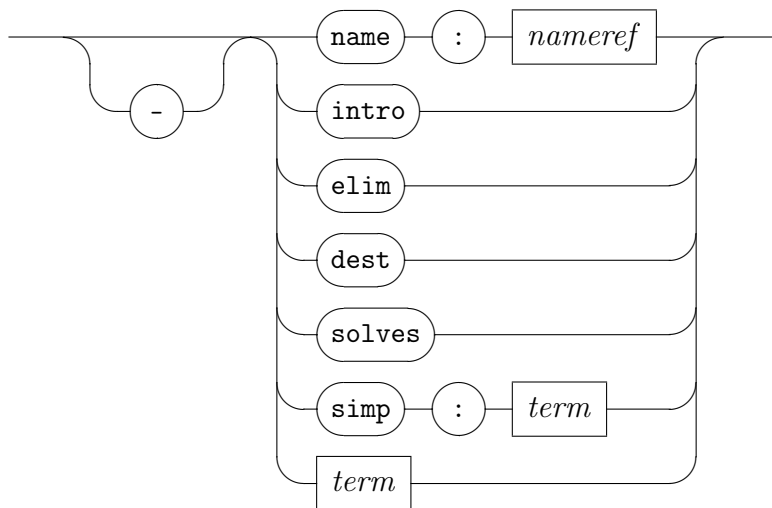
```

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_term_bindings* : context →



```



thmcriteria



constcriterion



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_theory prints the main logical content of the background theory; 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 of the background theory resulting from the last command; the “!” option indicates extra verbosity.

print_facts prints all local facts of the current context, both named and unnamed ones; the “!” option indicates extra verbosity.

print_term_bindings prints all term bindings that are present in the context.

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 [57]).

unused_thms $A_1 \dots A_m - B_1 \dots B_n$ displays all theorems that are proved in theories $B_1 \dots B_n$ or their parents but not in $A_1 \dots A_m$ or their parents and that are never used. 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.

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` like this:

```
declare [[show_main_goal = false]]
```

```
notepad  
begin  
  note [[show_main_goal = true]]  
end
```

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_options : context →
```



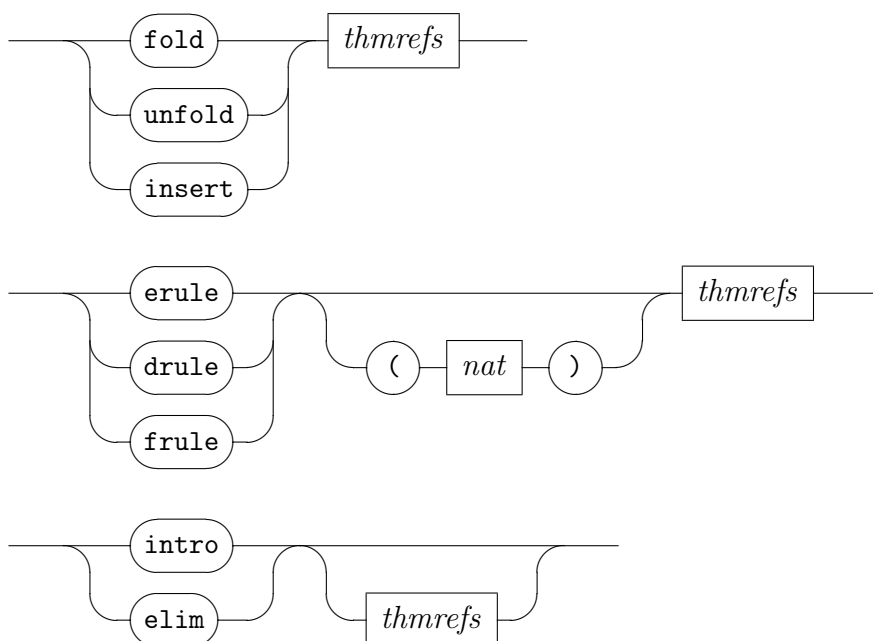
`print_options` prints the available configuration options, with names, types, and current values.

`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*
`intro` : *method*
`elim` : *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 [50]. The optional natural number argument (default 0) specifies additional assumption steps to be performed here.

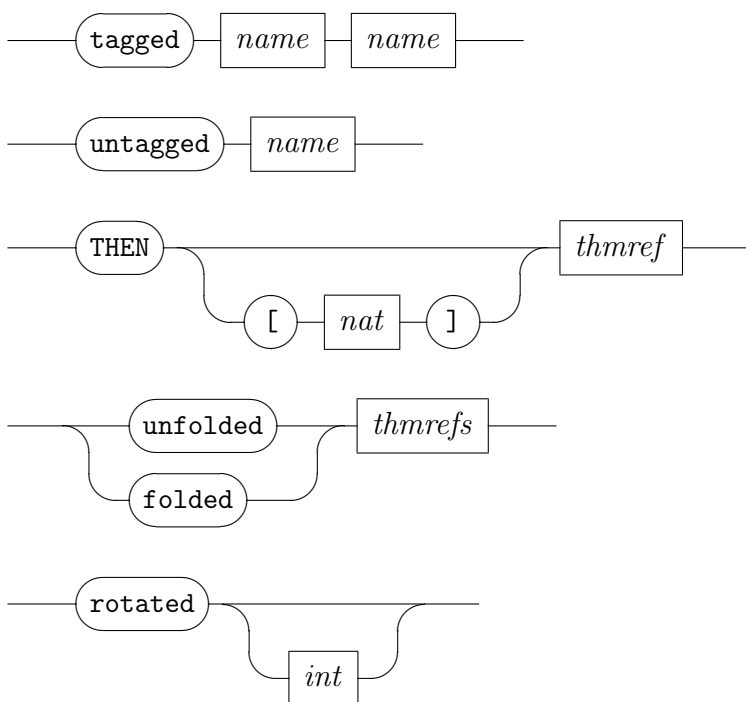
Note that these methods are improper ones, mainly serving for experimentation and tactic script emulation. Different modes of basic rule application are usually expressed in Isar at the proof language level, rather than via implicit proof state manipulations. For example, a proper single-step elimination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

`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.

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
unfolded : attribute
folded : attribute
abs_def : attribute
rotated : attribute
elim_format : 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 composes rules by resolution; it resolves with the first premise of *a* (an alternative position may be also specified). See also **RS** in [50].

unfolded a₁ ... a_n and *folded a₁ ... a_n* expand and fold back again the given definitions throughout a rule.

abs_def turns an equation of the form $f\ x\ y \equiv t$ into $f \equiv \lambda x\ y. t$, which ensures that *simp* or *unfold* steps always expand it. This also works for object-logic equality.

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 $PROP\ A \implies (PROP\ A \implies PROP\ B) \implies PROP\ B$.

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

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. Splitting is performed in the conclusion or some assumption of the subgoal, depending of the structure of the rule.

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

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.

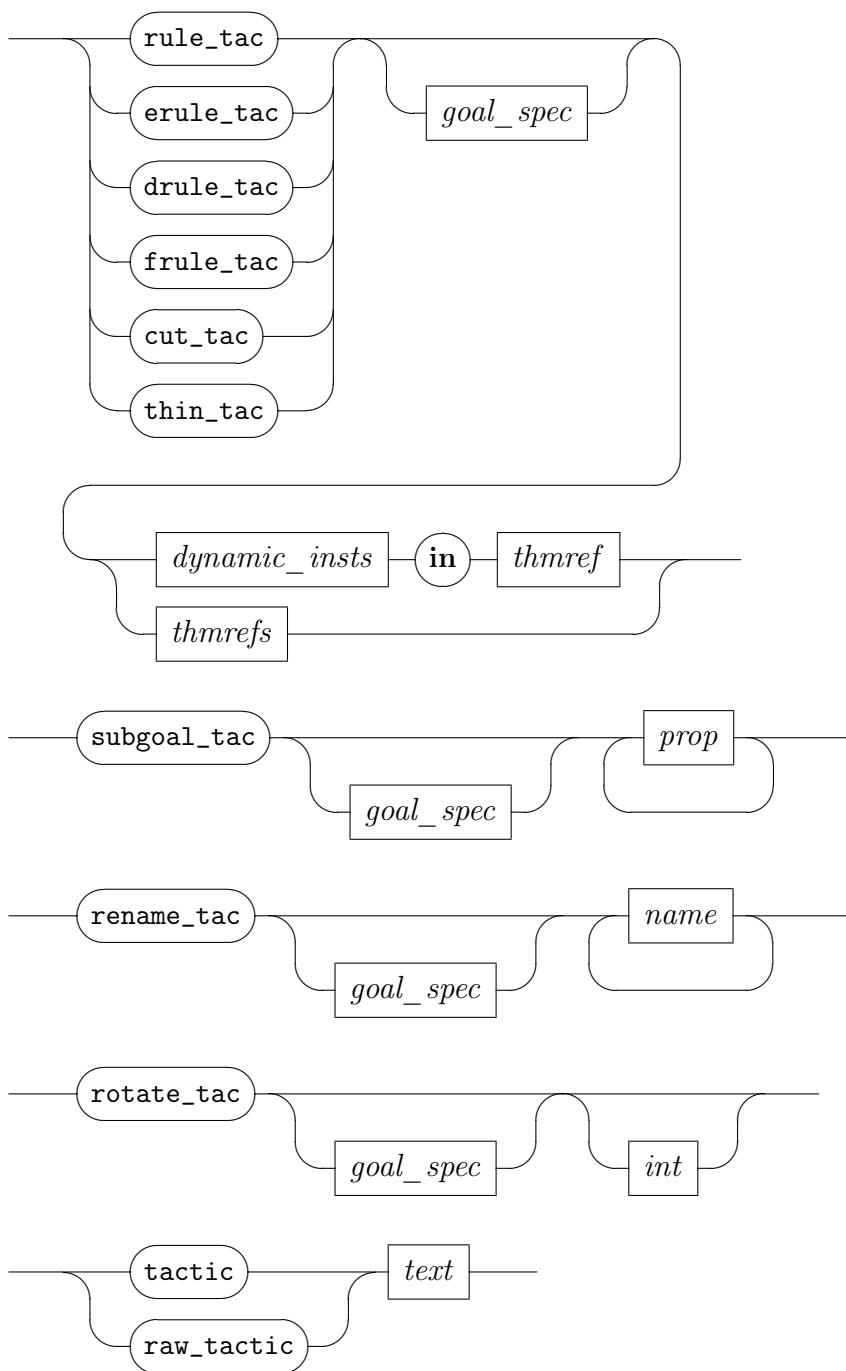
! Dynamic instantiations refer to universally quantified parameters of a subgoal
 • (the dynamic context) rather than fixed variables and term abbreviations of a (static) Isar context.

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
  
```



dynamic_insts



rule_tac etc. do resolution of rules with explicit instantiation. This works the same way as the ML tactics *res_inst_tac* etc. (see [50])

Multiple rules may be only given if there is no instantiation; then *rule_tac* is the same as *resolve_tac* in ML (see [50]).

cut_tac inserts facts into the proof state as assumption of a subgoal; instantiations may be given as well. Note that the scope of schematic variables is spread over the main goal statement and rule premises are turned into new subgoals. This is in contrast to the regular method *insert* which inserts closed rule statements.

thin_tac φ deletes the specified premise from a subgoal. Note that φ may contain schematic variables, to abbreviate the intended proposition; the first matching subgoal premise will be deleted. Removing useless premises from a subgoal increases its readability and can make search tactics run faster.

subgoal_tac $\varphi_1 \dots \varphi_n$ adds the propositions $\varphi_1 \dots \varphi_n$ as local premises to a subgoal, and poses the same as new subgoals (in the original context).

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 premises of a subgoal 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.

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

The Simplifier performs conditional and unconditional rewriting and uses contextual information: rule declarations in the background theory or local proof context are taken into account, as well as chained facts and subgoal premises (“local assumptions”). There are several general hooks that allow to modify the simplification strategy, or incorporate other proof tools that solve sub-problems, produce rewrite rules on demand etc.


The rewriting strategy is always strictly bottom up, except for congruence rules, which are applied while descending into a term. Conditions in conditional rewrite rules are solved recursively before the rewrite rule is applied.

The default Simplifier setup of major object logics (HOL, HOLCF, FOL, ZF) makes the Simplifier ready for immediate use, without engaging into the internal structures. Thus it serves as general-purpose proof tool with the main focus on equational reasoning, and a bit more than that.


9.3.1 Simplification methods

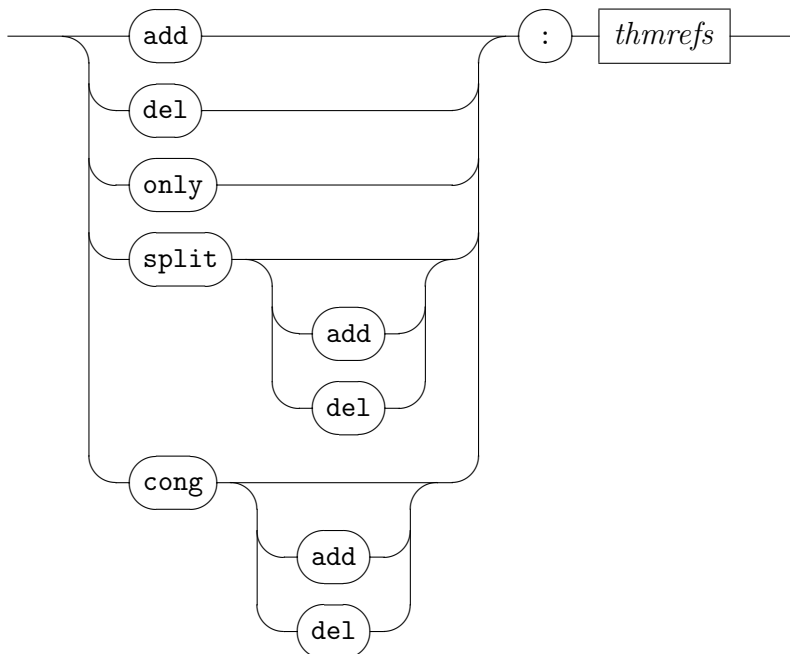
```

      simp      : method
      simp_all  : method
      simp_depth_limit : attribute default 100
  
```



opt



simpmod

simp invokes the Simplifier on the first subgoal, after inserting chained facts as additional goal premises; further rule declarations may be included via (*simp add: facts*). The proof method fails if the subgoal remains unchanged after simplification.

Note that the original goal premises and chained facts are subject to simplification themselves, while declarations via *add/del* merely follow the policies of the object-logic to extract rewrite rules from theorems, without further simplification. This may lead to slightly different behavior in either case, which might be required precisely like that in some boundary situations to perform the intended simplification step!

The *only* modifier first removes all other rewrite rules, looper tactics (including *split* rules), congruence rules, and then behaves like *add*. Implicit solvers remain, which means that trivial rules like reflexivity or introduction of *True* are available to solve the simplified subgoals, but also non-trivial tools like linear arithmetic in HOL. The latter may lead to some surprise of the meaning of “only” in Isabelle/HOL compared to English!

The *split* modifiers add or delete rules for the Splitter (see also §9.3.6 on the looper). 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).

There is also a separate *split* method available for single-step case splitting. The effect of repeatedly applying (*split thms*) can be imitated by “(*simp only: split: thms*)”.

The *cong* modifiers add or delete Simplifier congruence rules (see also §9.3.2); the default is to add.

simp_all is similar to *simp*, but acts on all goals, working backwards from the last to the first one as usual in Isabelle.¹

Chained facts are inserted into all subgoals, before the simplification process starts. Further rule declarations are the same as for *simp*.

The proof method fails if all subgoals remain unchanged after simplification.

simp_depth_limit limits the number of recursive invocations of the Simplifier during conditional rewriting.

By default the Simplifier methods above take local assumptions fully into account, using equational assumptions in the subsequent normalization process, or simplifying assumptions themselves. Further options allow to fine-tune the behavior of the Simplifier in this respect, corresponding to a variety of ML tactics as follows.²

¹The order is irrelevant for goals without schematic variables, so simplification might actually be performed in parallel here.

²Unlike the corresponding Isar proof methods, the ML tactics do not insist in changing the goal state.

Isar method	ML tactic	behavior
<i>(simp (no_asm))</i>	<code>simp_tac</code>	assumptions are ignored completely
<i>(simp (no_asm_simp))</i>	<code>asm_simp_tac</code>	assumptions are used in the simplification of the conclusion but are not themselves simplified
<i>(simp (no_asm_use))</i>	<code>full_simp_tac</code>	assumptions are simplified but are not used in the simplification of each other or the conclusion
<i>(simp)</i>	<code>asm_full_simp_tac</code>	assumptions are used in the simplification of the conclusion and to simplify other assumptions

<code>(simp (asm_lr))</code>	<code>asm_lr_simp_tac</code>	compatibility mode: an assumption is only used for simplifying assumptions which are to the right of it
------------------------------	------------------------------	---

Examples

We consider basic algebraic simplifications in Isabelle/HOL. The rather trivial goal $0 + (x + 0) = x + 0 + 0$ looks like a good candidate to be solved by a single call of *simp*:

lemma $0 + (x + 0) = x + 0 + 0$ **apply** *simp?* **oops**

The above attempt *fails*, because *0* and *op +* in the HOL library are declared as generic type class operations, without stating any algebraic laws yet. More specific types are required to get access to certain standard simplifications of the theory context, e.g. like this:

lemma fixes $x :: nat$ **shows** $0 + (x + 0) = x + 0 + 0$ **by** *simp*

lemma fixes $x :: int$ **shows** $0 + (x + 0) = x + 0 + 0$ **by** *simp*

lemma fixes $x :: 'a :: monoid_add$ **shows** $0 + (x + 0) = x + 0 + 0$ **by** *simp*

In many cases, assumptions of a subgoal are also needed in the simplification process. For example:

lemma fixes $x :: nat$ **shows** $x = 0 \implies x + x = 0$ **by** *simp*

lemma fixes $x :: nat$ **assumes** $x = 0$ **shows** $x + x = 0$ **apply** *simp* **oops**

lemma fixes $x :: nat$ **assumes** $x = 0$ **shows** $x + x = 0$ **using** *assms* **by** *simp*

As seen above, local assumptions that shall contribute to simplification need to be part of the subgoal already, or indicated explicitly for use by the subsequent method invocation. Both too little or too much information can make simplification fail, for different reasons.

In the next example the malicious assumption $\bigwedge x::nat. f x = g (f (g x))$ does not contribute to solve the problem, but makes the default *simp* method loop: the rewrite rule $f ?x \equiv g (f (g ?x))$ extracted from the assumption does not terminate. The Simplifier notices certain simple forms of nontermination, but not this one. The problem can be solved nonetheless, by ignoring assumptions via special options as explained before:

lemma $(\bigwedge x::nat. f x = g (f (g x))) \implies f 0 = f 0 + 0$
by (*simp* (*no_asm*))

The latter form is typical for long unstructured proof scripts, where the control over the goal content is limited. In structured proofs it is usually better to avoid pushing too many facts into the goal state in the first place. Assumptions in the Isar proof context do not intrude the reasoning if not used explicitly. This is illustrated for a toplevel statement and a local proof body as follows:

```
lemma
  assumes  $\bigwedge x::nat. f\ x = g\ (f\ (g\ x))$ 
  shows  $f\ 0 = f\ 0 + 0$  by simp
```

```
notepad
begin
  assume  $\bigwedge x::nat. f\ x = g\ (f\ (g\ x))$ 
  have  $f\ 0 = f\ 0 + 0$  by simp
end
```

Because assumptions may simplify each other, there can be very subtle cases of nontermination. For example, the regular *simp* method applied to $P\ (f\ x) \implies y = x \implies f\ x = f\ y \implies Q$ gives rise to the infinite reduction sequence

$$P\ (f\ x) \xrightarrow{f\ x \equiv f\ y} P\ (f\ y) \xrightarrow{y \equiv x} P\ (f\ x) \xrightarrow{f\ x \equiv f\ y} \dots$$


whereas applying the same to $y = x \implies f\ x = f\ y \implies P\ (f\ x) \implies Q$ terminates (without solving the goal):

```
lemma  $y = x \implies f\ x = f\ y \implies P\ (f\ x) \implies Q$ 
  apply simp
  oops
```

See also §9.3.4 for options to enable Simplifier trace mode, which often helps to diagnose problems with rewrite systems.

9.3.2 Declaring rules

```
simp      : attribute
split    : attribute
cong     : attribute
print_simpset* : context  $\rightarrow$ 
```



simp declares rewrite rules, by adding or deleting them from the simpset within the theory or proof context. Rewrite rules are theorems expressing some form of equality, for example:

$$\begin{aligned} \text{Suc } ?m + ?n &\equiv \text{Suc } ?m + ?n \\ ?P \wedge ?P &\longleftrightarrow ?P \\ ?A \cup ?B &\equiv \{x. x \in ?A \vee x \in ?B\} \end{aligned}$$

Conditional rewrites such as $?m < ?n \implies ?m \text{ div } ?n = 0$ are also permitted; the conditions can be arbitrary formulas.

Internally, all rewrite rules are translated into Pure equalities, theorems with conclusion $lhs \equiv rhs$. The simpset contains a function for extracting equalities from arbitrary theorems, which is usually installed when the object-logic is configured initially. For example, $\neg ?x \in \{\}$ could be turned into $?x \in \{\} \equiv \text{False}$. Theorems that are declared as *simp* and local assumptions within a goal are treated uniformly in this respect.

The Simplifier accepts the following formats for the *lhs* term:

1. First-order patterns, considering the sublanguage of application of constant operators to variable operands, without λ -abstractions or functional variables. For example:

$$\begin{aligned} (?x + ?y) + ?z &\equiv ?x + (?y + ?z) \\ f (f ?x ?y) ?z &\equiv f ?x (f ?y ?z) \end{aligned}$$
2. Higher-order patterns in the sense of [29]. These are terms in β -normal form (this will always be the case unless you have done something strange) where each occurrence of an unknown is of the form $?F x_1 \dots x_n$, where the x_i are distinct bound variables. For example, $(\forall x. ?P x \wedge ?Q x) \equiv (\forall x. ?P x) \wedge (\forall x. ?Q x)$ or its symmetric form, since the *rhs* is also a higher-order pattern.
3. Physical first-order patterns over raw λ -term structure without $\alpha\beta\eta$ -equality; abstractions and bound variables are treated like quasi-constant term material.

For example, the rule $?f \ ?x \in \text{range } ?f = \text{True}$ rewrites the term $g \ a \in \text{range } g$ to True , but will fail to match $g \ (h \ b) \in \text{range } (\lambda x. g \ (h \ x))$. However, offending subterms (in our case $?f \ ?x$, which is not a pattern) can be replaced by adding new variables and conditions like this: $?y = ?f \ ?x \implies ?y \in \text{range } ?f = \text{True}$ is acceptable as a conditional rewrite rule of the second category since conditions can be arbitrary terms.

split declares case split rules.

cong declares congruence rules to the Simplifier context.

Congruence rules are equalities of the form

$$\dots \implies f \ ?x_1 \ \dots \ ?x_n = f \ ?y_1 \ \dots \ ?y_n$$

This controls the simplification of the arguments of f . For example, some arguments can be simplified under additional assumptions:

$$\begin{aligned} ?P_1 \longleftrightarrow ?Q_1 \implies (?Q_1 \implies ?P_2 \longleftrightarrow ?Q_2) \implies \\ (?P_1 \longrightarrow ?P_2) \longleftrightarrow (?Q_1 \longrightarrow ?Q_2) \end{aligned}$$

Given this rule, the Simplifier assumes $?Q_1$ and extracts rewrite rules from it when simplifying $?P_2$. Such local assumptions are effective for rewriting formulae such as $x = 0 \longrightarrow y + x = y$.

The following congruence rule for bounded quantifiers also supplies contextual information — about the bound variable:

$$\begin{aligned} (?A = ?B) \implies (\bigwedge x. x \in ?B \implies ?P \ x \longleftrightarrow ?Q \ x) \implies \\ (\forall x \in ?A. ?P \ x) \longleftrightarrow (\forall x \in ?B. ?Q \ x) \end{aligned}$$

This congruence rule for conditional expressions can supply contextual information for simplifying the arms:

$$\begin{aligned} ?p = ?q \implies (?q \implies ?a = ?c) \implies (\neg ?q \implies ?b = ?d) \implies \\ (\text{if } ?p \text{ then } ?a \text{ else } ?b) = (\text{if } ?q \text{ then } ?c \text{ else } ?d) \end{aligned}$$

A congruence rule can also *prevent* simplification of some arguments. Here is an alternative congruence rule for conditional expressions that conforms to non-strict functional evaluation:

$$?p = ?q \implies (\text{if } ?p \text{ then } ?a \text{ else } ?b) = (\text{if } ?q \text{ then } ?a \text{ else } ?b)$$

Only the first argument is simplified; the others remain unchanged. This can make simplification much faster, but may require an extra case split over the condition $?q$ to prove the goal.

`print_simpset` prints the collection of rules declared to the Simplifier, which is also known as “simpset” internally.

For historical reasons, simpsets may occur independently from the current context, but are conceptually dependent on it. When the Simplifier is invoked via one of its main entry points in the Isar source language (as proof method §9.3.1 or rule attribute §9.3.1), its simpset is derived from the current proof context, and carries a back-reference to that for other tools that might get invoked internally (e.g. simplification procedures §9.3.5). A mismatch of the context of the simpset and the context of the problem being simplified may lead to unexpected results.

The implicit simpset of the theory context is propagated monotonically through the theory hierarchy: forming a new theory, the union of the simpsets of its imports are taken as starting point. Also note that definitional packages like `datatype`, `primrec`, `fun` routinely declare Simplifier rules to the target context, while plain `definition` is an exception in *not* declaring anything.

It is up to the user to manipulate the current simpset further by explicitly adding or deleting theorems as simplification rules, or installing other tools via simplification procedures (§9.3.5). Good simpsets are hard to design. Rules that obviously simplify, like $?n + 0 \equiv ?n$ are good candidates for the implicit simpset, unless a special non-normalizing behavior of certain operations is intended. More specific rules (such as distributive laws, which duplicate subterms) should be added only for specific proof steps. Conversely, sometimes a rule needs to be deleted just for some part of a proof. The need of frequent additions or deletions may indicate a poorly designed simpset.

! The union of simpsets from theory imports (as described above) is not always a good starting point for the new theory. If some ancestors have deleted simplification rules because they are no longer wanted, while others have left those rules in, then the union will contain the unwanted rules, and thus have to be deleted again in the theory body.

9.3.3 Ordered rewriting with permutative rules

A rewrite rule is *permutative* if the left-hand side and right-hand side are the equal up to renaming of variables. The most common permutative rule is

commutativity: $?x + ?y = ?y + ?x$. Other examples include $(?x - ?y) - ?z = (?x - ?z) - ?y$ in arithmetic and $insert\ ?x\ (insert\ ?y\ ?A) = insert\ ?y\ (insert\ ?x\ ?A)$ for sets. Such rules are common enough to merit special attention.

Because ordinary rewriting loops given such rules, the Simplifier employs a special strategy, called *ordered rewriting*. Permutative rules are detected and only applied if the rewriting step decreases the redex wrt. a given term ordering. For example, commutativity rewrites $b + a$ to $a + b$, but then stops, because the redex cannot be decreased further in the sense of the term ordering.

The default is lexicographic ordering of term structure, but this could be also changed locally for special applications via `Simplifier.set_termless` in Isabelle/ML.

Permutative rewrite rules are declared to the Simplifier just like other rewrite rules. Their special status is recognized automatically, and their application is guarded by the term ordering accordingly.

Rewriting with AC operators

Ordered rewriting is particularly effective in the case of associative-commutative operators. (Associativity by itself is not permutative.) When dealing with an AC-operator f , keep the following points in mind:

- The associative law must always be oriented from left to right, namely $f\ (f\ x\ y)\ z = f\ x\ (f\ y\ z)$. The opposite orientation, if used with commutativity, leads to looping in conjunction with the standard term order.
- To complete your set of rewrite rules, you must add not just associativity (A) and commutativity (C) but also a derived rule *left-commutativity* (LC): $f\ x\ (f\ y\ z) = f\ y\ (f\ x\ z)$.

Ordered rewriting with the combination of A, C, and LC sorts a term lexicographically — the rewriting engine imitates bubble-sort.

```

locale AC_example =
  fixes f :: 'a ⇒ 'a ⇒ 'a (infix · 60)
  assumes assoc: (x · y) · z = x · (y · z)
  assumes commute: x · y = y · x
begin

lemma left_commute: x · (y · z) = y · (x · z)

```



```

proof –
  have  $(x \cdot y) \cdot z = (y \cdot x) \cdot z$  by (simp only: commute)
  then show ?thesis by (simp only: assoc)
qed

```

```

lemmas AC_rules = assoc commute left_commute

```

Thus the Simplifier is able to establish equalities with arbitrary permutations of subterms, by normalizing to a common standard form. For example:

```

lemma  $(b \cdot c) \cdot a = xxx$ 
  apply (simp only: AC_rules)

1.  $a \cdot (b \cdot c) = xxx$ 

oops

lemma  $(b \cdot c) \cdot a = a \cdot (b \cdot c)$  by (simp only: AC_rules)
lemma  $(b \cdot c) \cdot a = c \cdot (b \cdot a)$  by (simp only: AC_rules)
lemma  $(b \cdot c) \cdot a = (c \cdot b) \cdot a$  by (simp only: AC_rules)

```

```

end

```

Martin and Nipkow [25] discuss the theory and give many examples; other algebraic structures are amenable to ordered rewriting, such as Boolean rings. The Boyer-Moore theorem prover [11] also employs ordered rewriting.

Re-orienting equalities

Another application of ordered rewriting uses the derived rule *eq_commute*: $(?a = ?b) = (?b = ?a)$ to reverse equations.

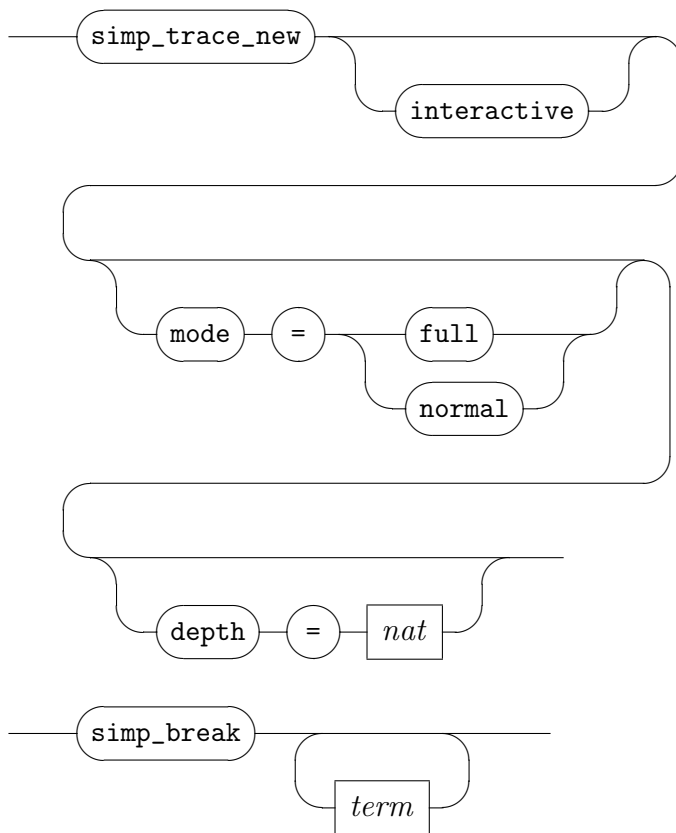
This is occasionally useful to re-orient local assumptions according to the term ordering, when other built-in mechanisms of reorientation and mutual simplification fail to apply.

9.3.4 Simplifier tracing and debugging

```

      simp_trace      : attribute  default false
simp_trace_depth_limit : attribute  default 1
      simp_debug     : attribute  default false
      simp_trace_new : attribute
      simp_break    : attribute

```



These attributes and configurations options control various aspects of Simplifier tracing and debugging.

`simp_trace` makes the Simplifier output internal operations. This includes rewrite steps, but also bookkeeping like modifications of the simpset.

`simp_trace_depth_limit` limits the effect of `simp_trace` to the given depth of recursive Simplifier invocations (when solving conditions of rewrite rules).

`simp_debug` makes the Simplifier output some extra information about internal operations. This includes any attempted invocation of simplification procedures.

`simp_trace_new` controls Simplifier tracing within Isabelle/PIDE applications, notably Isabelle/jEdit [51]. This provides a hierarchical representation of the rewriting steps performed by the Simplifier.

Users can configure the behaviour by specifying breakpoints, verbosity and enabling or disabling the interactive mode. In normal verbosity

(the default), only rule applications matching a breakpoint will be shown to the user. In full verbosity, all rule applications will be logged. Interactive mode interrupts the normal flow of the Simplifier and defers the decision how to continue to the user via some GUI dialog.

simp_break declares term or theorem breakpoints for *simp_trace_new* as described above. Term breakpoints are patterns which are checked for matches on the redex of a rule application. Theorem breakpoints trigger when the corresponding theorem is applied in a rewrite step. For example:

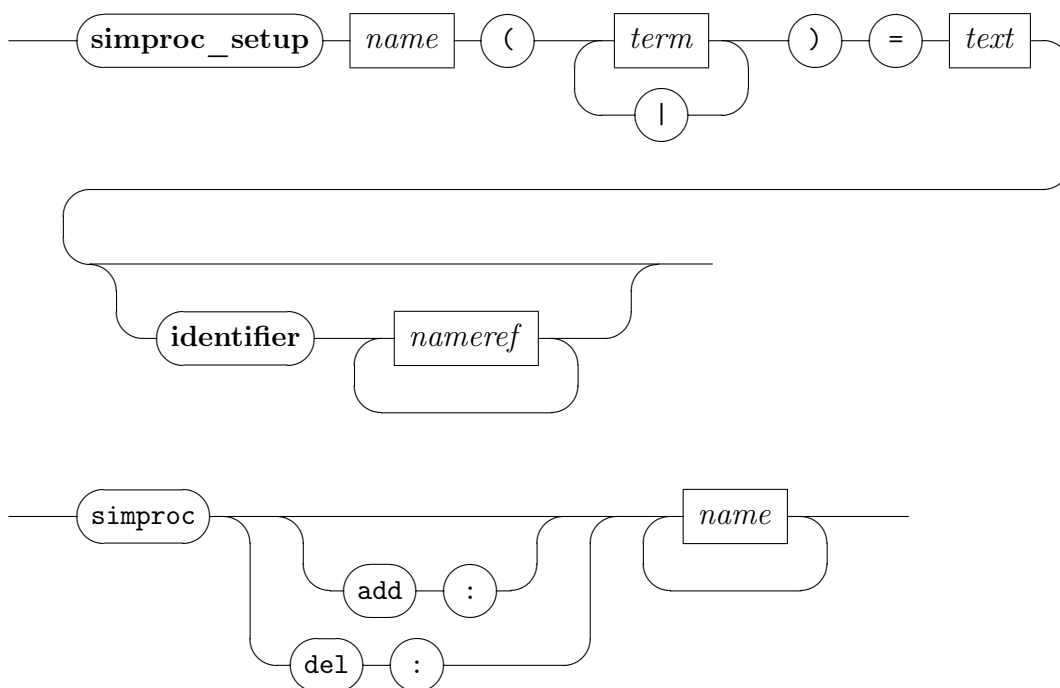
```
declare conjI [simp_break]
declare [[simp_break ?x  $\wedge$  ?y]]
```

9.3.5 Simplification procedures

Simplification procedures are ML functions that produce proven rewrite rules on demand. They are associated with higher-order patterns that approximate the left-hand sides of equations. The Simplifier first matches the current redex against one of the LHS patterns; if this succeeds, the corresponding ML function is invoked, passing the Simplifier context and redex term. Thus rules may be specifically fashioned for particular situations, resulting in a more powerful mechanism than term rewriting by a fixed set of rules.

Any successful result needs to be a (possibly conditional) rewrite rule $t \equiv u$ that is applicable to the current redex. The rule will be applied just as any ordinary rewrite rule. It is expected to be already in *internal form*, bypassing the automatic preprocessing of object-level equivalences.

```
simproc_setup : local_theory  $\rightarrow$  local_theory
      simproc : attribute
```



`simplproc_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 `simplproc_setup` already adds the new `simproc` to the subsequent context.

Example

The following simplification procedure for $(?u::unit) = ()$ in HOL performs fine-grained control over rule application, beyond higher-order pattern matching. Declaring *unit_eq* as *simp* directly would make the Simplifier loop! Note that a version of this simplification procedure is already active in Isabelle/HOL.

```

simproc_setup unit ("x::unit") = {*
  fn _ => fn _ => fn ct =>
    if HOLogic.is_unit (term_of ct) then NONE
    else SOME (mk_meta_eq @{thm unit_eq})
*}

```

Since the Simplifier applies simplification procedures frequently, it is important to make the failure check in ML reasonably fast.

9.3.6 Configurable Simplifier strategies

The core term-rewriting engine of the Simplifier is normally used in combination with some add-on components that modify the strategy and allow to integrate other non-Simplifier proof tools. These may be reconfigured in ML as explained below. Even if the default strategies of object-logics like Isabelle/HOL are used unchanged, it helps to understand how the standard Simplifier strategies work.

The subgoaler

```

Simplifier.set_subgoaler: (Proof.context -> int -> tactic) ->
  Proof.context -> Proof.context
Simplifier.premis_of: Proof.context -> thm list

```

The subgoaler is the tactic used to solve subgoals arising out of conditional rewrite rules or congruence rules. The default should be simplification itself. In rare situations, this strategy may need to be changed. For example, if the premise of a conditional rule is an instance of its conclusion, as in *Suc* $?m < ?n \implies ?m < ?n$, the default strategy could loop.

`Simplifier.set_subgoaler tac ctxt` sets the subgoaler of the context to *tac*. The tactic will be applied to the context of the running Simplifier instance.

`Simplifier.premis_of ctxt` retrieves the current set of premises from the context. This may be non-empty only if the Simplifier has been told to utilize local assumptions in the first place (cf. the options in §9.3.1).

As an example, consider the following alternative subgoaler:

```
ML {*
  fun subgoaler_tac ctxt =
    assume_tac ORELSE'
    resolve_tac (Simplifier.premis_of ctxt) ORELSE'
    asm_simp_tac ctxt
  *}

```

This tactic first tries to solve the subgoal by assumption or by resolving with one of the premises, calling simplification only if that fails.

The solver

```
type solver
Simplifier.mk_solver: string ->
  (Proof.context -> int -> tactic) -> solver
infix setSolver: Proof.context * solver -> Proof.context
infix addSolver: Proof.context * solver -> Proof.context
infix setSSolver: Proof.context * solver -> Proof.context
infix addSSolver: Proof.context * solver -> Proof.context

```

A solver is a tactic that attempts to solve a subgoal after simplification. Its core functionality is to prove trivial subgoals such as *True* and $t = t$, but object-logics might be more ambitious. For example, Isabelle/HOL performs a restricted version of linear arithmetic here.

Solvers are packaged up in abstract type `solver`, with `Simplifier.mk_solver` as the only operation to create a solver.

Rewriting does not instantiate unknowns. For example, rewriting alone cannot prove $a \in ?A$ since this requires instantiating $?A$. The solver, however, is an arbitrary tactic and may instantiate unknowns as it pleases. This is the only way the Simplifier can handle a conditional rewrite rule whose condition contains extra variables. When a simplification tactic is to be combined with other provers, especially with the Classical Reasoner, it is important whether it can be considered safe or not. For this reason a simpset contains two solvers: `safe` and `unsafe`.

The standard simplification strategy solely uses the `unsafe` solver, which is appropriate in most cases. For special applications where the simplification

process is not allowed to instantiate unknowns within the goal, simplification starts with the safe solver, but may still apply the ordinary unsafe one in nested simplifications for conditional rules or congruences. Note that in this way the overall tactic is not totally safe: it may instantiate unknowns that appear also in other subgoals.

`Simplifier.mk_solver name tac` turns *tac* into a solver; the *name* is only attached as a comment and has no further significance.

`ctxt setSSolver solver` installs *solver* as the safe solver of *ctxt*.

`ctxt addSSolver solver` adds *solver* as an additional safe solver; it will be tried after the solvers which had already been present in *ctxt*.

`ctxt setSolver solver` installs *solver* as the unsafe solver of *ctxt*.

`ctxt addSolver solver` adds *solver* as an additional unsafe solver; it will be tried after the solvers which had already been present in *ctxt*.

The solver tactic is invoked with the context of the running `Simplifier`. Further operations may be used to retrieve relevant information, such as the list of local `Simplifier` premises via `Simplifier.premis_of` — this list may be non-empty only if the `Simplifier` runs in a mode that utilizes local assumptions (see also §9.3.1). The solver is also presented the full goal including its assumptions in any case. Thus it can use these (e.g. by calling `assume_tac`), even if the `Simplifier` proper happens to ignore local premises at the moment.

As explained before, the subgoal is also used to solve the premises of congruence rules. These are usually of the form $s = ?x$, where *s* needs to be simplified and *?x* needs to be instantiated with the result. Typically, the subgoal will invoke the `Simplifier` at some point, which will eventually call the solver. For this reason, solver tactics must be prepared to solve goals of the form $t = ?x$, usually by reflexivity. In particular, reflexivity should be tried before any of the fancy automated proof tools.

It may even happen that due to simplification the subgoal is no longer an equality. For example, $False \longleftrightarrow ?Q$ could be rewritten to $\neg ?Q$. To cover this case, the solver could try resolving with the theorem $\neg False$ of the object-logic.

! If a premise of a congruence rule cannot be proved, then the congruence is ignored. This should only happen if the rule is *conditional* — that is, contains premises not of the form $t = ?x$. Otherwise it indicates that some congruence rule, or possibly the subgoal or solver, is faulty.

The looper

```

infix setloop: Proof.context *
  (Proof.context -> int -> tactic) -> Proof.context
infix addloop: Proof.context *
  (string * (Proof.context -> int -> tactic))
  -> Proof.context
infix delloop: Proof.context * string -> Proof.context
Splitter.add_split: thm -> Proof.context -> Proof.context
Splitter.del_split: thm -> Proof.context -> Proof.context

```

The looper is a list of tactics that are applied after simplification, in case the solver failed to solve the simplified goal. If the looper succeeds, the simplification process is started all over again. Each of the subgoals generated by the looper is attacked in turn, in reverse order.

A typical looper is *case splitting*: the expansion of a conditional. Another possibility is to apply an elimination rule on the assumptions. More adventurous loopers could start an induction.

ctxt setloop tac installs *tac* as the only looper tactic of *ctxt*.

ctxt addloop (name, tac) adds *tac* as an additional looper tactic with name *name*, which is significant for managing the collection of loopers. The tactic will be tried after the looper tactics that had already been present in *ctxt*.

ctxt delloop name deletes the looper tactic that was associated with *name* from *ctxt*.

`Splitter.add_split thm ctxt` adds split tactics for *thm* as additional looper tactics of *ctxt*.

`Splitter.del_split thm ctxt` deletes the split tactic corresponding to *thm* from the looper tactics of *ctxt*.

The splitter replaces applications of a given function; the right-hand side of the replacement can be anything. For example, here is a splitting rule for conditional expressions:

$$?P \text{ (if } ?Q \text{ ?x ?y)} \longleftrightarrow (?Q \longrightarrow ?P \text{ ?x}) \wedge (\neg ?Q \longrightarrow ?P \text{ ?y})$$

Another example is the elimination operator for Cartesian products (which happens to be called *split* in Isabelle/HOL):

$$?P \text{ (split } ?f \text{ } ?p) \longleftrightarrow (\forall a \ b. \ ?p = (a, b) \longrightarrow ?P \ (f \ a \ b))$$


For technical reasons, there is a distinction between case splitting in the conclusion and in the premises of a subgoal. The former is done by `Splitter.split_tac` with rules like `split_if` or `option.split`, which do not split the subgoal, while the latter is done by `Splitter.split_asm_tac` with rules like `split_if_asm` or `option.split_asm`, which split the subgoal. The function `Splitter.add_split` automatically takes care of which tactic to call, analyzing the form of the rules given as argument; it is the same operation behind `split` attribute or method modifier syntax in the Isar source language.

Case splits should be allowed only when necessary; they are expensive and hard to control. Case-splitting on if-expressions in the conclusion is usually beneficial, so it is enabled by default in Isabelle/HOL and Isabelle/FOL/ZF.


! With `Splitter.split_asm_tac` as looper component, the Simplifier may split subgoals! This might cause unexpected problems in tactic expressions that silently assume 0 or 1 subgoals after simplification.

9.3.7 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 looper tactics (§9.3.6) are *not* involved here. The *simplified* attribute should be only rarely required under normal circumstances.

9.4 The Classical Reasoner

9.4.1 Basic concepts

Although Isabelle is generic, many users will be working in some extension of classical first-order logic. Isabelle/ZF is built upon theory FOL, while Isabelle/HOL conceptually contains first-order logic as a fragment. Theorem-proving in predicate logic is undecidable, but many automated strategies have been developed to assist in this task.

Isabelle’s classical reasoner is a generic package that accepts certain information about a logic and delivers a suite of automatic proof tools, based on rules that are classified and declared in the context. These proof procedures are slow and simplistic compared with high-end automated theorem provers, but they can save considerable time and effort in practice. They can prove theorems such as Pelletier’s [42] problems 40 and 41 in a few milliseconds (including full proof reconstruction):

lemma $(\exists y. \forall x. F x y \longleftrightarrow F x x) \longrightarrow \neg (\forall x. \exists y. \forall z. F z y \longleftrightarrow \neg F z x)$
 by *blast*

lemma $(\forall z. \exists y. \forall x. f x y \longleftrightarrow f x z \wedge \neg f x x) \longrightarrow \neg (\exists z. \forall x. f x z)$
 by *blast*

The proof tools are generic. They are not restricted to first-order logic, and have been heavily used in the development of the Isabelle/HOL library and applications. The tactics can be traced, and their components can be called directly; in this manner, any proof can be viewed interactively.

The sequent calculus

Isabelle supports natural deduction, which is easy to use for interactive proof. But natural deduction does not easily lend itself to automation, and has a bias towards intuitionism. For certain proofs in classical logic, it can not be called natural. The *sequent calculus*, a generalization of natural deduction, is easier to automate.

A **sequent** has the form $\Gamma \vdash \Delta$, where Γ and Δ are sets of formulae.³ The sequent $P_1, \dots, P_m \vdash Q_1, \dots, Q_n$ is **valid** if $P_1 \wedge \dots \wedge P_m$ implies $Q_1 \vee \dots \vee Q_n$. Thus P_1, \dots, P_m represent assumptions, each of which is true, while Q_1, \dots, Q_n represent alternative goals. A sequent is **basic** if its left and right sides have a common formula, as in $P, Q \vdash Q, R$; basic sequents are trivially valid.

Sequent rules are classified as **right** or **left**, indicating which side of the \vdash symbol they operate on. Rules that operate on the right side are analogous to natural deduction's introduction rules, and left rules are analogous to elimination rules. The sequent calculus analogue of $(\rightarrow I)$ is the rule

$$\frac{P, \Gamma \vdash \Delta, Q}{\Gamma \vdash \Delta, P \rightarrow Q} (\rightarrow R)$$

Applying the rule backwards, this breaks down some implication on the right side of a sequent; Γ and Δ stand for the sets of formulae that are unaffected by the inference. The analogue of the pair $(\vee I1)$ and $(\vee I2)$ is the single rule

$$\frac{\Gamma \vdash \Delta, P, Q}{\Gamma \vdash \Delta, P \vee Q} (\vee R)$$

This breaks down some disjunction on the right side, replacing it by both disjuncts. Thus, the sequent calculus is a kind of multiple-conclusion logic.

To illustrate the use of multiple formulae on the right, let us prove the classical theorem $(P \rightarrow Q) \vee (Q \rightarrow P)$. Working backwards, we reduce this formula to a basic sequent:

$$\frac{\frac{\frac{P, Q \vdash Q, P}{P \vdash Q, (Q \rightarrow P)} (\rightarrow R)}{\vdash (P \rightarrow Q), (Q \rightarrow P)} (\rightarrow R)}{\vdash (P \rightarrow Q) \vee (Q \rightarrow P)} (\vee R)$$

This example is typical of the sequent calculus: start with the desired theorem and apply rules backwards in a fairly arbitrary manner. This yields a surprisingly effective proof procedure. Quantifiers add only few complications, since Isabelle handles parameters and schematic variables. See [41, Chapter 10] for further discussion.

³For first-order logic, sequents can equivalently be made from lists or multisets of formulae.

Simulating sequents by natural deduction

Isabelle can represent sequents directly, as in the object-logic LK. But natural deduction is easier to work with, and most object-logics employ it. Fortunately, we can simulate the sequent $P_1, \dots, P_m \vdash Q_1, \dots, Q_n$ by the Isabelle formula $P_1 \Longrightarrow \dots \Longrightarrow P_m \Longrightarrow \neg Q_2 \Longrightarrow \dots \Longrightarrow \neg Q_n \Longrightarrow Q_1$ where the order of the assumptions and the choice of Q_1 are arbitrary. Elim-resolution plays a key role in simulating sequent proofs.

We can easily handle reasoning on the left. Elim-resolution with the rules $(\vee E)$, $(\perp E)$ and $(\exists E)$ achieves a similar effect as the corresponding sequent rules. For the other connectives, we use sequent-style elimination rules instead of destruction rules such as $(\wedge E1, 2)$ and $(\forall E)$. But note that the rule $(\neg L)$ has no effect under our representation of sequents!

$$\frac{\Gamma \vdash \Delta, P}{\neg P, \Gamma \vdash \Delta} (\neg L)$$

What about reasoning on the right? Introduction rules can only affect the formula in the conclusion, namely Q_1 . The other right-side formulae are represented as negated assumptions, $\neg Q_2, \dots, \neg Q_n$. In order to operate on one of these, it must first be exchanged with Q_1 . Elim-resolution with the *swap* rule has this effect: $\neg P \Longrightarrow (\neg R \Longrightarrow P) \Longrightarrow R$

To ensure that swaps occur only when necessary, each introduction rule is converted into a swapped form: it is resolved with the second premise of (*swap*). The swapped form of $(\wedge I)$, which might be called $(\neg \wedge E)$, is

$$\neg (P \wedge Q) \Longrightarrow (\neg R \Longrightarrow P) \Longrightarrow (\neg R \Longrightarrow Q) \Longrightarrow R$$

Similarly, the swapped form of $(\rightarrow I)$ is

$$\neg (P \rightarrow Q) \Longrightarrow (\neg R \Longrightarrow P \Longrightarrow Q) \Longrightarrow R$$

Swapped introduction rules are applied using elim-resolution, which deletes the negated formula. Our representation of sequents also requires the use of ordinary introduction rules. If we had no regard for readability of intermediate goal states, we could treat the right side more uniformly by representing sequents as

$$P_1 \Longrightarrow \dots \Longrightarrow P_m \Longrightarrow \neg Q_1 \Longrightarrow \dots \Longrightarrow \neg Q_n \Longrightarrow \perp$$

Extra rules for the sequent calculus

As mentioned, destruction rules such as $(\wedge E1, 2)$ and $(\forall E)$ must be replaced by sequent-style elimination rules. In addition, we need rules to embody the classical equivalence between $P \longrightarrow Q$ and $\neg P \vee Q$. The introduction rules $(\vee I1, 2)$ are replaced by a rule that simulates $(\vee R)$:

$$(\neg Q \Longrightarrow P) \Longrightarrow P \vee Q$$

The destruction rule $(\longrightarrow E)$ is replaced by

$$(P \longrightarrow Q) \Longrightarrow (\neg P \Longrightarrow R) \Longrightarrow (Q \Longrightarrow R) \Longrightarrow R$$

Quantifier replication also requires special rules. In classical logic, $\exists x. P x$ is equivalent to $\neg (\forall x. \neg P x)$; the rules $(\exists R)$ and $(\forall L)$ are dual:

$$\frac{\Gamma \vdash \Delta, \exists x. P x, P t}{\Gamma \vdash \Delta, \exists x. P x} (\exists R) \quad \frac{P t, \forall x. P x, \Gamma \vdash \Delta}{\forall x. P x, \Gamma \vdash \Delta} (\forall L)$$

Thus both kinds of quantifier may be replicated. Theorems requiring multiple uses of a universal formula are easy to invent; consider

$$(\forall x. P x \longrightarrow P (f x)) \wedge P a \longrightarrow P (f^n a)$$

for any $n > 1$. Natural examples of the multiple use of an existential formula are rare; a standard one is $\exists x. \forall y. P x \longrightarrow P y$.

Forgoing quantifier replication loses completeness, but gains decidability, since the search space becomes finite. Many useful theorems can be proved without replication, and the search generally delivers its verdict in a reasonable time. To adopt this approach, represent the sequent rules $(\exists R)$, $(\exists L)$ and $(\forall R)$ by $(\exists I)$, $(\exists E)$ and $(\forall I)$, respectively, and put $(\forall E)$ into elimination form:

$$\forall x. P x \Longrightarrow (P t \Longrightarrow Q) \Longrightarrow Q$$

Elim-resolution with this rule will delete the universal formula after a single use. To replicate universal quantifiers, replace the rule by

$$\forall x. P x \Longrightarrow (P t \Longrightarrow \forall x. P x \Longrightarrow Q) \Longrightarrow Q$$

To replicate existential quantifiers, replace $(\exists I)$ by

$$(\neg (\exists x. P x) \Longrightarrow P t) \Longrightarrow \exists x. P x$$

All introduction rules mentioned above are also useful in swapped form.

Replication makes the search space infinite; we must apply the rules with care. The classical reasoner distinguishes between safe and unsafe rules, applying the latter only when there is no alternative. Depth-first search may well go down a blind alley; best-first search is better behaved in an infinite search space. However, quantifier replication is too expensive to prove any but the simplest theorems.

9.4.2 Rule declarations

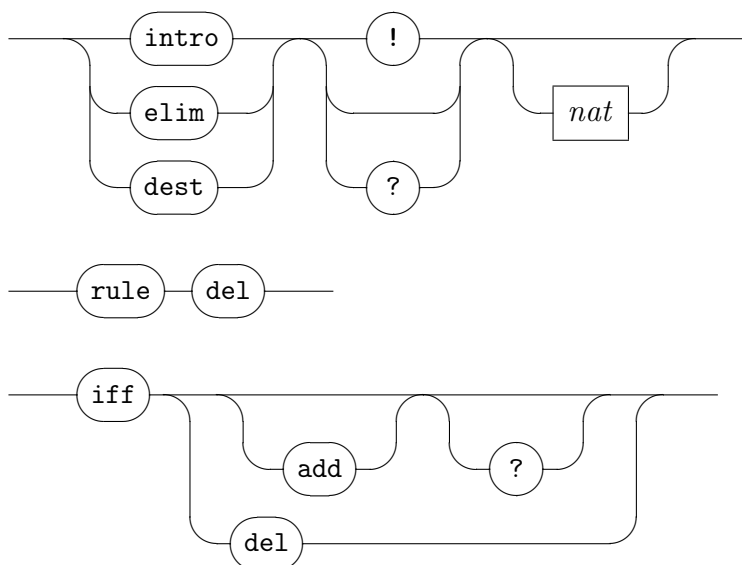
The proof tools of the Classical Reasoner depend on collections of rules declared in the context, which are classified as introduction, elimination or destruction and as *safe* or *unsafe*. In general, safe rules can be attempted blindly, while unsafe rules must be used with care. A safe rule must never reduce a provable goal to an unprovable set of subgoals.

The rule $P \implies P \vee Q$ is unsafe because it reduces $P \vee Q$ to P , which might turn out as premature choice of an unprovable subgoal. Any rule is unsafe whose premises contain new unknowns. The elimination rule $\forall x. P x \implies (P t \implies Q) \implies Q$ is unsafe, since it is applied via elim-resolution, which discards the assumption $\forall x. P x$ and replaces it by the weaker assumption $P t$. The rule $P t \implies \exists x. P x$ is unsafe for similar reasons. The quantifier duplication rule $\forall x. P x \implies (P t \implies \forall x. P x \implies Q) \implies Q$ is unsafe in a different sense: since it keeps the assumption $\forall x. P x$, it is prone to looping. In classical first-order logic, all rules are safe except those mentioned above. The safe / unsafe distinction is vague, and may be regarded merely as a way of giving some rules priority over others. One could argue that $(\forall E)$ is unsafe, because repeated application of it could generate exponentially many subgoals. Induction rules are unsafe because inductive proofs are difficult to set up automatically. Any inference is unsafe that instantiates an unknown in the proof state — thus matching must be used, rather than unification. Even proof by assumption is unsafe if it instantiates unknowns shared with other subgoals.

```

print_claset* : context →
    intro : attribute
    elim : attribute
    dest : attribute
    rule : attribute
    iff : attribute
    swapped : attribute

```



`print_claset` prints the collection of rules declared to the Classical Reasoner, i.e. the `claset` within the context.

`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 the automated reasoning tools, but determines the search order of single rule steps.

Introduction rules are those that can be applied using ordinary resolution. Their swapped forms are generated internally, which will be applied using elim-resolution. Elimination rules are applied using elim-resolution. Rules are sorted by the number of new subgoals they will yield; rules that generate the fewest subgoals will be tried first. Otherwise, later declarations take precedence over earlier ones.

Rules already present in the context with the same classification are ignored. A warning is printed if the rule has already been added with some other classification, but the rule is added anyway as requested.

`rule del` deletes all occurrences of a rule from the classical context, regardless of its classification as introduction / elimination / destruction and safe / unsafe.


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.

swapped turns an introduction rule into an elimination, by resolving with the classical swap principle $\neg P \implies (\neg R \implies P) \implies R$ in the second position. This is mainly for illustrative purposes: the Classical Reasoner already swaps rules internally as explained above.

9.4.3 Structured methods

rule : *method*
contradiction : *method*



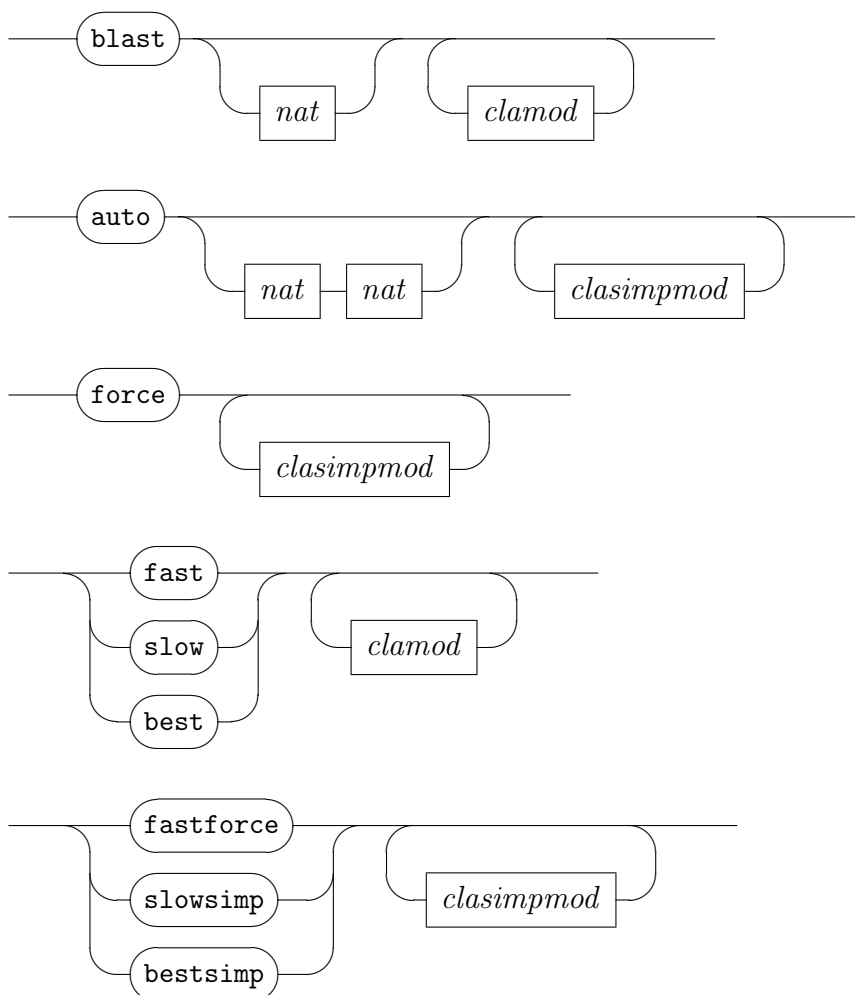
rule as offered by the Classical Reasoner is a refinement over the Pure one (see §6.3.3). Both versions work the same, but the classical version observes the classical rule context in addition to that of Isabelle/Pure.


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.


9.4.4 Fully automated methods

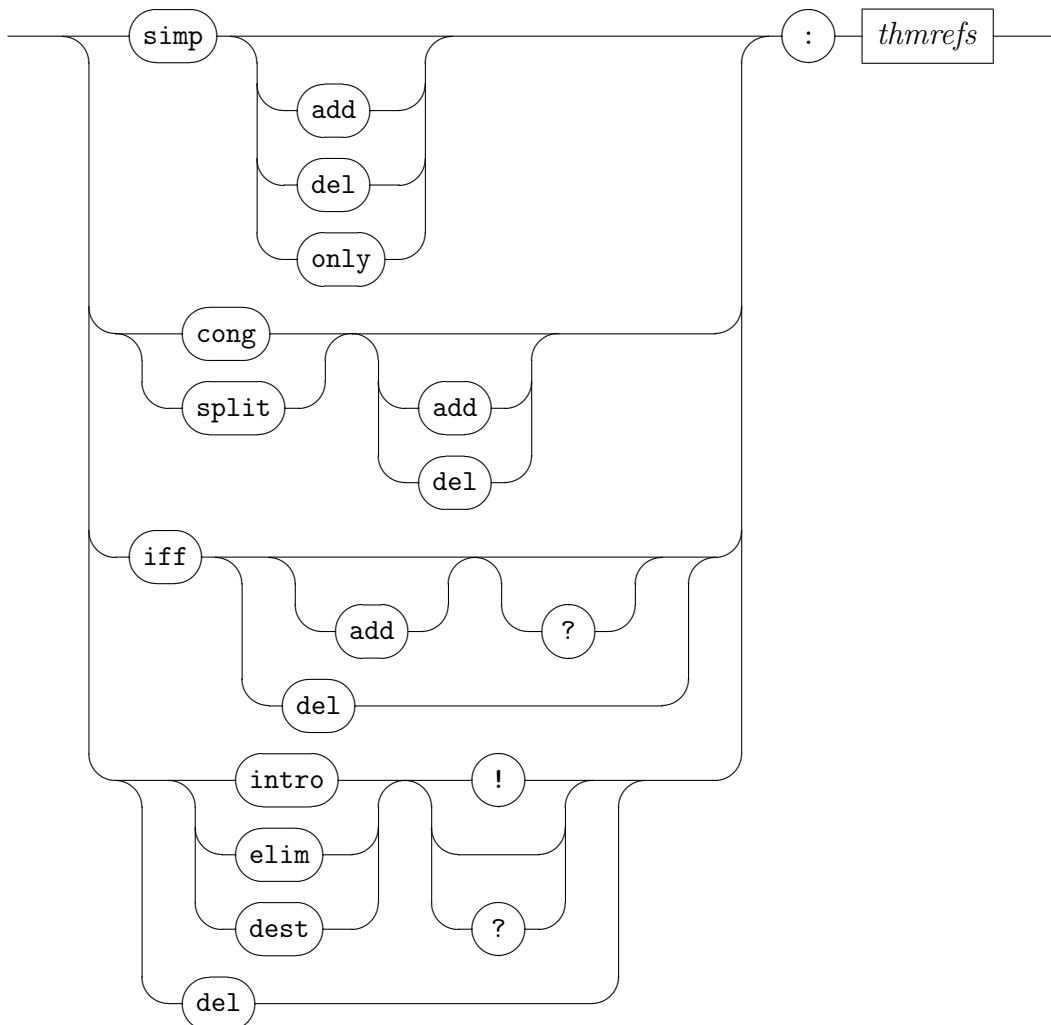
blast : method
auto : method
force : method
fast : method
slow : method
best : method
fastforce : method
slowsimp : method
bestsimp : method
deepen : method





clamod



clasimpmod

blast is a separate classical tableau prover that uses the same classical rule declarations as explained before.

Proof search is coded directly in ML using special data structures. A successful proof is then reconstructed using regular Isabelle inferences. It is faster and more powerful than the other classical reasoning tools, but has major limitations too.

- It does not use the classical wrapper tacticals, such as the integration with the Simplifier of *fastforce*.
- It does not perform higher-order unification, as needed by the

rule $?f \ ?x \in \text{range } ?f$ in HOL. There are often alternatives to such rules, for example $?b = ?f \ ?x \implies ?b \in \text{range } ?f$.

- Function variables may only be applied to parameters of the subgoal. (This restriction arises because the prover does not use higher-order unification.) If other function variables are present then the prover will fail with the message `Function Var's argument not a bound variable`.
- Its proof strategy is more general than *fast* but can be slower. If *blast* fails or seems to be running forever, try *fast* and the other proof tools described below.

The optional integer argument specifies a bound for the number of unsafe steps used in a proof. By default, *blast* starts with a bound of 0 and increases it successively to 20. In contrast, (*blast lim*) tries to prove the goal using a search bound of *lim*. Sometimes a slow proof using *blast* can be made much faster by supplying the successful search bound to this proof method instead.

auto combines classical reasoning with simplification. It is intended for situations where there are a lot of mostly trivial subgoals; it proves all the easy ones, leaving the ones it cannot prove. Occasionally, attempting to prove the hard ones may take a long time.

The optional depth arguments in (*auto m n*) refer to its builtin classical reasoning procedures: *m* (default 4) is for *blast*, which is tried first, and *n* (default 2) is for a slower but more general alternative that also takes wrappers into account.

force is intended to prove the first subgoal completely, using many fancy proof tools and performing a rather exhaustive search. As a result, proof attempts may take rather long or diverge easily.

fast, *best*, *slow* attempt to prove the first subgoal using sequent-style reasoning as explained before. Unlike *blast*, they construct proofs directly in Isabelle.

There is a difference in search strategy and back-tracking: *fast* uses depth-first search and *best* uses best-first search (guided by a heuristic function: normally the total size of the proof state).

Method *slow* is like *fast*, but conducts a broader search: it may, when backtracking from a failed proof attempt, undo even the step of proving a subgoal by assumption.

fastforce, *slowsimp*, *bestsimp* are like *fast*, *slow*, *best*, respectively, but use the Simplifier as additional wrapper. The name *fastforce*, reflects the behaviour of this popular method better without requiring an understanding of its implementation.

deepen works by exhaustive search up to a certain depth. The start depth is 4 (unless specified explicitly), and the depth is increased iteratively up to 10. Unsafe rules are modified to preserve the formula they act on, so that it be used repeatedly. This method can prove more goals than *fast*, but is much slower, for example if the assumptions have many universal quantifiers.

Any of the above methods support additional modifiers of the context of classical (and simplifier) rules, but the ones related to the Simplifier are explicitly prefixed by *simp* here. The semantics of these ad-hoc rule declarations is analogous to the attributes given before. Facts provided by forward chaining are inserted into the goal before commencing proof search.


9.4.5 Partially automated methods

These proof methods may help in situations when the fully-automated tools fail. The result is a simpler subgoal that can be tackled by other means, such as by manual instantiation of quantifiers.

```

safe : method
clarify : method
clarsimp : method

```



safe repeatedly performs safe steps on all subgoals. It is deterministic, with at most one outcome.

clarify performs a series of safe steps without splitting subgoals; see also *clarify_step*.

clarsimp acts like *clarify*, but also does simplification. Note that if the Simplifier context includes a splitter for the premises, the subgoal may still be split.

9.4.6 Single-step tactics

```

safe_step : method
inst_step : method
  step : method
slow_step : method
clarify_step : method

```

These are the primitive tactics behind the automated proof methods of the Classical Reasoner. By calling them yourself, you can execute these procedures one step at a time.

safe_step performs a safe step on the first subgoal. The safe wrapper tacticals are applied to a tactic that may include proof by assumption or Modus Ponens (taking care not to instantiate unknowns), or substitution.

inst_step is like *safe_step*, but allows unknowns to be instantiated.

step is the basic step of the proof procedure, it operates on the first subgoal. The unsafe wrapper tacticals are applied to a tactic that tries *safe*, *inst_step*, or applies an unsafe rule from the context.

slow_step resembles *step*, but allows backtracking between using safe rules with instantiation (*inst_step*) and using unsafe rules. The resulting search space is larger.

clarify_step performs a safe step on the first subgoal; no splitting step is applied. For example, the subgoal $A \wedge B$ is left as a conjunction. Proof by assumption, Modus Ponens, etc., may be performed provided they do not instantiate unknowns. Assumptions of the form $x = t$ may be eliminated. The safe wrapper tactical is applied.

9.4.7 Modifying the search step

```

type wrapper = (int -> tactic) -> (int -> tactic)
infix addWrapper: Proof.context *
  (string * (Proof.context -> wrapper)) -> Proof.context
infix addSbefore: Proof.context *
  (string * (Proof.context -> int -> tactic)) -> Proof.context
infix addSafter: Proof.context *
  (string * (Proof.context -> int -> tactic)) -> Proof.context
infix delWrapper: Proof.context * string -> Proof.context
infix addWrapper: Proof.context *
  (string * (Proof.context -> wrapper)) -> Proof.context
infix addbefore: Proof.context *
  (string * (Proof.context -> int -> tactic)) -> Proof.context
infix addafter: Proof.context *
  (string * (Proof.context -> int -> tactic)) -> Proof.context
infix delWrapper: Proof.context * string -> Proof.context
addSss: Proof.context -> Proof.context
addss: Proof.context -> Proof.context

```

The proof strategy of the Classical Reasoner is simple. Perform as many safe inferences as possible; or else, apply certain safe rules, allowing instantiation of unknowns; or else, apply an unsafe rule. The tactics also eliminate assumptions of the form $x = t$ by substitution if they have been set up to do so. They may perform a form of Modus Ponens: if there are assumptions $P \longrightarrow Q$ and P , then replace $P \longrightarrow Q$ by Q .

The classical reasoning tools — except *blast* — allow to modify this basic proof strategy by applying two lists of arbitrary *wrapper tacticals* to it. The first wrapper list, which is considered to contain safe wrappers only, affects *safe_step* and all the tactics that call it. The second one, which may contain unsafe wrappers, affects the unsafe parts of *step*, *slow_step*, and the tactics that call them. A wrapper transforms each step of the search, for example by attempting other tactics before or after the original step tactic. All members of a wrapper list are applied in turn to the respective step tactic.

Initially the two wrapper lists are empty, which means no modification of the step tactics. Safe and unsafe wrappers are added to a claset with the functions given below, supplying them with wrapper names. These names may be used to selectively delete wrappers.

ctxt addSWrapper (*name*, *wrapper*) adds a new wrapper, which should yield a safe tactic, to modify the existing safe step tactic.

ctxt addSbefore (*name*, *tac*) adds the given tactic as a safe wrapper, such that it is tried *before* each safe step of the search.

ctxt addSafter (*name*, *tac*) adds the given tactic as a safe wrapper, such that it is tried when a safe step of the search would fail.

ctxt delSWrapper *name* deletes the safe wrapper with the given name.

ctxt addWrapper (*name*, *wrapper*) adds a new wrapper to modify the existing (unsafe) step tactic.

ctxt addbefore (*name*, *tac*) adds the given tactic as an unsafe wrapper, such that its result is concatenated *before* the result of each unsafe step.

ctxt addafter (*name*, *tac*) adds the given tactic as an unsafe wrapper, such that its result is concatenated *after* the result of each unsafe step.

ctxt delWrapper *name* deletes the unsafe wrapper with the given name.

addSss adds the simpset of the context to its classical set. The assumptions and goal will be simplified, in a rather safe way, after each safe step of the search.

addss adds the simpset of the context to its classical set. The assumptions and goal will be simplified, before the each unsafe step of the search.

9.5 Object-logic setup

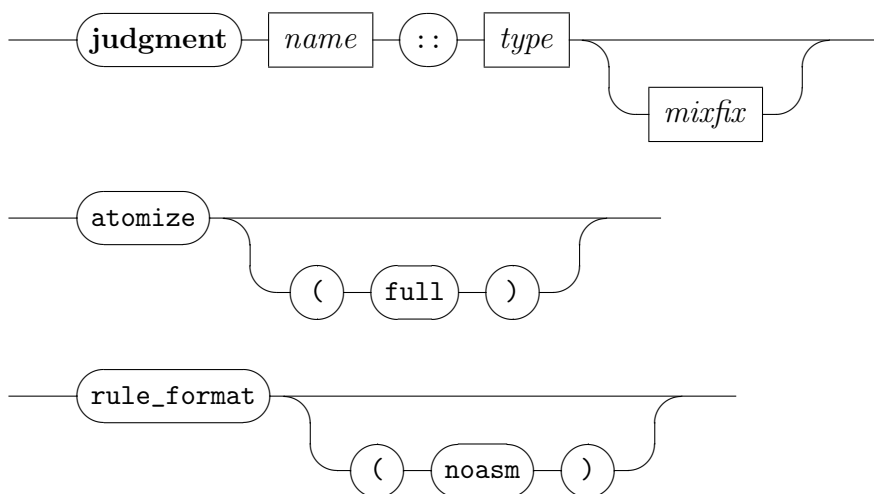
judgment : *theory* \rightarrow *theory*
atomize : *method*
atomize : *attribute*
rule_format : *attribute*
rulify : *attribute*

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 \wedge , \implies , and \equiv . Meta-level conjunction should be covered as well (this is particularly important for locales, see §5.7).

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 .

9.6 Tracing higher-order unification

```

unify_trace_simp : attribute default false
unify_trace_types : attribute default false
unify_trace_bound : attribute default 50
unify_search_bound : attribute default 60

```

Higher-order unification works well in most practical situations, but sometimes needs extra care to identify problems. These tracing options may help.

unify_trace_simp controls tracing of the simplification phase of higher-order unification.

unify_trace_types controls warnings of incompleteness, when unification is not considering all possible instantiations of schematic type variables.

unify_trace_bound determines the depth where unification starts to print tracing information once it reaches depth; 0 for full tracing. At the default value, tracing information is almost never printed in practice.

unify_search_bound prevents unification from searching past the given depth. Because of this bound, higher-order unification cannot return an infinite sequence, though it can return an exponentially long one. The search rarely approaches the default value in practice. If the search is cut off, unification prints a warning “Unification bound exceeded”.

- ! Options for unification cannot be modified in a local context. Only the global theory content is taken into account.

Part III

Isabelle/HOL

Higher-Order Logic

Isabelle/HOL is based on Higher-Order Logic, a polymorphic version of Church’s Simple Theory of Types. HOL can be best understood as a simply-typed version of classical set theory. The logic was first implemented in Gordon’s HOL system [17]. It extends Church’s original logic [14] by explicit type variables (naive polymorphism) and a sound axiomatization scheme for new types based on subsets of existing types.

Andrews’s book [1] is a full description of the original Church-style higher-order logic, with proofs of correctness and completeness wrt. certain set-theoretic interpretations. The particular extensions of Gordon-style HOL are explained semantically in two chapters of the 1993 HOL book [44].

Experience with HOL over decades has demonstrated that higher-order logic is widely applicable in many areas of mathematics and computer science. In a sense, Higher-Order Logic is simpler than First-Order Logic, because there are fewer restrictions and special cases. Note that HOL is *weaker* than FOL with axioms for ZF set theory, which is traditionally considered the standard foundation of regular mathematics, but for most applications this does not matter. If you prefer ML to Lisp, you will probably prefer HOL to ZF.

The syntax of HOL follows λ -calculus and functional programming. Function application is curried. To apply the function f of type $\tau_1 \Rightarrow \tau_2 \Rightarrow \tau_3$ to the arguments a and b in HOL, you simply write $f a b$ (as in ML or Haskell). There is no “apply” operator; the existing application of the Pure λ -calculus is re-used. Note that in HOL $f (a, b)$ means “ f applied to the pair (a, b) ” (which is notation for *Pair a b*). The latter typically introduces extra formal efforts that can be avoided by currying functions by default. Explicit tuples are as infrequent in HOL formalizations as in good ML or Haskell programs.

Isabelle/HOL has a distinct feel, compared to other object-logics like Isabelle/ZF. It identifies object-level types with meta-level types, taking advantage of the default type-inference mechanism of Isabelle/Pure. HOL fully identifies object-level functions with meta-level functions, with native abstraction and application.

These identifications allow Isabelle to support HOL particularly nicely, but they also mean that HOL requires some sophistication from the user. In particular, an understanding of Hindley-Milner type-inference with type-classes, which are both used extensively in the standard libraries and applications. Beginners can set *show_types* or even *show_sorts* to get more explicit information about the result of type-inference.

Derived specification elements

11.1 Inductive and coinductive definitions

```

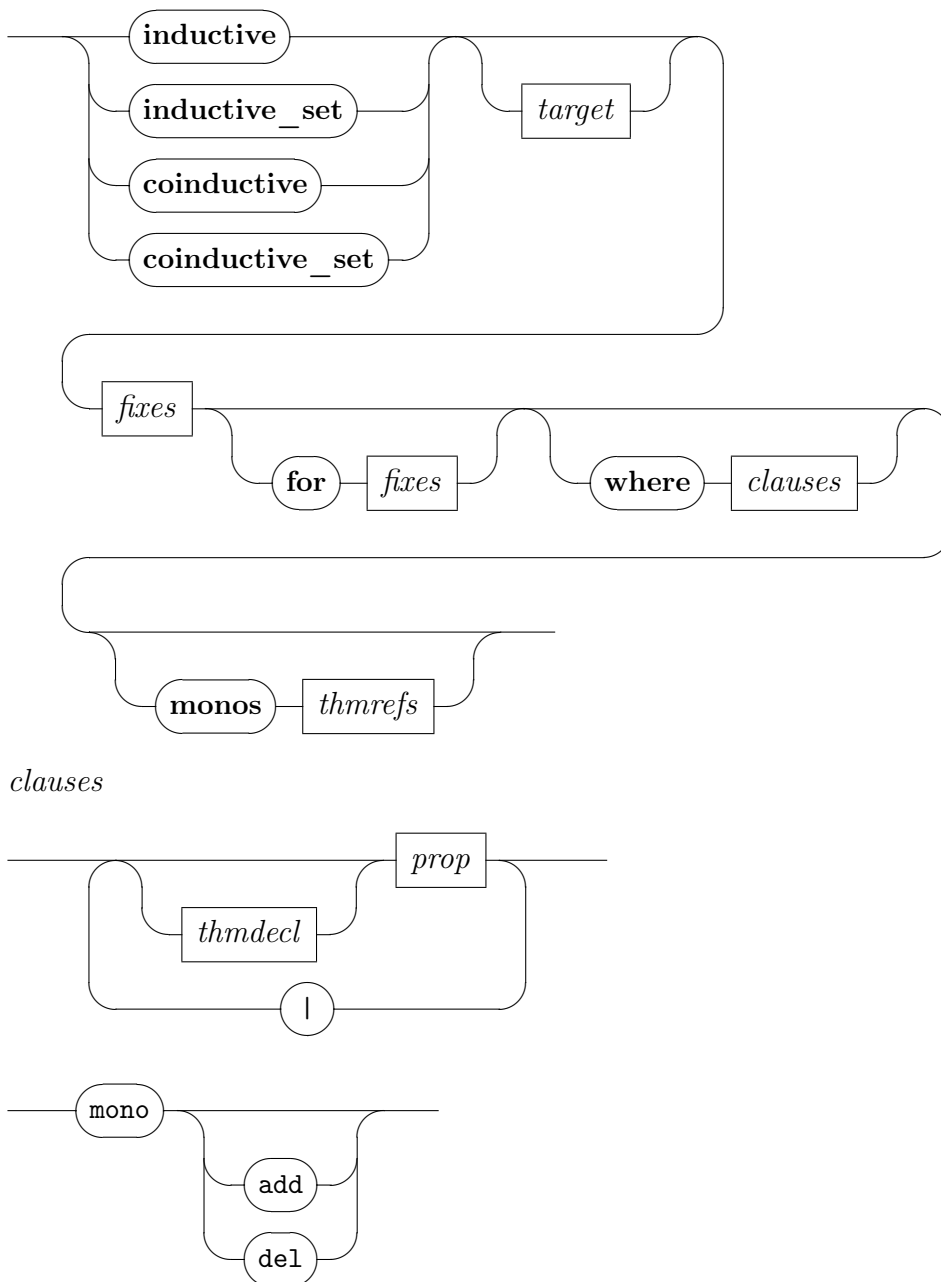
inductive      : local_theory → local_theory
inductive_set : local_theory → local_theory
coinductive   : local_theory → local_theory
coinductive_set : local_theory → local_theory
print_inductives* : context →
                    mono : attribute

```

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 or 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.

Both inductive and coinductive definitions are based on the Knaster-Tarski fixed-point theorem for complete lattices. The collection of introduction rules given by the user determines a functor on subsets of set-theoretic relations. The required monotonicity of the recursion scheme is proven as a prerequisite to the fixed-point definition and the resulting consequences. This works by pushing inclusion through logical connectives and any other operator that might be wrapped around recursive occurrences of the defined relation: there must be a monotonicity theorem of the form $A \leq B \implies \mathcal{M} A \leq \mathcal{M} B$, for each premise $\mathcal{M} R t$ in an introduction rule. The default rule declarations of Isabelle/HOL already take care of most common situations.



inductive and **coinductive** define (co)inductive predicates from the introduction rules.

The propositions given as *clauses* in the **where** part are either rules of the usual \bigwedge/\implies format (with arbitrary nesting), or equalities using \equiv . The latter specifies extra-logical abbreviations in the sense of **abbreviation**. Introducing abstract syntax simultaneously with the

actual introduction rules is occasionally useful for complex specifications.

The optional **for** part contains a list of parameters of the (co)inductive predicates that remain fixed throughout the definition, in contrast to arguments of the relation that may vary in each occurrence within the given *clauses*.

The optional **monos** declaration contains additional *monotonicity theorems*, which are required for each operator applied to a recursive set in the introduction rules.

inductive_set and **coinductive_set** are wrappers for to the previous commands for native HO \bar{L} predicates. This allows to define (co)inductive sets, where multiple arguments are simulated via tuples.

print_inductives prints (co)inductive definitions and monotonicity rules.

mono declares monotonicity rules in the context. These rule are involved in the automated monotonicity proof of the above inductive and coinductive definitions.

11.1.1 Derived rules

A (co)inductive definition of R provides the following main 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;

$R.simps$ is the equation unrolling the fixpoint of the predicate one step.

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 $R_1 \dots R_n.inducts$.

11.1.2 Monotonicity theorems

The context maintains a default set of theorems that are used in monotonicity proofs. New rules can be declared via the *mono* attribute. See the main Isabelle/HOL sources for some examples. The general format of such monotonicity theorems is as follows:

- Theorems of the form $A \leq B \implies \mathcal{M} A \leq \mathcal{M} B$, for proving monotonicity of inductive definitions whose introduction rules have premises involving terms such as $\mathcal{M} R t$.
- 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.

$$(P \longrightarrow Q) \longleftrightarrow \neg P \vee Q \qquad \text{Ball } A P \equiv \forall x. x \in A \longrightarrow P x$$

Examples

The finite powerset operator can be defined inductively like this:

inductive_set *Fin* :: 'a set \Rightarrow 'a set set **for** *A* :: 'a set
where

empty: $\{\} \in \text{Fin } A$
| *insert*: $a \in A \implies B \in \text{Fin } A \implies \text{insert } a B \in \text{Fin } A$

The accessible part of a relation is defined as follows:

inductive *acc* :: ('a \Rightarrow 'a \Rightarrow bool) \Rightarrow 'a \Rightarrow bool
for *r* :: 'a \Rightarrow 'a \Rightarrow bool (**infix** < 50)
where *acc*: $(\bigwedge y. y \prec x \implies \text{acc } r y) \implies \text{acc } r x$

Common logical connectives can be easily characterized as non-recursive inductive definitions with parameters, but without arguments.

```
inductive AND for A B :: bool
where A  $\implies$  B  $\implies$  AND A B
```

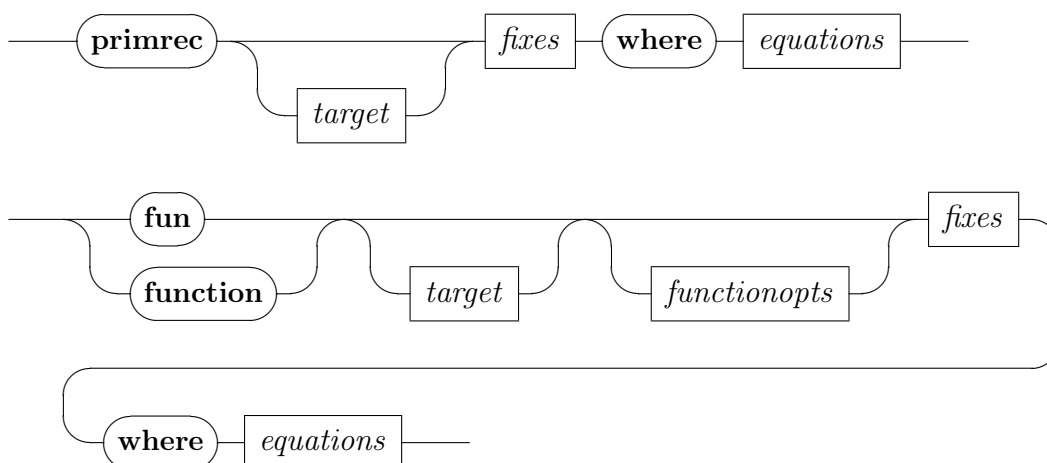
```
inductive OR for A B :: bool
where A  $\implies$  OR A B
  | B  $\implies$  OR A B
```

```
inductive EXISTS for B :: 'a  $\Rightarrow$  bool'
where B a  $\implies$  EXISTS B
```


Here the *cases* or *induct* rules produced by the **inductive** package coincide with the expected elimination rules for Natural Deduction. Already in the original article by Gerhard Gentzen [16] there is a hint that each connective can be characterized by its introductions, and the elimination can be constructed systematically.

11.2 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)
fun_cases : local_theory  $\rightarrow$  local_theory
```



equations



functionopts



primrec defines primitive recursive functions over datatypes (see also **datatype** and **rep_datatype**). The given *equations* specify reduction rules that are produced by instantiating the generic combinator for primitive recursion that is available for each datatype.

Each equation needs to be of the form:

$$f x_1 \dots x_m (C y_1 \dots y_k) z_1 \dots z_n = rhs$$

such that C is a datatype constructor, rhs contains only the free variables on the left-hand side (or from the context), and all recursive occurrences of f in rhs are of the form $f \dots y_i \dots$ for some i . At most one reduction rule for each constructor can be given. The order does

not matter. For missing constructors, the function is defined to return a default value, but this equation is made difficult to access for users.

The reduction rules are declared as *simp* by default, which enables standard proof methods like *simp* and *auto* to normalize expressions of *f* applied to datatype constructions, by simulating symbolic computation via rewriting.

function defines functions by general wellfounded recursion. A detailed description with examples can be found in [21]. 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 [21] 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.

fun_cases generates specialized elimination rules for function equations. It expects one or more function equations and produces rules that eliminate the given equalities, following the cases given in the function definition.

Recursive definitions introduced by the **function** command accommodate reasoning by induction (cf. *induct*): rule *f.induct* 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 where 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.

Example: evaluation of expressions

Subsequently, we define mutual datatypes for arithmetic and boolean expressions, and use **primrec** for evaluation functions that follow the same recursive structure.

```

datatype 'a aexp =
  | IF 'a bexp 'a aexp 'a aexp
  | Sum 'a aexp 'a aexp
  | Diff 'a aexp 'a aexp
  | Var 'a
  | Num nat
and 'a bexp =
  | Less 'a aexp 'a aexp
  | And 'a bexp 'a bexp
  | Neg 'a bexp

```

Evaluation of arithmetic and boolean expressions

```

primrec evala :: ('a ⇒ nat) ⇒ 'a aexp ⇒ nat
and evalb :: ('a ⇒ nat) ⇒ 'a bexp ⇒ bool
where
  evala env (IF b a1 a2) = (if evalb env b then evala env a1 else evala env a2)
| evala env (Sum a1 a2) = evala env a1 + evala env a2
| evala env (Diff a1 a2) = evala env a1 - evala env a2
| evala env (Var v) = env v
| evala env (Num n) = n
| evalb env (Less a1 a2) = (evala env a1 < evala env a2)
| evalb env (And b1 b2) = (evalb env b1 ∧ evalb env b2)
| evalb env (Neg b) = (¬ evalb env b)

```

Since the value of an expression depends on the value of its variables, the functions *evala* and *evalb* take an additional parameter, an *environment* that maps variables to their values.

Substitution on expressions can be defined similarly. The mapping *f* of type $'a \Rightarrow 'a \text{ aexp}$ given as a parameter is lifted canonically on the types $'a \text{ aexp}$ and $'a \text{ bexp}$, respectively.

primrec *subst* :: $('a \Rightarrow 'b \text{ aexp}) \Rightarrow 'a \text{ aexp} \Rightarrow 'b \text{ aexp}$
and *substb* :: $('a \Rightarrow 'b \text{ aexp}) \Rightarrow 'a \text{ bexp} \Rightarrow 'b \text{ bexp}$

where

subst *f* (*IF* *b* *a1* *a2*) = *IF* (*substb* *f* *b*) (*subst* *f* *a1*) (*subst* *f* *a2*)
| *subst* *f* (*Sum* *a1* *a2*) = *Sum* (*subst* *f* *a1*) (*subst* *f* *a2*)
| *subst* *f* (*Diff* *a1* *a2*) = *Diff* (*subst* *f* *a1*) (*subst* *f* *a2*)
| *subst* *f* (*Var* *v*) = *f* *v*
| *subst* *f* (*Num* *n*) = *Num* *n*
| *substb* *f* (*Less* *a1* *a2*) = *Less* (*subst* *f* *a1*) (*subst* *f* *a2*)
| *substb* *f* (*And* *b1* *b2*) = *And* (*substb* *f* *b1*) (*substb* *f* *b2*)
| *substb* *f* (*Neg* *b*) = *Neg* (*substb* *f* *b*)

In textbooks about semantics one often finds substitution theorems, which express the relationship between substitution and evaluation. For $'a \text{ aexp}$ and $'a \text{ bexp}$, we can prove such a theorem by mutual induction, followed by simplification.

lemma *subst_one*:

evala *env* (*subst* (*Var* (*v* := *a'*)) *a*) = *evala* (*env* (*v* := *evala* *env* *a'*)) *a*
evalb *env* (*substb* (*Var* (*v* := *a'*)) *b*) = *evalb* (*env* (*v* := *evala* *env* *a'*)) *b*
by (*induct* *a* **and** *b*) *simp_all*

lemma *subst_all*:

evala *env* (*subst* *s* *a*) = *evala* ($\lambda x. \text{evala } \text{env } (s \ x)$) *a*
evalb *env* (*substb* *s* *b*) = *evalb* ($\lambda x. \text{evala } \text{env } (s \ x)$) *b*
by (*induct* *a* **and** *b*) *simp_all*

Example: a substitution function for terms

Functions on datatypes with nested recursion are also defined by mutual primitive recursion.

datatype $('a, 'b) \text{ term} = \text{Var } 'a \mid \text{App } 'b ('a, 'b) \text{ term list}$

A substitution function on type $('a, 'b) \text{ term}$ can be defined as follows, by working simultaneously on $('a, 'b) \text{ term list}$:

primrec *subst_term* :: $('a \Rightarrow ('a, 'b) \text{ term}) \Rightarrow ('a, 'b) \text{ term} \Rightarrow ('a, 'b) \text{ term}$ **and**

```

subst_term_list :: ('a ⇒ ('a, 'b) term) ⇒ ('a, 'b) term list ⇒ ('a, 'b) term list
where
  subst_term f (Var a) = f a
| subst_term f (App b ts) = App b (subst_term_list f ts)
| subst_term_list f [] = []
| subst_term_list f (t # ts) = subst_term f t # subst_term_list f ts

```

The recursion scheme follows the structure of the unfolded definition of type $(\text{'a}, \text{'b}) \text{ term}$. To prove properties of this substitution function, mutual induction is needed:

```

lemma subst_term (subst_term f1 ◦ f2) t = subst_term f1 (subst_term f2 t)
and
  subst_term_list (subst_term f1 ◦ f2) ts = subst_term_list f1 (subst_term_list
f2 ts)
  by (induct t and ts) simp_all

```

Example: a map function for infinitely branching trees

Defining functions on infinitely branching datatypes by primitive recursion is just as easy.

```

datatype 'a tree = Atom 'a | Branch nat ⇒ 'a tree

```

```

primrec map_tree :: ('a ⇒ 'b) ⇒ 'a tree ⇒ 'b tree
where
  map_tree f (Atom a) = Atom (f a)
| map_tree f (Branch ts) = Branch (λx. map_tree f (ts x))

```

Note that all occurrences of functions such as ts above must be applied to an argument. In particular, $map_tree\ f \circ ts$ is not allowed here.

Here is a simple composition lemma for map_tree :

```

lemma map_tree g (map_tree f t) = map_tree (g ◦ f) t
  by (induct t) simp_all

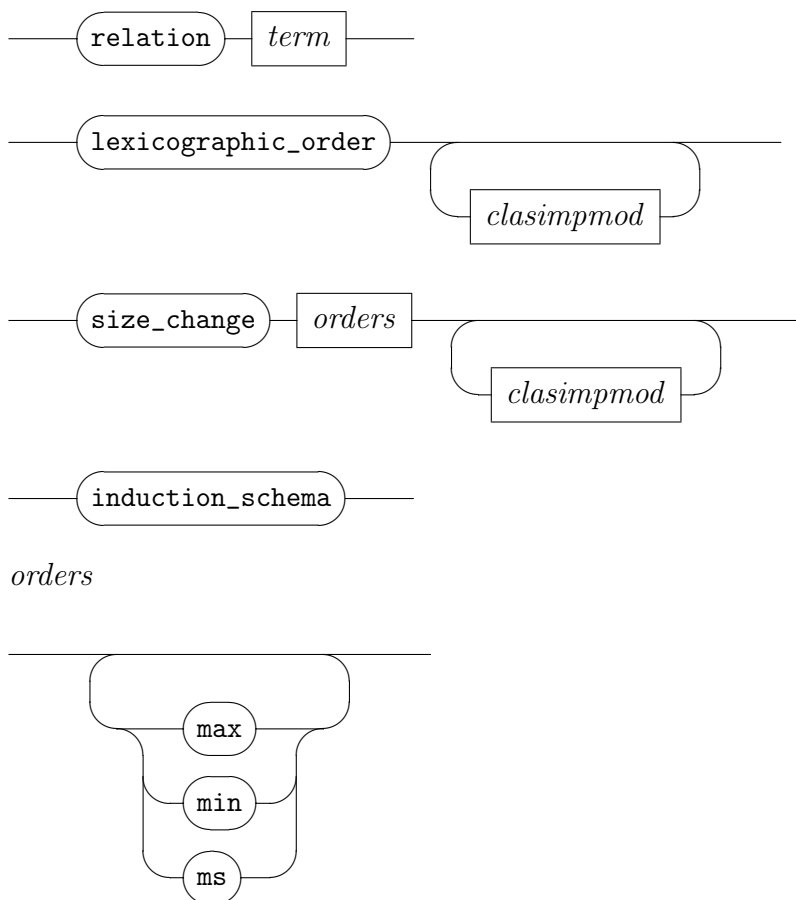
```

11.2.1 Proof methods related to recursive definitions

```

  pat_completeness : method
  relation : method
  lexicographic_order : method
  size_change : method
  induction_schema : method

```



pat_completeness is a specialized method to solve goals regarding the completeness of pattern matching, as required by the **function** package (cf. [21]).

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 *clasimpmod* modifiers are accepted (as for *auto*).

In case of failure, extensive information is printed, which can help to analyse the situation (cf. [21]).


size_change also works on termination goals, using a variation of the size-change principle, together with a graph decomposition technique (see [22] 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 *clasimpmod* modifiers are accepted (as for *auto*).

induction_schema derives user-specified induction rules from well-founded induction and completeness of patterns. This factors out some operations that are done internally by the function package and makes them available separately. See `~/src/HOL/ex/Induction_Schema.thy` for examples.

11.2.2 Functions with explicit partiality

partial_function : *local_theory* \rightarrow *local_theory*
partial_function_mono : *attribute*



partial_function (*mode*) 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 monotonicity proofs of partial function definitions.

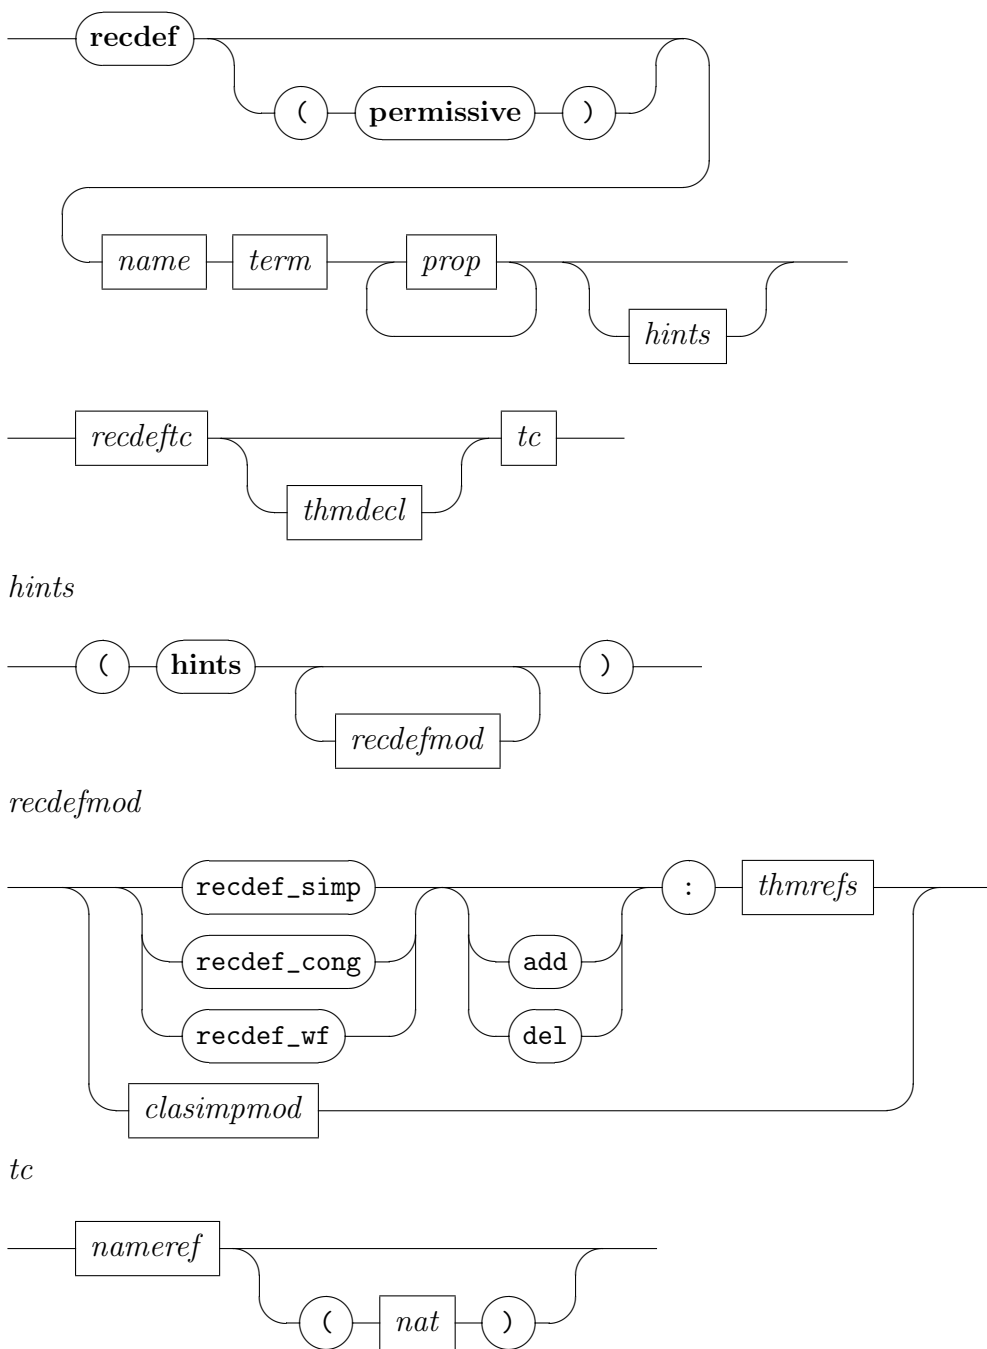
11.2.3 Old-style recursive function definitions (TFL)

```

recdef   : theory → theory)
recdef_tc* : theory → proof(prove)

```

The old TFL commands **recdef** and **recdef_tc** for defining recursive are mostly obsolete; **function** or **fun** should be used instead.



recdef defines general well-founded recursive functions (using the TFL package), see also [31]. 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 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

```





11.3 Datatypes

```

datatype   : theory → theory
rep_datatype : theory → proof(prove)

```



spec*cons*

datatype defines inductive datatypes in HOL.

rep_datatype represents existing types as datatypes.

For foundational reasons, some basic types such as *nat*, $'a \times 'b$, $'a + 'b$, *bool* and *unit* are introduced by more primitive means using **typedef**. To recover the rich infrastructure of **datatype** (e.g. rules for *cases* and *induct* and the primitive recursion combinators), such types may be represented as actual datatypes later. This is done by specifying the constructors of the desired type, and giving a proof of the induction rule, distinctness and injectivity of constructors.

For example, see `~/src/HOL/Sum_Type.thy` for the representation of the primitive sum type as fully-featured datatype.

The generated rules for *induct* and *cases* provide case names according to the given constructors, while parameters are named after the types (see also §6.6).

See [31] 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 §12.12; these admit to refer directly to the internal structure of subgoals (including internally bound parameters).

Examples

We define a type of finite sequences, with slightly different names than the existing *'a list* that is already in *Main*:

```
datatype 'a seq = Empty | Seq 'a 'a seq
```

We can now prove some simple lemma by structural induction:

```

lemma Seq x xs ≠ xs
proof (induct xs arbitrary: x)
  case Empty

```

This case can be proved using the simplifier: the freeness properties of the datatype are already declared as *simp* rules.

```

  show Seq x Empty ≠ Empty
    by simp
next
  case (Seq y ys)

```

The step case is proved similarly.

```

  show Seq x (Seq y ys) ≠ Seq y ys
    using 'Seq y ys ≠ ys' by simp
qed

```

Here is a more succinct version of the same proof:

```

lemma Seq x xs ≠ xs
  by (induct xs arbitrary: x) simp_all

```

11.4 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 [28] for more details on object-oriented verification and record subtyping in HOL.

11.4.1 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	$(x = a, y = b)$	$(x :: A, y :: B)$
schematic	$(x = a, y = b, \dots = m)$	$(x :: A, y :: B, \dots :: M)$

The ASCII representation of $(x = a)$ 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 $() :: \textit{unit}$ element. In fact, $\langle x = a, y = b \rangle$ is just an abbreviation for $\langle x = a, y = b, \dots = () \rangle$.

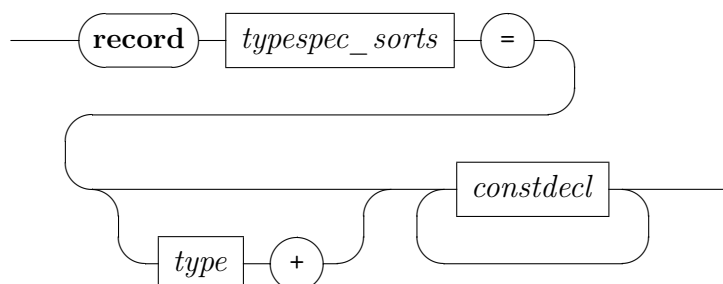
Two key observations make extensible records in a simply typed language like HOL work out:

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 [32] for further instructions on using records in practice.

11.4.2 Record specifications

record : *theory* \rightarrow *theory*



constdecl



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 $(c_1 :: \sigma_1, \dots, c_n :: \sigma_n)$, likewise is $(\alpha_1, \dots, \alpha_m, \zeta) t_scheme$ made an abbreviation for $(c_1 :: \sigma_1, \dots, c_n :: \sigma_n, \dots :: \zeta)$.

11.4.3 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*”):

$$\begin{aligned} c_i &:: (\bar{c} :: \bar{\sigma}, \dots :: \zeta) \Rightarrow \sigma_i \\ c_i_update &:: \sigma_i \Rightarrow (\bar{c} :: \bar{\sigma}, \dots :: \zeta) \Rightarrow (\bar{c} :: \bar{\sigma}, \dots :: \zeta) \end{aligned}$$

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:

$$t.make \quad :: \quad \sigma_1 \Rightarrow \dots \sigma_n \Rightarrow (\bar{c} :: \bar{\sigma})$$

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:

$$\begin{aligned} c_i & \quad :: \quad (\bar{b} :: \bar{\varrho}, \bar{c} :: \bar{\sigma}, \dots :: \zeta) \Rightarrow \sigma_i \\ c_i_update & \quad :: \quad \sigma_i \Rightarrow (\bar{b} :: \bar{\varrho}, \bar{c} :: \bar{\sigma}, \dots :: \zeta) \Rightarrow (\bar{b} :: \bar{\varrho}, \bar{c} :: \bar{\sigma}, \dots :: \zeta) \\ t.make & \quad :: \quad \varrho_1 \Rightarrow \dots \varrho_k \Rightarrow \sigma_1 \Rightarrow \dots \sigma_n \Rightarrow (\bar{b} :: \bar{\varrho}, \bar{c} :: \bar{\sigma}) \end{aligned}$$

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.

$$\begin{aligned} t.fields & \quad :: \quad \sigma_1 \Rightarrow \dots \sigma_n \Rightarrow (\bar{c} :: \bar{\sigma}) \\ t.extend & \quad :: \quad (\bar{b} :: \bar{\varrho}, \bar{c} :: \bar{\sigma}) \Rightarrow \zeta \Rightarrow (\bar{b} :: \bar{\varrho}, \bar{c} :: \bar{\sigma}, \dots :: \zeta) \\ t.truncate & \quad :: \quad (\bar{b} :: \bar{\varrho}, \bar{c} :: \bar{\sigma}, \dots :: \zeta) \Rightarrow (\bar{b} :: \bar{\varrho}, \bar{c} :: \bar{\sigma}) \end{aligned}$$

Note that $t.make$ and $t.fields$ coincide for root records.

11.4.4 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*.

Examples

See `~/src/HOL/ex/Records.thy`, for example.

11.5 Typedef axiomatization

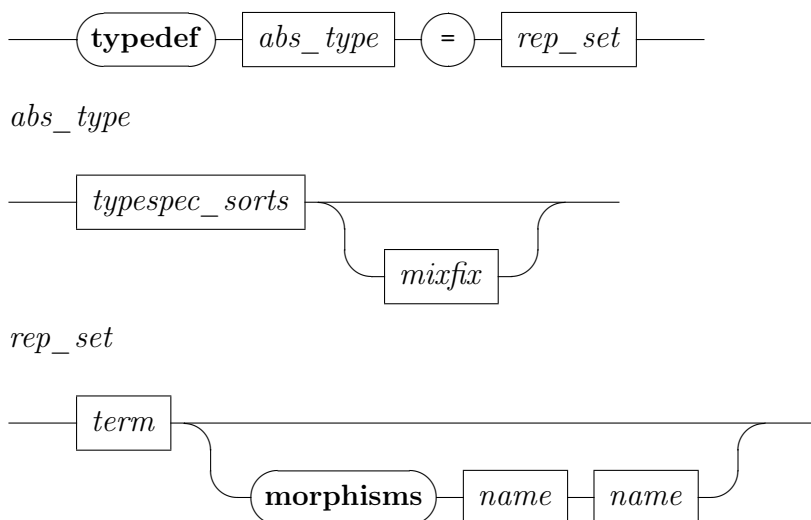
`typedef` : *local_theory* \rightarrow *proof*(*prove*)

A Gordon/HOL-style type definition is a certain axiom scheme that identifies a new type with a subset of an existing type. More precisely, the new type is defined by exhibiting an existing type τ , a set $A :: \tau$ *set*, and a theorem that proves $\exists x. x \in A$. Thus A is a non-empty subset of τ , and the new type denotes this subset. New functions are postulated that establish an isomorphism between the new type and the subset. In general, the type

τ may involve type variables $\alpha_1, \dots, \alpha_n$ which means that the type definition produces a type constructor $(\alpha_1, \dots, \alpha_n) t$ depending on those type arguments.

The axiomatization can be considered a “definition” in the sense of the particular set-theoretic interpretation of HOL [44], where the universe of types is required to be downwards-closed wrt. arbitrary non-empty subsets. Thus genuinely new types introduced by **typedef** stay within the range of HOL models by construction.

In contrast, the command **type_synonym** from Isabelle/Pure merely introduces syntactic abbreviations, without any logical significance. Thus it is more faithful to the idea of a genuine type definition, but less powerful in practice.



typedef $(\alpha_1, \dots, \alpha_n) t = A$ produces an axiomatization (§5.5) for a type definition in the background theory of the current context, depending on a non-emptiness result of the set A that needs to be proven here. The set A may contain type variables $\alpha_1, \dots, \alpha_n$ as specified on the LHS, but no term variables.

Even though a local theory specification, the newly introduced type constructor cannot depend on parameters or assumptions of the context: this is structurally impossible in HOL. In contrast, the non-emptiness proof may use local assumptions in unusual situations, which could result in different interpretations in target contexts: the meaning of the bijection between the representing set A and the new type t may then change in different application contexts.

For **typedef** $t = A$ the newly introduced type t is accompanied by a pair of morphisms to relate it to the representing set over the old type. By default, the injection from type to set is called Rep_t and its inverse Abs_t : An explicit **morphisms** specification allows to provide alternative names.

The core axiomatization uses the locale predicate $type_definition$ as defined in Isabelle/HOL. Various basic consequences of that are instantiated accordingly, re-using the locale facts with names derived from the new type constructor. Thus the generic $type_definition.Rep$ is turned into the specific Rep_t , for example.

Theorems $type_definition.Rep$, $type_definition.Rep_inverse$, and $type_definition.Abs_inverse$ provide the most basic characterization as a corresponding injection/surjection pair (in both directions). The derived rules $type_definition.Rep_inject$ and $type_definition.Abs_inject$ provide a more convenient version of injectivity, suitable for automated proof tools (e.g. in declarations involving *simp* or *iff*). Furthermore, the rules $type_definition.Rep_cases$ / $type_definition.Rep_induct$, and $type_definition.Abs_cases$ / $type_definition.Abs_induct$ provide alternative views on surjectivity. These rules are already declared as set or type rules for the generic *cases* and *induct* methods, respectively.

Examples

Type definitions permit the introduction of abstract data types in a safe way, namely by providing models based on already existing types. Given some abstract axiomatic description P of a type, this involves two steps:

1. Find an appropriate type τ and subset A which has the desired properties P , and make a type definition based on this representation.
2. Prove that P holds for τ by lifting P from the representation.

You can later forget about the representation and work solely in terms of the abstract properties P .

The following trivial example pulls a three-element type into existence within the formal logical environment of HOL.

```
typedef three = {(True, True), (True, False), (False, True)}
by blast
```

definition $One = Abs_three (True, True)$

definition $Two = Abs_three (True, False)$

definition $Three = Abs_three (False, True)$

lemma $three_distinct: One \neq Two \ One \neq Three \ Two \neq Three$

by ($simp_all \ add: One_def \ Two_def \ Three_def \ Abs_three_inject$)

lemma $three_cases:$

fixes $x :: three$ **obtains** $x = One \mid x = Two \mid x = Three$

by ($cases \ x$) ($auto \ simp: One_def \ Two_def \ Three_def \ Abs_three_inject$)


Note that such trivial constructions are better done with derived specification mechanisms such as **datatype**:

datatype $three' = One' \mid Two' \mid Three'$

This avoids re-doing basic definitions and proofs from the primitive **typedef** above.

11.6 Functorial structure of types

functor : $local_theory \rightarrow proof(prove)$



functor *prefix*: m 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 $\alpha_1 \Rightarrow \beta_1, \beta_1 \Rightarrow \alpha_1, \dots, \alpha_n \Rightarrow \beta_n, \beta_n \Rightarrow \alpha_n$.


11.7 Quotient types

```

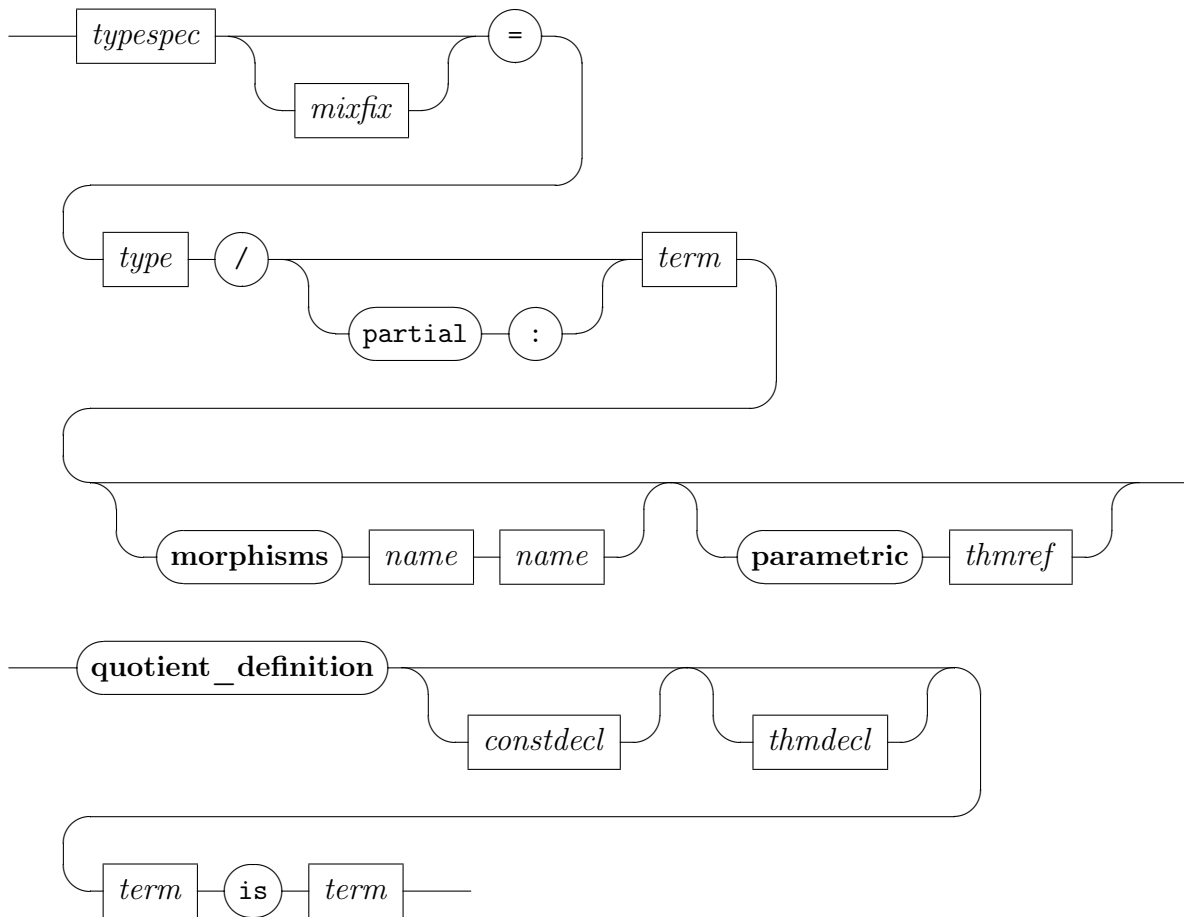
quotient_type : local_theory → proof(prove)
quotient_definition : local_theory → proof(prove)
print_quotmapsQ3 : context →
print_quotientsQ3 : context →
print_quotconsts : context →
    lifting : method
    lifting_setup : method
    descending : method
    descending_setup : method
    partiality_descending : method
partiality_descending_setup : method
    regularize : method
    injection : method
    cleaning : method
    quot_thm : attribute
    quot_lifted : attribute
    quot_respect : attribute
    quot_preserve : attribute

```

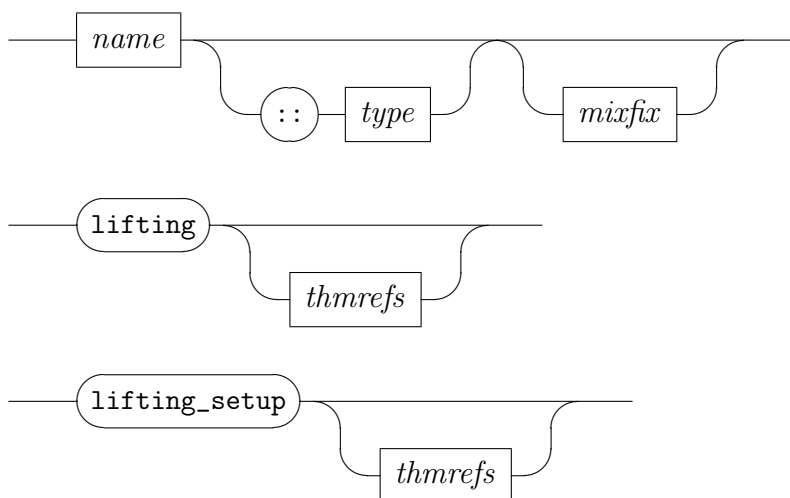
The quotient package defines a new quotient type given a raw type and a partial equivalence relation. The package also historically includes automation for transporting definitions and theorems. But most of this automation was superseded by the Lifting and Transfer packages. The user should consider using these two new packages for lifting definitions and transporting theorems.



spec



constdecl



quotient_type defines a new quotient type τ . The injection from a quotient type to a raw type is called *rep_* τ , its inverse *abs_* τ unless explicit **morphisms** specification provides alternative names. **quotient_type** requires the user to prove that the relation is an equivalence relation (predicate *equivp*), unless the user specifies explicitly *partial* in which case the obligation is *part_equivp*. A quotient defined with *partial* is weaker in the sense that less things can be proved automatically.

The command internally proves a Quotient theorem and sets up the Lifting package by the command **setup_lifting**. Thus the Lifting and Transfer packages can be used also with quotient types defined by **quotient_type** without any extra set-up. The parametricity theorem for the equivalence relation R can be provided as an extra argument of the command and is passed to the corresponding internal call of **setup_lifting**. This theorem allows the Lifting package to generate a stronger transfer rule for equality.

The most of the rest of the package was superseded by the Lifting and Transfer packages. The user should consider using these two new packages for lifting definitions and transporting theorems.

quotient_definition defines a constant on the quotient type.

print_quotmapsQ3 prints quotient map functions.

print_quotientsQ3 prints quotients.

print_quotconsts prints quotient constants.

lifting and *lifting_setup* methods match the current goal with the given raw theorem to be lifted producing three new subgoals: regularization, injection and cleaning subgoals. *lifting* tries to apply the heuristics for automatically solving these three subgoals and leaves only the subgoals unsolved by the heuristics to the user as opposed to *lifting_setup* which leaves the three subgoals unsolved.

descending and *descending_setup* try to guess a raw statement that would lift to the current subgoal. Such statement is assumed as a new subgoal and *descending* continues in the same way as *lifting* does. *descending* tries to solve the arising regularization, injection and cleaning subgoals with the analogous method *descending_setup* which leaves the four unsolved subgoals.

partiality_descending finds the regularized theorem that would lift to the current subgoal, lifts it and leaves as a subgoal. This method can be used with partial equivalence quotients where the non regularized statements would not be true. *partiality_descending_setup* leaves the injection and cleaning subgoals unchanged.

regularize applies the regularization heuristics to the current subgoal.

injection applies the injection heuristics to the current goal using the stored quotient respectfulness theorems.

cleaning applies the injection cleaning heuristics to the current subgoal using the stored quotient preservation theorems.


quot_lifted attribute tries to automatically transport the theorem to the quotient type. The attribute uses all the defined quotients types and quotient constants often producing undesired results or theorems that cannot be lifted.

quot_respect and *quot_preserve* attributes declare a theorem as a respectfulness and preservation theorem respectively. These are stored in the local theory store and used by the *injection* and *cleaning* methods respectively.


quot_thm declares that a certain theorem is a quotient extension theorem. Quotient extension theorems allow for quotienting inside container types. Given a polymorphic type that serves as a container, a map function defined for this container using **functor** and a relation map defined for the container type, the quotient extension theorem should be $Quotient3\ R\ Abs\ Rep \implies Quotient3\ (rel_map\ R)\ (map\ Abs)\ (map\ Rep)$. Quotient extension theorems are stored in a database and are used all the steps of lifting theorems.

11.8 Definition by specification

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.

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.


11.9 Adhoc overloading of constants

```

adhoc_overloading  : local_theory → local_theory
no_adhoc_overloading : local_theory → local_theory
      show_variants : attribute default false

```

Adhoc overloading allows to overload a constant depending on its type. Typically this involves the introduction of an uninterpreted constant (used for input and output) and the addition of some variants (used internally). For examples see `~/src/HOL/ex/Adhoc_Overloading_Examples.thy` and `~/src/HOL/Library/Monad_Syntax.thy`.



adhoc_overloading $c v_1 \dots v_n$ associates variants with an existing constant.


no_adhoc_overloading is similar to **adhoc_overloading**, but removes the specified variants from the present context.

show_variants controls printing of variants of overloaded constants. If enabled, the internally used variants are printed instead of their respective overloaded constants. This is occasionally useful to check whether the system agrees with a user's expectations about derived variants.

Proof tools

12.1 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.

12.2 Transfer package

transfer : *method*
transfer' : *method*
transfer_prover : *method*
Transfer.transferred : *attribute*
untransferred : *attribute*
transfer_rule : *attribute*
transfer_domain_rule : *attribute*
relator_eq : *attribute*
relator_domain : *attribute*

transfer method replaces the current subgoal with a logically equivalent one that uses different types and constants. The replacement of types and constants is guided by the database of transfer rules. Goals are generalized over all free variables by default; this is necessary for variables whose types change, but can be overridden for specific variables with e.g. *transfer fixing: x y z*.

transfer' is a variant of *transfer* that allows replacing a subgoal with one that is logically stronger (rather than equivalent). For example, a subgoal involving equality on a quotient type could be replaced with a subgoal involving equality (instead of the corresponding equivalence relation) on the underlying raw type.

transfer_prover method assists with proving a transfer rule for a new constant, provided the constant is defined in terms of other constants that already have transfer rules. It should be applied after unfolding the constant definitions.

untransferred proves the same equivalent theorem as *transfer* internally does.

Transfer.transferred works in the opposite direction than *transfer'*. E.g., given the transfer relation $ZN\ x\ n \equiv (x = int\ n)$, corresponding transfer rules and the theorem $\forall x::int \in \{0..\}. x < x + 1$, the attribute would prove $\forall n::nat. n < n + 1$. The attribute is still in experimental phase of development.

transfer_rule attribute maintains a collection of transfer rules, which relate constants at two different types. Typical transfer rules may relate different type instances of the same polymorphic constant, or they may relate an operation on a raw type to a corresponding operation on an abstract type (quotient or subtype). For example:

$$\begin{aligned} & ((A ==> B) ==> list_all2\ A ==> list_all2\ B)\ map\ map \\ & (cr_int ==> cr_int ==> cr_int)\ (\lambda(x,y)\ (u,v).\ (x+u,\ y+v)) \\ & plus \end{aligned}$$

Lemmas involving predicates on relations can also be registered using the same attribute. For example:

$$\begin{aligned} & bi_unique\ A \implies (list_all2\ A ==> op\ =)\ distinct\ distinct \\ & \llbracket bi_unique\ A; bi_unique\ B \rrbracket \implies bi_unique\ (rel_prod\ A\ B) \end{aligned}$$

Preservation of predicates on relations (*bi_unique*, *bi_total*, *right_unique*, *right_total*, *left_unique*, *left_total*) with the respect to a relator is proved automatically if the involved type is BNF[10] without dead variables.

transfer_domain_rule attribute maintains a collection of rules, which specify a domain of a transfer relation by a predicate. E.g., given the transfer relation $ZN\ x\ n \equiv (x = int\ n)$, one can register the following transfer domain rule: *Domainp* $ZN = (\lambda x. x \geq 0)$. The rules allow the package

to produce more readable transferred goals, e.g., when quantifiers are transferred.

relator_eq attribute collects identity laws for relators of various type constructors, e.g. $rel_set (op =) = (op =)$. The *transfer* method uses these lemmas to infer transfer rules for non-polymorphic constants on the fly. For examples see `~/src/HOL/Lifting_Set.thy` or `~/src/HOL/Lifting.thy`. This property is proved automatically if the involved type is BNF without dead variables.

relator_domain attribute collects rules describing domains of relators by predicates. E.g., $Domainp (rel_set T) = (\lambda A. Ball A (Domainp T))$. This allows the package to lift transfer domain rules through type constructors. For examples see `~/src/HOL/Lifting_Set.thy` or `~/src/HOL/Lifting.thy`. This property is proved automatically if the involved type is BNF without dead variables.

Theoretical background can be found in [20].

12.3 Lifting package

The Lifting package allows users to lift terms of the raw type to the abstract type, which is a necessary step in building a library for an abstract type. Lifting defines a new constant by combining coercion functions (`Abs` and `Rep`) with the raw term. It also proves an appropriate transfer rule for the Transfer package and, if possible, an equation for the code generator.

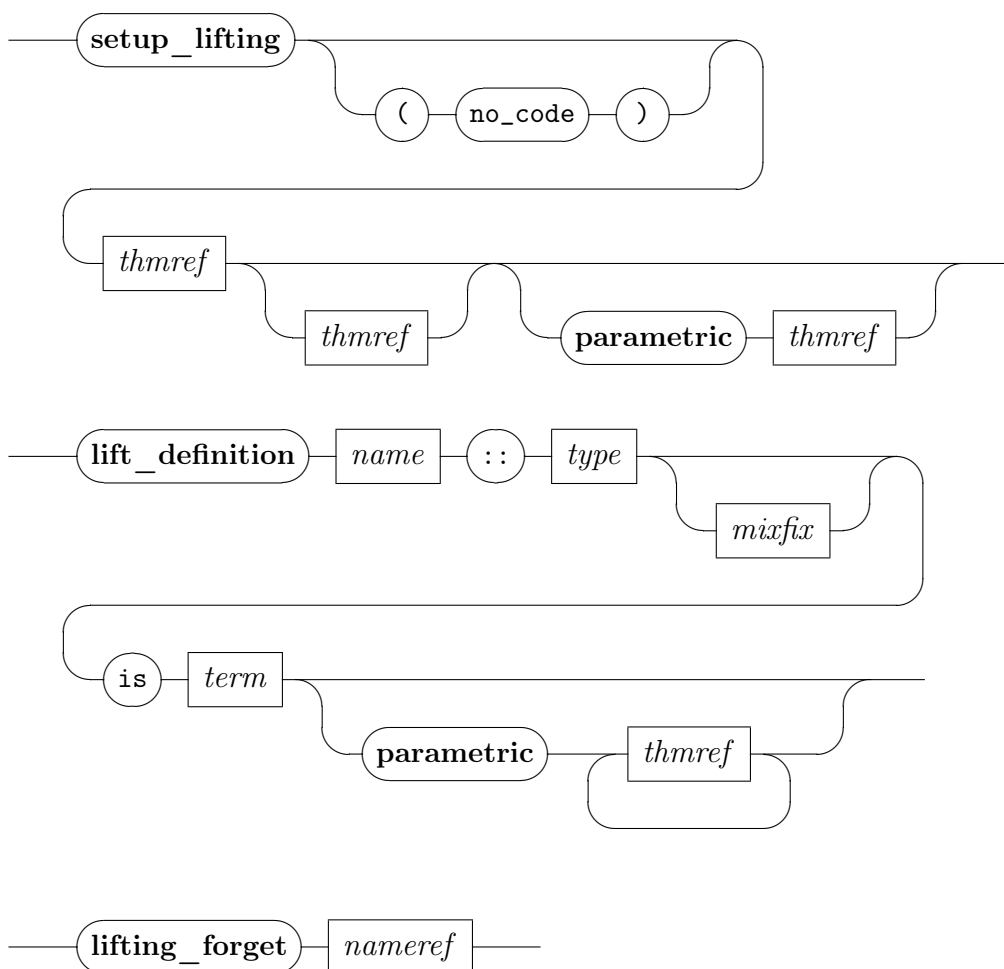
The Lifting package provides two main commands: `setup_lifting` for initializing the package to work with a new type, and `lift_definition` for lifting constants. The Lifting package works with all four kinds of type abstraction: type copies, subtypes, total quotients and partial quotients.


Theoretical background can be found in [20].

```

setup_lifting : local_theory → local_theory
lift_definition : local_theory → proof(prove)
lifting_forget : local_theory → local_theory
lifting_update : local_theory → local_theory
print_quot_maps : context →
print_quotients : context →
    quot_map : attribute
    relator_eq_onp : attribute
    relator_mono : attribute
    relator_distr : attribute
    quot_del : attribute
    lifting_restore : attribute

```





setup_lifting Sets up the Lifting package to work with a user-defined type. The command supports two modes. The first one is a low-level mode when the user must provide as a first argument of **setup_lifting** a quotient theorem *Quotient R Abs Rep T*. The package configures a transfer rule for equality, a domain transfer rules and sets up the **lift_definition** command to work with the abstract type. An optional theorem *reflp R*, which certifies that the equivalence relation R is total, can be provided as a second argument. This allows the package to generate stronger transfer rules. And finally, the parametricity theorem for R can be provided as a third argument. This allows the package to generate a stronger transfer rule for equality.

Users generally will not prove the *Quotient* theorem manually for new types, as special commands exist to automate the process.

When a new subtype is defined by **typedef**, **lift_definition** can be used in its second mode, where only the *type_definition Rep Abs A* theorem is used as an argument of the command. The command internally proves the corresponding Quotient theorem and registers it with **setup_lifting** using its first mode.

For quotients, the command **quotient_type** can be used. The command defines a new quotient type and similarly to the previous case, the corresponding Quotient theorem is proved and registered by **setup_lifting**.

The command **setup_lifting** also sets up the code generator for the new type. Later on, when a new constant is defined by **lift_definition**, the Lifting package proves and registers a code equation (if there is one) for the new constant. If the option *no_code* is specified, the Lifting package does not set up the code generator and as a consequence no code equations involving an abstract type are registered by **lift_definition**.

lift_definition $f :: \tau$ **is** t Defines a new function f with an abstract type τ in terms of a corresponding operation t on a representation type. More

formally, if $t :: \sigma$, then the command builds a term F as a corresponding combination of abstraction and representation functions such that $F :: \sigma \Rightarrow \tau$ and defines f is as $f \equiv F t$. The term t does not have to be necessarily a constant but it can be any term.

The command opens a proof environment and the user must discharge a respectfulness proof obligation. For a type copy, i.e., a typedef with *UNIV*, the obligation is discharged automatically. The proof goal is presented in a user-friendly, readable form. A respectfulness theorem in the standard format $f.rsp$ and a transfer rule $f.transfer$ for the Transfer package are generated by the package.

The user can specify a parametricity theorems for t after the keyword **parametric**, which allows the command to generate parametric transfer rules for f .

For each constant defined through trivial quotients (type copies or subtypes) $f.rep_eq$ is generated. The equation is a code certificate that defines f using the representation function.

For each constant $f.abs_eq$ is generated. The equation is unconditional for total quotients. The equation defines f using the abstraction function.

Integration with [*code abstract*]: For subtypes (e.g., corresponding to a datatype invariant, such as *dlist*), **lift_definition** uses a code certificate theorem $f.rep_eq$ as a code equation.

Integration with [*code equation*]: For total quotients, **lift_definition** uses $f.abs_eq$ as a code equation.

lifting_forget and **lifting_update** These two commands serve for storing and deleting the set-up of the Lifting package and corresponding transfer rules defined by this package. This is useful for hiding of type construction details of an abstract type when the construction is finished but it still allows additions to this construction when this is later necessary.

Whenever the Lifting package is set up with a new abstract type τ by **lift_definition**, the package defines a new bundle that is called $\tau.lifting$. This bundle already includes set-up for the Lifting package. The new transfer rules introduced by **lift_definition** can be stored in the bundle by the command **lifting_update** $\tau.lifting$.

The command **lifting_forget** $\tau.lifting$ deletes set-up of the Lifting package for τ and deletes all the transfer rules that were introduced by **lift_definition** using τ as an abstract type.

The stored set-up in a bundle can be reintroduced by the Isar commands for including a bundle (**include**, **includes** and **including**).

print _quot _maps prints stored quotient map theorems.

print _quotients prints stored quotient theorems.

quot_map registers a quotient map theorem, a theorem showing how to "lift" quotients over type constructors. E.g., $Quotient\ R\ Abs\ Rep\ T \implies Quotient\ (rel_set\ R)\ (image\ Abs)\ (image\ Rep)\ (rel_set\ T)$. For examples see `~/src/HOL/Lifting_Set.thy` or `~/src/HOL/Lifting.thy`. This property is proved automatically if the involved type is BNF without dead variables.

relator_eq_onp registers a theorem that shows that a relator applied to an equality restricted by a predicate P (i.e., $eq_onp\ P$) is equal to a predicator applied to the P . The combinator *eq_onp* is used for internal encoding of proper subtypes. Such theorems allows the package to hide *eq_onp* from a user in a user-readable form of a respectfulness theorem. For examples see `~/src/HOL/Lifting_Set.thy` or `~/src/HOL/Lifting.thy`. This property is proved automatically if the involved type is BNF without dead variables.

relator_mono registers a property describing a monotonicity of a relator. E.g., $A \leq B \implies rel_set\ A \leq rel_set\ B$. This property is needed for proving a stronger transfer rule in **lift_definition** when a parametricity theorem for the raw term is specified and also for the reflexivity prover. For examples see `~/src/HOL/Lifting_Set.thy` or `~/src/HOL/Lifting.thy`. This property is proved automatically if the involved type is BNF without dead variables.

relator_distr registers a property describing a distributivity of the relation composition and a relator. E.g., $rel_set\ R \circ rel_set\ S = rel_set\ (R \circ S)$. This property is needed for proving a stronger transfer rule in **lift_definition** when a parametricity theorem for the raw term is specified. When this equality does not hold unconditionally (e.g., for the function type), the user can specified each direction separately and also register multiple theorems with different set of assumptions. This attribute can be used only after the monotonicity property was already registered by *relator_mono*. For examples see `~/src/HOL/Lifting_Set.thy` or `~/src/HOL/Lifting.thy`. This property is proved automatically if the involved type is BNF without dead variables.

quot_del deletes a corresponding Quotient theorem from the Lifting infrastructure and thus de-register the corresponding quotient. This effectively causes that **lift_definition** will not do any lifting for the corresponding type. This attribute is rather used for low-level manipulation with set-up of the Lifting package because **lifting_forget** is preferred for normal usage.


lifting_restore Quotient_thm pcr_def pcr_cr_eq_thm registers the Quotient theorem *Quotient_thm* in the Lifting infrastructure and thus sets up lifting for an abstract type τ (that is defined by *Quotient_thm*). Optional theorems *pcr_def* and *pcr_cr_eq_thm* can be specified to register the parametrized correspondence relation for τ . E.g., for *'a dlist*, *pcr_def* is *pcr_dlist A \equiv list_all2 A $\circ\circ$ cr_dlist* and *pcr_cr_eq_thm* is *pcr_dlist op= = op=*. This attribute is rather used for low-level manipulation with set-up of the Lifting package because using of the bundle τ .*lifting* together with the commands **lifting_forget** and **lifting_update** is preferred for normal usage.


Integration with the BNF package[10]: As already mentioned, the theorems that are registered by the following attributes are proved and registered automatically if the involved type is BNF without dead variables: *quot_map*, *relator_eq_onp*, *relator_mono*, *relator_distr*. Also the definition of a relator and predicator is provided automatically. Moreover, if the BNF represents a datatype, simplification rules for a predicator are again proved automatically.

12.4 Coercive subtyping

coercion : attribute
coercion_enabled : attribute
coercion_map : attribute

Coercive subtyping allows the user to omit explicit type conversions, also called *coercions*. Type inference will add them as necessary when parsing a term. See [48] for details.





coercion f registers a new coercion function $f :: \sigma_1 \Rightarrow \sigma_2$ where σ_1 and σ_2 are type constructors without arguments. Coercions are composed by the inference algorithm if needed. Note that the type inference algorithm is complete only if the registered coercions form a lattice.

coercion_map map registers a new map function to lift coercions through type constructors. The function *map* must conform to the following type pattern

$$\text{map} :: f_1 \Rightarrow \dots \Rightarrow f_n \Rightarrow (\alpha_1, \dots, \alpha_n) t \Rightarrow (\beta_1, \dots, \beta_n) t$$

where t is a type constructor and f_i is of type $\alpha_i \Rightarrow \beta_i$ or $\beta_i \Rightarrow \alpha_i$. Registering a map function overwrites any existing map function for this particular type constructor.

coercion_enabled enables the coercion inference algorithm.

12.5 Arithmetic proof support

arith : *method*
arith : *attribute*
arith_split : *attribute*

arith decides linear arithmetic problems (on types *nat*, *int*, *real*). Any current facts are inserted into the goal before running the procedure.


arith declares facts that are supplied to the arithmetic provers implicitly.

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.

12.6 Intuitionistic proof search

iprover : *method*



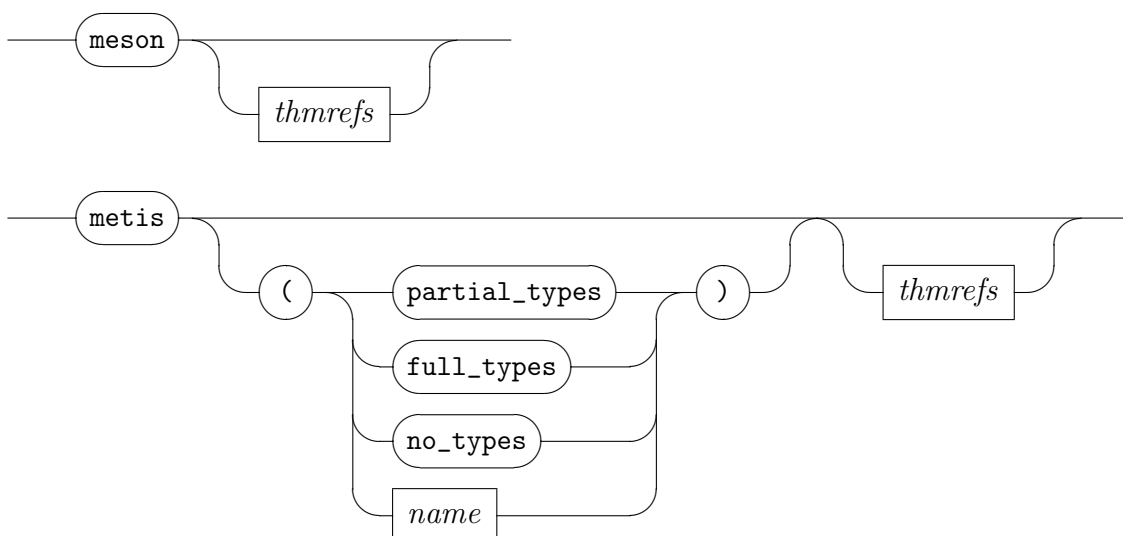
iprover performs intuitionistic proof search, depending on specifically declared rules from the context, or given as explicit arguments. Chained facts are inserted into the goal before commencing proof search.

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.

12.7 Model Elimination and Resolution

meson : *method*

metis : *method*




meson implements Loveland’s model elimination procedure [24]. See `~/src/HOL/ex/Meson_Test.thy` for examples.

metis combines ordered resolution and ordered paramodulation to find first-order (or mildly higher-order) proofs. The first optional argument specifies a type encoding; see the Sledgehammer manual [8] for details. The directory `~/src/HOL/Metis_Examples` contains several small theories developed to a large extent using *metis*.

12.8 Algebraic reasoning via Gröbner bases

algebra : *method*
algebra : *attribute*



algebra performs algebraic reasoning via Gröbner bases, see also [13] and [12, §3.2]. The method handles deals with two main classes of problems:

1. Universal problems over multivariate polynomials in a (semi)-ring/field/idom; the capabilities of the method are augmented according to properties of these structures. For this problem class the method is only complete for algebraically closed fields, since the underlying method is based on Hilbert’s Nullstellensatz, where the equivalence only holds for algebraically closed fields.
 The problems can contain equations $p = 0$ or inequations $q \neq 0$ anywhere within a universal problem statement.
2. All-exists problems of the following restricted (but useful) form:

$$\begin{aligned}
& \forall x_1 \dots x_n. \\
& e_1(x_1, \dots, x_n) = 0 \wedge \dots \wedge e_m(x_1, \dots, x_n) = 0 \longrightarrow \\
& (\exists y_1 \dots y_k. \\
& \quad p_{11}(x_1, \dots, x_n) * y_1 + \dots + p_{1k}(x_1, \dots, x_n) * y_k = 0 \wedge \\
& \quad \dots \wedge \\
& \quad p_{t1}(x_1, \dots, x_n) * y_1 + \dots + p_{tk}(x_1, \dots, x_n) * y_k = 0)
\end{aligned}$$

Here e_1, \dots, e_n and the p_{ij} are multivariate polynomials only in the variables mentioned as arguments.

The proof method is preceded by a simplification step, which may be modified by using the form (*algebra add: ths₁ del: ths₂*). This acts like declarations for the Simplifier (§9.3) on a private simpset for this tool.

algebra (as attribute) manages the default collection of pre-simplification rules of the above proof method.

Example

The subsequent example is from geometry: collinearity is invariant by rotation.

```
type_synonym point = int × int
```

```
fun collinear :: point ⇒ point ⇒ point ⇒ bool where
  collinear (Ax, Ay) (Bx, By) (Cx, Cy) ⟷
  (Ax - Bx) * (By - Cy) = (Ay - By) * (Bx - Cx)
```


```
lemma collinear_inv_rotation:
```

```
  assumes collinear (Ax, Ay) (Bx, By) (Cx, Cy) and c2 + s2 = 1
  shows collinear (Ax * c - Ay * s, Ay * c + Ax * s)
    (Bx * c - By * s, By * c + Bx * s) (Cx * c - Cy * s, Cy * c + Cx * s)
  using assms by (algebra add: collinear.simps)
```

See also `~/src/HOL/ex/Groebner_Examples.thy`.

12.9 Coherent Logic

```
coherent : method
```



coherent solves problems of *Coherent Logic* [7], which covers applications in confluence theory, lattice theory and projective geometry. See `~/src/HOL/ex/Coherent.thy` for some examples.

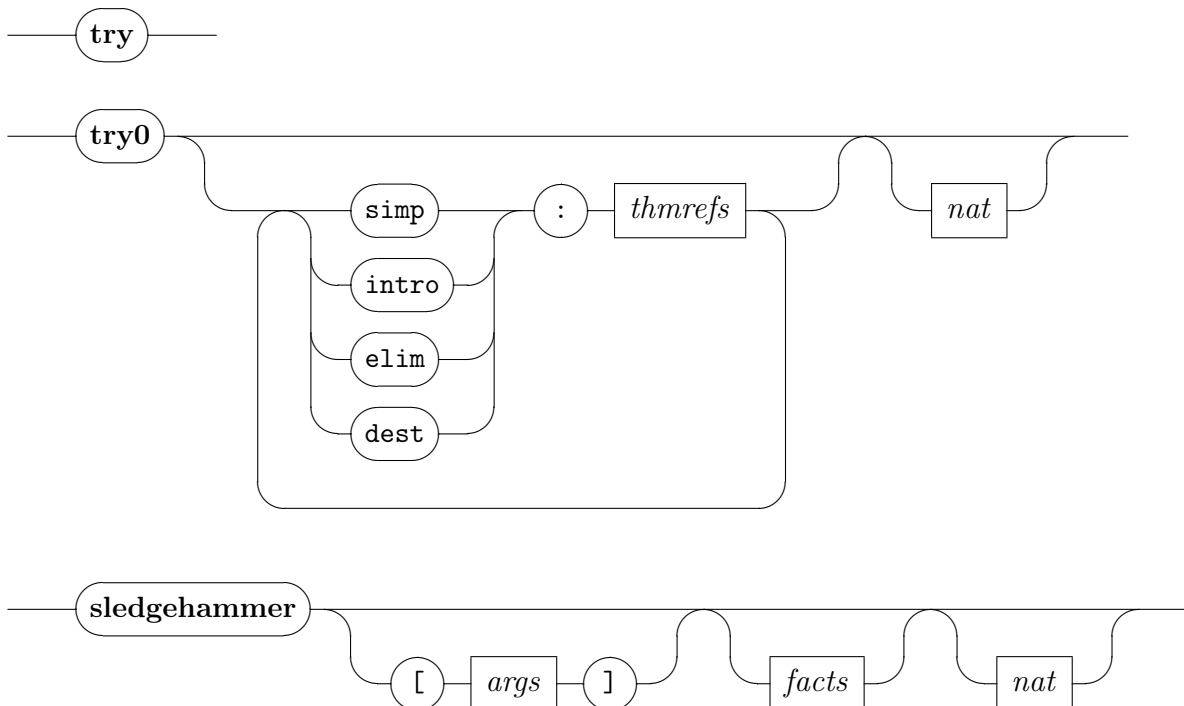
12.10 Proving propositions

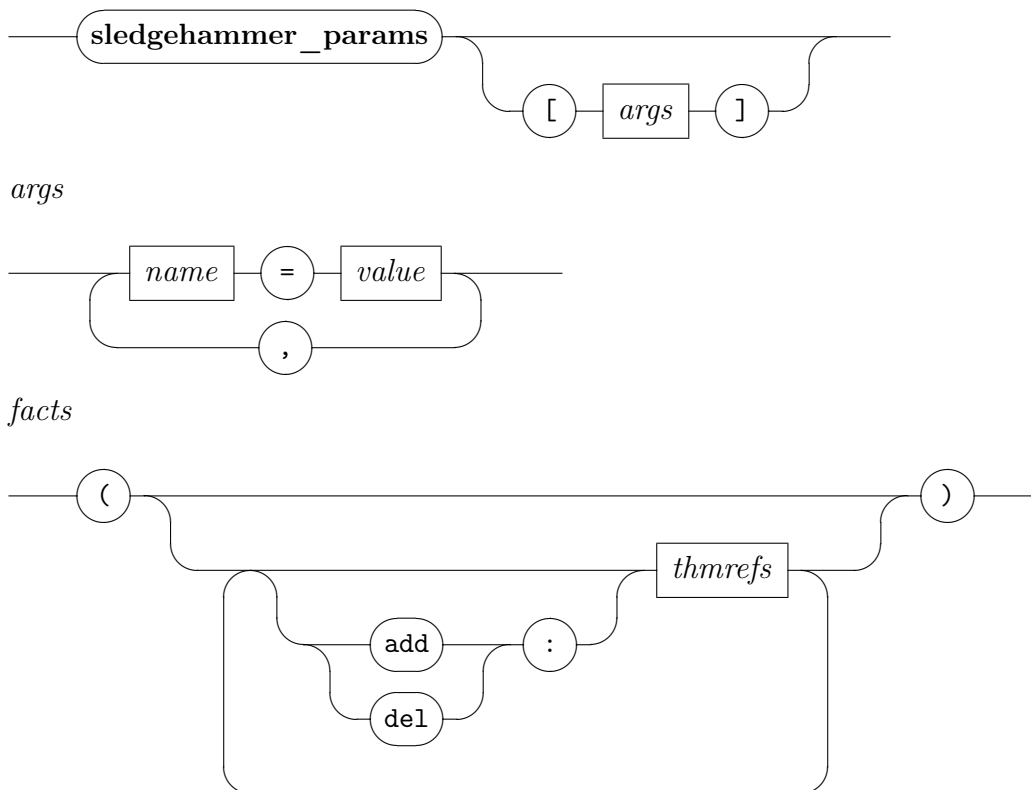
In addition to the standard proof methods, a number of diagnosis tools search for proofs and provide an Isar proof snippet on success. These tools are available via the following commands.

```

solve_direct* : proof →
try* : proof →
try0* : proof →
sledgehammer* : proof →
sledgehammer_params : theory → theory

```





solve_direct checks whether the current subgoals can be solved directly by an existing theorem. Duplicate lemmas can be detected in this way.

try0 attempts to prove a subgoal using a combination of standard proof methods (*auto*, *simp*, *blast*, etc.). Additional facts supplied via *simp:*, *intro:*, *elim:*, and *dest:* are passed to the appropriate proof methods.

try attempts to prove or disprove a subgoal using a combination of provers and disprovers (**solve_direct**, **quickcheck**, **try0**, **sledgehammer**, **nitpick**).

sledgehammer attempts to prove a subgoal using external automatic provers (resolution provers and SMT solvers). See the Sledgehammer manual [8] for details.

sledgehammer_params changes **sledgehammer** configuration options persistently.

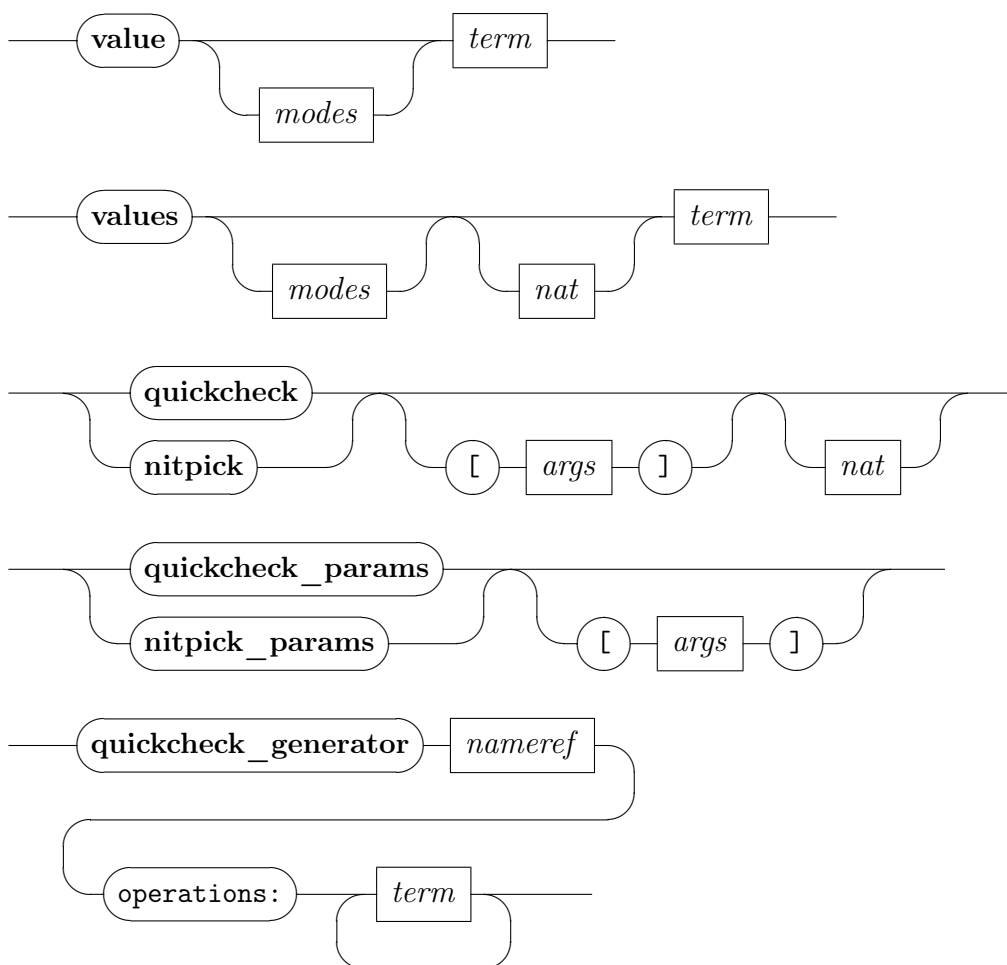
12.11 Checking and refuting propositions


Identifying incorrect propositions usually involves evaluation of particular assignments and systematic counterexample search. This is supported by the following commands.

```


value*      : context →
values*    : context →
quickcheck* : proof →
nitpick*   : proof →
quickcheck_params : theory → theory
nitpick_params   : theory → theory
quickcheck_generator : theory → theory
find_unused_assms : context →

```






modes



args



value t evaluates and prints a term; optionally *modes* can be specified, which are appended to the current print mode; see §7.1.3. Evaluation is tried first using ML, falling back to normalization by evaluation if this fails.

values t enumerates a set comprehension by evaluation and prints its values up to the given number of solutions; optionally *modes* can be specified, which are appended to the current print mode; see §7.1.3.

quickcheck tests the current goal for counterexamples using a series of assignments for its free variables; by default the first subgoal is tested, another can be selected explicitly using an optional goal index. Assignments can be chosen exhausting the search space up to a given size, or using a fixed number of random assignments in the search space, or exploring the search space symbolically using narrowing. By default, quickcheck uses exhaustive testing. A number of configuration options are supported for **quickcheck**, notably:

tester specifies which testing approach to apply. There are three testers, *exhaustive*, *random*, and *narrowing*. An unknown configuration option is treated as an argument to *tester*, making *tester* = optional. When multiple testers are given, these are applied in parallel. If no tester is specified, quickcheck uses the testers that are set active, i.e., configurations *quickcheck_exhaustive_active*, *quickcheck_random_active*, *quickcheck_narrowing_active* are set to true.

- size* specifies the maximum size of the search space for assignment values.
- genuine_only* sets quickcheck only to return genuine counterexample, but not potentially spurious counterexamples due to underspecified functions.
- abort_potential* sets quickcheck to abort once it found a potentially spurious counterexample and to not continue to search for a further genuine counterexample. For this option to be effective, the *genuine_only* option must be set to false.
- eval* takes a term or a list of terms and evaluates these terms under the variable assignment found by quickcheck. This option is currently only supported by the default (exhaustive) tester.
- iterations* sets how many sets of assignments are generated for each particular size.
- no_assms* specifies whether assumptions in structured proofs should be ignored.
- locale* specifies how to process conjectures in a locale context, i.e., they can be interpreted or expanded. The option is a whitespace-separated list of the two words *interpret* and *expand*. The list determines the order they are employed. The default setting is to first use interpretations and then test the expanded conjecture. The option is only provided as attribute declaration, but not as parameter to the command.
- 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.
- use_subtype* if set quickcheck automatically lifts conjectures to registered subtypes if possible, and tests the lifted conjecture.
- quiet* if set quickcheck does not output anything while testing.
- verbose* if set quickcheck informs about the current size and cardinality while testing.
- 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.

Using the following type classes, the testers generate values and convert them back into Isabelle terms for displaying counterexamples.

exhaustive The parameters of the type classes *exhaustive* and *full_exhaustive* implement the testing. They take a testing function as a parameter, which takes a value of type *'a* and optionally produces a counterexample, and a size parameter for the test values. In *full_exhaustive*, the testing function parameter additionally expects a lazy term reconstruction in the type *Code_Evaluation.term* of the tested value.

The canonical implementation for *exhaustive* testers calls the given testing function on all values up to the given size and stops as soon as a counterexample is found.

random The operation *Quickcheck_Random.random* of the type class *random* generates a pseudo-random value of the given size and a lazy term reconstruction of the value in the type *Code_Evaluation.term*. A pseudo-randomness generator is defined in theory *Random*.

narrowing implements Haskell's Lazy Smallcheck [46] using the type classes *narrowing* and *partial_term_of*. Variables in the current goal are initially represented as symbolic variables. If the execution of the goal tries to evaluate one of them, the test engine replaces it with refinements provided by *narrowing*. Narrowing views every value as a sum-of-products which is expressed using the operations *Quickcheck_Narrowing.cons* (embedding a value), *Quickcheck_Narrowing.apply* (product) and *Quickcheck_Narrowing.sum* (sum). The refinement should enable further evaluation of the goal.

For example, *narrowing* for the list type *'a :: narrowing list* can be recursively defined as *Quickcheck_Narrowing.sum* (*Quickcheck_Narrowing.cons* []) (*Quickcheck_Narrowing.apply* (*Quickcheck_Narrowing.apply* (*Quickcheck_Narrowing.cons* (op #)) *narrowing*) *narrowing*). If a symbolic variable of type *_ list* is evaluated, it is replaced by (i) the empty list [] and (ii) by a non-empty list whose head and tail can then be recursively refined if needed.

To reconstruct counterexamples, the operation *partial_term_of* transforms *narrowing*'s deep representation of terms to the type *Code_Evaluation.term*. The deep representation models symbolic variables as *Quickcheck_Narrowing.Narrowing_variable*, which

are normally converted to *Code_Evaluation.Free*, and refined values as *Quickcheck_Narrowing.Narrowing_constructor i args*, where $i :: integer$ denotes the index in the sum of refinements. In the above example for lists, 0 corresponds to `[]` and 1 to `op #`.

The command `code_datatype` sets up *partial_term_of* such that the i -th refinement is interpreted as the i -th constructor, but it does not ensures consistency with *narrowing*.

`quickcheck_params` changes `quickcheck` configuration options persistently.

`quickcheck_generator` creates random and exhaustive value generators for a given type and operations. It generates values by using the operations as if they were constructors of that type.

`nitpick` tests the current goal for counterexamples using a reduction to first-order relational logic. See the Nitpick manual [9] for details.

`nitpick_params` changes `nitpick` configuration options persistently.


`find_unused_assms` finds potentially superfluous assumptions in theorems using quickcheck. It takes the theory name to be checked for superfluous assumptions as optional argument. If not provided, it checks the current theory. Options to the internal quickcheck invocations can be changed with common configuration declarations.

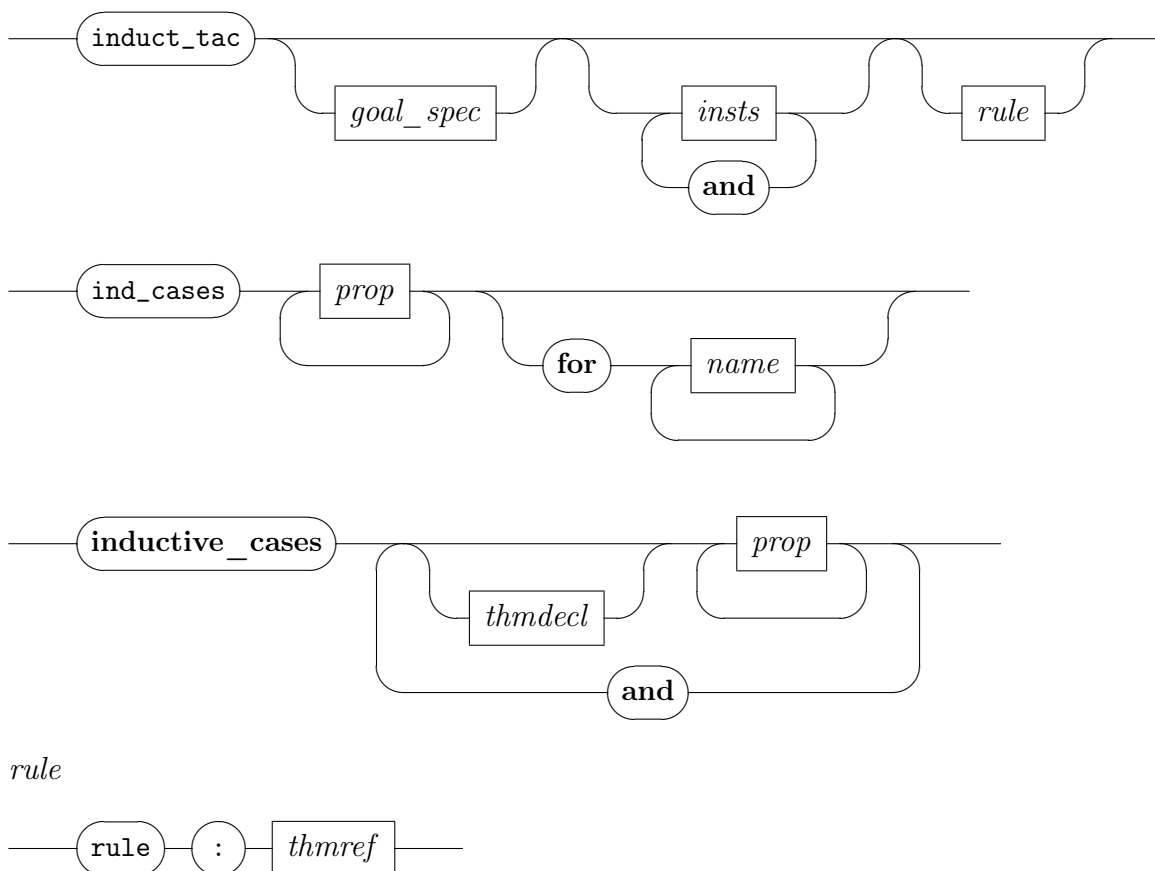
12.12 Unstructured case analysis and induction

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.

Executable code

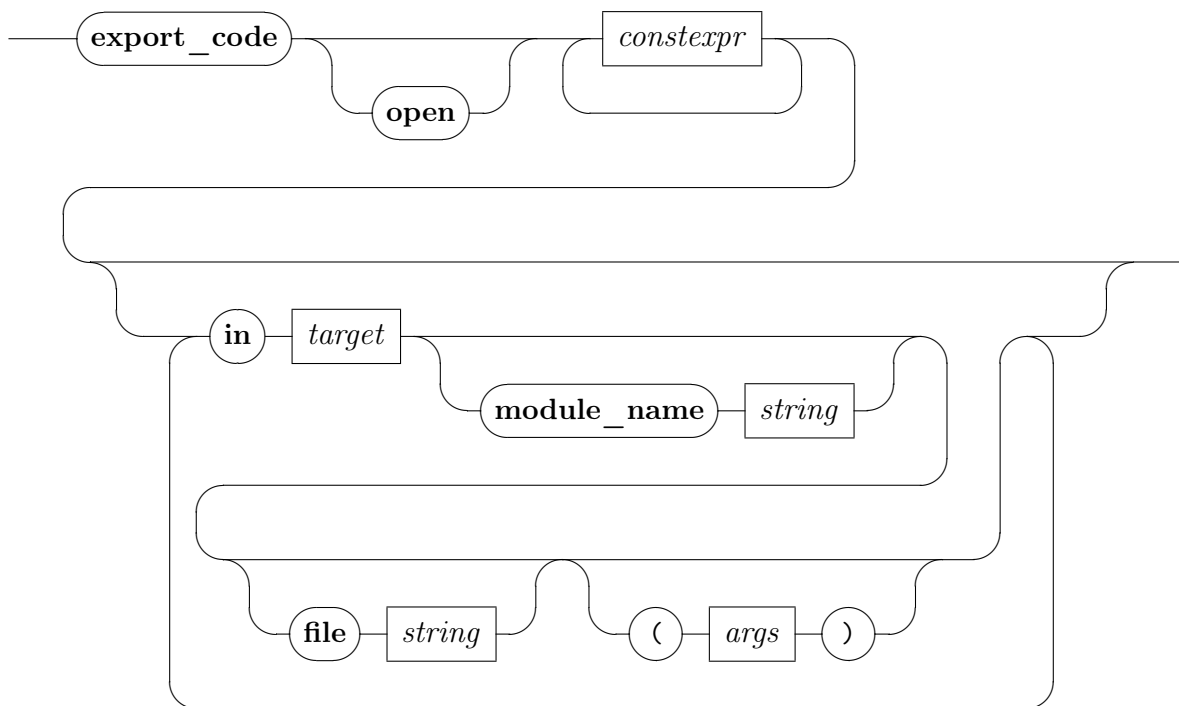
For validation purposes, it is often useful to *execute* specifications. In principle, execution could be simulated by Isabelle’s inference kernel, i.e. by a combination of resolution and simplification. Unfortunately, this approach is rather inefficient. A more efficient way of executing specifications is to translate them into a functional programming language such as ML.

Isabelle provides a generic framework to support code generation from executable specifications. Isabelle/HOL instantiates these mechanisms in a way that is amenable to end-user applications. Code can be generated for functional programs (including overloading using type classes) targeting SML [27], OCaml [23], Haskell [43] and Scala [35]. 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 [18] for an introduction.


```

export_code* : context →
               code : attribute
code_datatype : theory → theory
print_codesetup* : context →
                   code_unfold : attribute
                   code_post : attribute
                   code_abbrev : attribute
print_codeproc* : context →
                  code_thms* : context →
                  code_deps* : context →
                  code_reserved : theory → theory
                  code_printing : theory → theory
code_identifier : theory → theory
code_monad : theory → theory
code_reflect : theory → theory
code_pred : theory → proof(prove)


```




const



constexpr



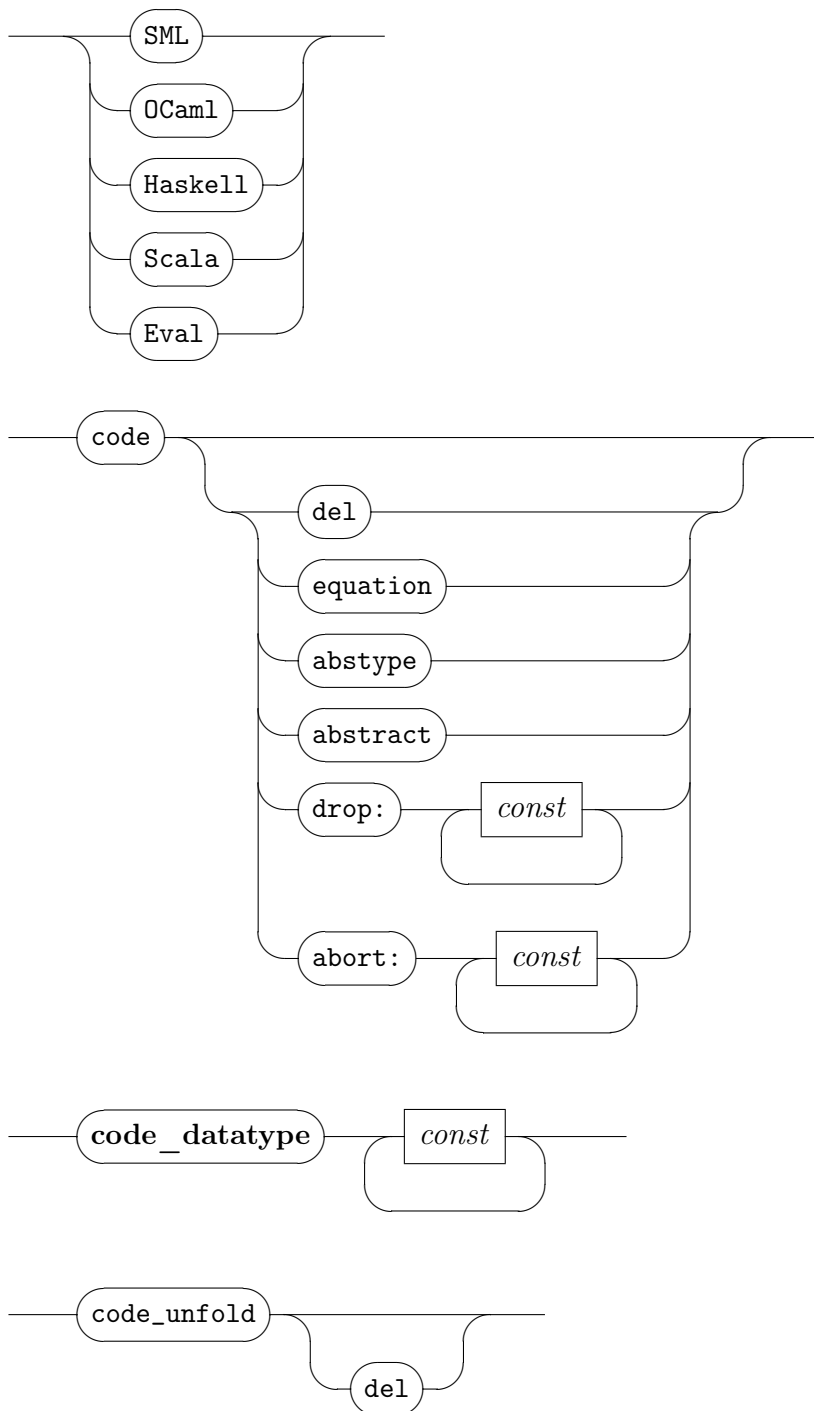
typeconstructor

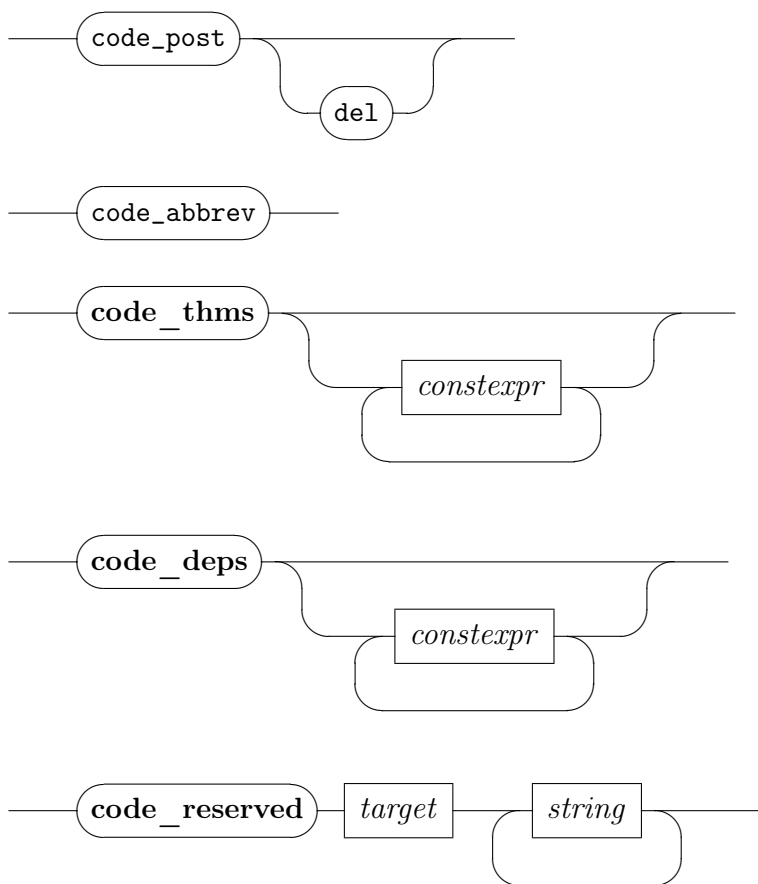


class




target







symbol_const




symbol_typeconstructor




symbol_class




symbol_class_relation




symbol_class_instance




symbol_module



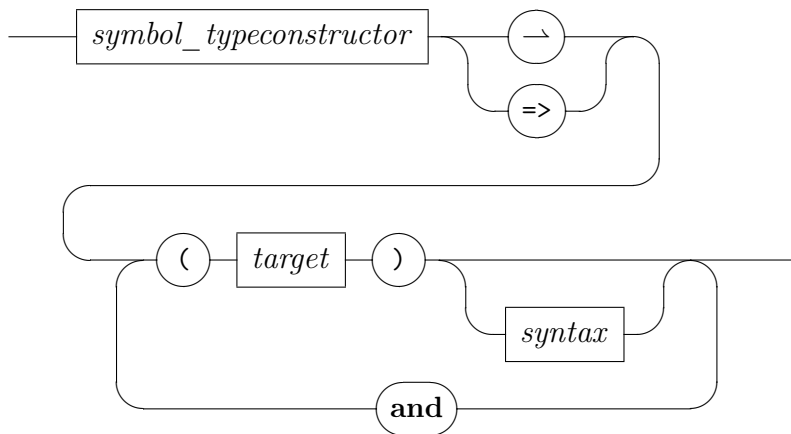
syntax




printing_const




printing_typeconstructor



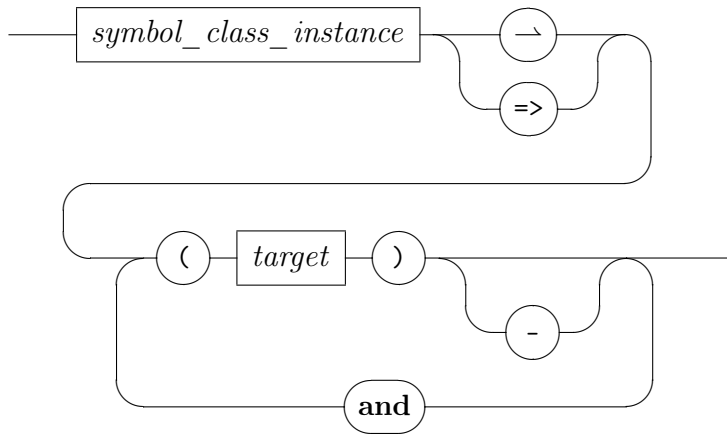
printing_class




printing_class_relation

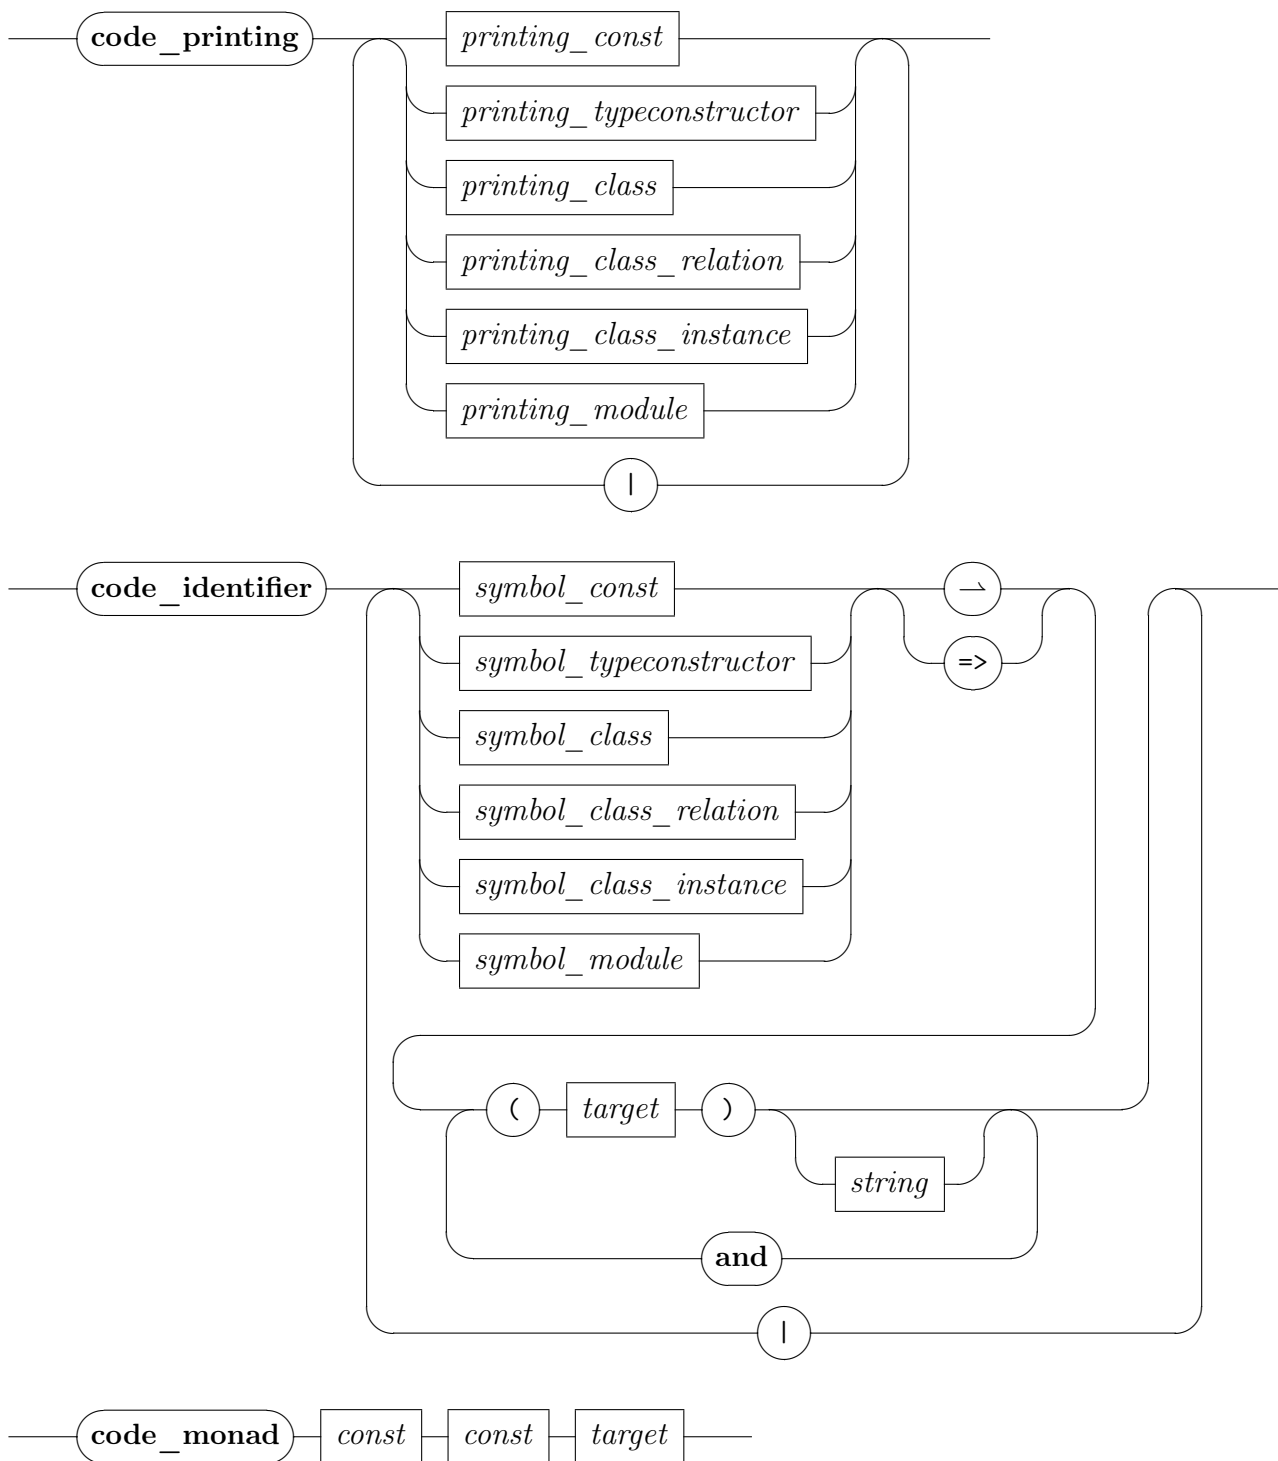


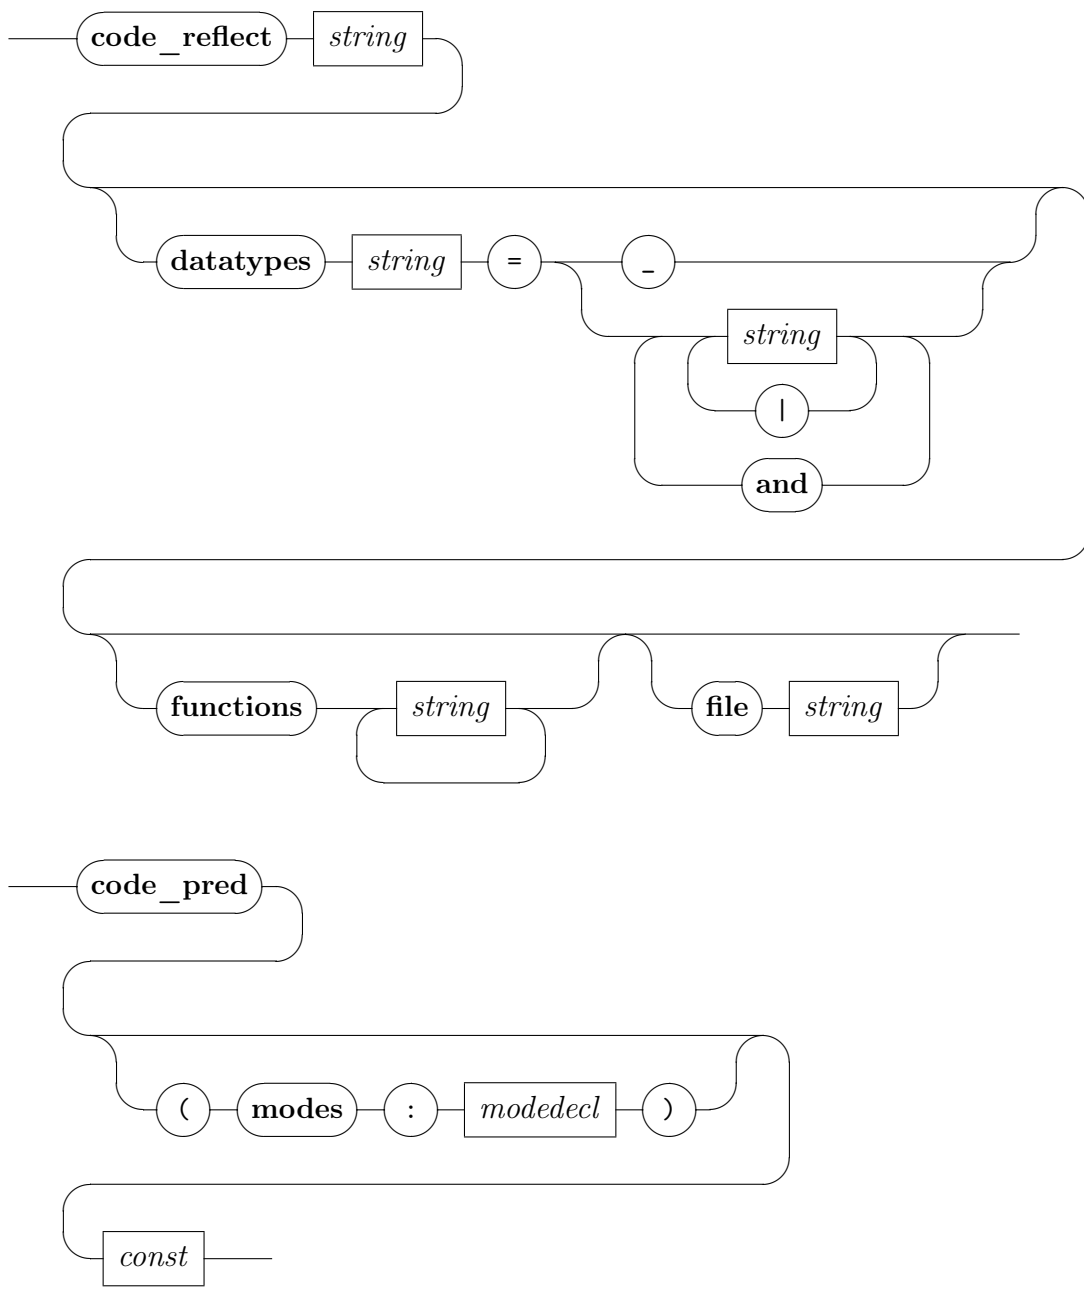
printing_class_instance




printing_module








modedecl



modes



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, exported identifiers are minimized per module. This can be suppressed by prepending **open** before the list of constants.

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 *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 declare code equations for code generation. Variant *code equation* declares a conventional equation as code equation. Variants *code abstype*

and *code abstract* declare abstract datatype certificates or code equations on abstract datatype representations respectively. Vanilla *code* falls back to *code equation* or *code abstype* depending on the syntactic shape of the underlying equation. Variant *code del* deselects a code equation for code generation.

Variants *code drop:* and *code abort:* take a list of constant as arguments and drop all code equations declared for them. In the case of text abort, these constants then are not required to have a definition by means of code equations; if needed these are implemented by program abort (exception) instead.

Usually packages introducing code equations provide a reasonable default setup for selection.

code_datatype specifies a constructor set for a logical type.

print_codesetup gives an overview on selected code equations and code generator datatypes.

code_unfold declares (or with option “*del*” removes) theorems which during preprocessing are applied as rewrite rules to any code equation or evaluation input.

code_post declares (or with option “*del*” removes) theorems which are applied as rewrite rules to any result of an evaluation.

code_abbrev declares (or with option “*del*” removes) equations which are applied as rewrite rules to any result of an evaluation and symmetrically during preprocessing to any code equation or evaluation input.

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_reserved declares a list of names as reserved for a given target, preventing it to be shadowed by any generated code.

code_printing associates a series of symbols (constants, type constructors, classes, class relations, instances, module names) with target-specific serializations; omitting a serialization deletes an existing serialization.

code_monad provides an auxiliary mechanism to generate monadic code for Haskell.

code_identifier associates a series of symbols (constants, type constructors, classes, class relations, instances, module names) with target-specific hints how these symbols shall be named. These hints gain precedence over names for symbols with no hints at all. Conflicting hints are subject to name disambiguation. *Warning:* It is at the discretion of the user to ensure that name prefixes of identifiers in compound statements like type classes or datatypes are still the same.

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.

code_pred creates code equations for a predicate given a set of introduction rules. Optional mode annotations determine which arguments are supposed to be input or output. If alternative introduction rules are declared, one must prove a corresponding elimination rule.

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

$proof$	=	$prfx^*$ proof $method^?$ $stmt^*$ qed $method^?$
		$prfx^*$ done
$prfx$	=	apply $method$
		using $facts$
		unfolding $facts$
$stmt$	=	{ $stmt^*$ }
		next
		note $name = facts$
		let $term = term$
		write $name$ ($mixfix$)
		fix var^+
		assume $name: props$
		then [?] $goal$
$goal$	=	have $name: props$ $proof$
		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 = this</i>
also _{$n+1$}	\approx	note <i>calculation = trans [OF calculation this]</i>
finally	\approx	also from <i>calculation</i>
moreover	\approx	note <i>calculation = 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	... fix x assume $a: \varphi$
case c	\approx	fix x assume $c: \varphi$
sorry	\approx	by cheating

A.1.4 Diagnostic commands

print_state	print proof state
print_statement	print fact in long statement form
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

Rules

<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 $\langle name \rangle$ (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 $\langle isasymname \rangle$ for each $\langle name \rangle$ in the `isabellesym.sty` file. Most of these symbols are displayed properly in Isabelle/jEdit and L^AT_EX generated from Isabelle.

Moreover, any single symbol (or ASCII character) may be prefixed by $\langle \sup \rangle$ for superscript and $\langle \sub \rangle$ for subscript, such as $A \langle \sup \rangle \langle star \rangle$ for A^* and $A \langle \sub \rangle 1$ for A_1 . Sub- and superscripts that span a region of text can be marked up with $\langle \bsub \rangle \dots \langle \esub \rangle$ and $\langle \bsup \rangle \dots \langle \esup \rangle$ respectively, but note that there are limitations in the typographic rendering quality of this form. Furthermore, all ASCII characters and most other symbols may be printed in bold by prefixing $\langle \bold \rangle$ such as $\langle \bold \rangle \langle alpha \rangle$ for α . Note that $\langle \sup \rangle$, $\langle \sub \rangle$, $\langle \bold \rangle$ cannot be combined.

Further details of Isabelle document preparation are covered in chapter 4.


$\langle zero \rangle$	0	$\langle one \rangle$	1
$\langle two \rangle$	2	$\langle three \rangle$	3
$\langle four \rangle$	4	$\langle five \rangle$	5
$\langle six \rangle$	6	$\langle seven \rangle$	7
$\langle eight \rangle$	8	$\langle nine \rangle$	9
$\langle A \rangle$	<i>A</i>	$\langle B \rangle$	<i>B</i>
$\langle C \rangle$	<i>C</i>	$\langle D \rangle$	<i>D</i>
$\langle E \rangle$	<i>E</i>	$\langle F \rangle$	<i>F</i>
$\langle G \rangle$	<i>G</i>	$\langle H \rangle$	<i>H</i>
$\langle I \rangle$	<i>I</i>	$\langle J \rangle$	<i>J</i>
$\langle K \rangle$	<i>K</i>	$\langle L \rangle$	<i>L</i>
$\langle M \rangle$	<i>M</i>	$\langle N \rangle$	<i>N</i>
$\langle O \rangle$	<i>O</i>	$\langle P \rangle$	<i>P</i>

<code>\<Q></code>	Q	<code>\<R></code>	\mathcal{R}
<code>\<S></code>	S	<code>\<T></code>	\mathcal{T}
<code>\<U></code>	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>	s	<code>\<tt></code>	t
<code>\<uu></code>	u	<code>\<vv></code>	v
<code>\<ww></code>	w	<code>\<xx></code>	x
<code>\<yy></code>	y	<code>\<zz></code>	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>	\Leftrightarrow	<code>\<Longleftarrow></code>	\Longleftarrow
<code>\<Rightarrow></code>	\Rightarrow	<code>\<Longrightarrow></code>	\Longrightarrow
<code>\<leftrightharpoon></code>	\leftrightarrow	<code>\<longleftrightharpoon></code>	\longleftrightarrow
<code>\<Leftrightarrow></code>	\Leftrightarrow	<code>\<Longleftrightharpoon></code>	\Longleftrightarrow
<code>\<mapsto></code>	\mapsto	<code>\<longmapsto></code>	\longmapsto
<code>\<midarrow></code>	\dashv	<code>\<Midarrow></code>	\equiv
<code>\<hookleftarrow></code>	\hookleftarrow	<code>\<hookrightarrow></code>	\hookrightarrow
<code>\<leftharpoon></code>	\leftharpoonright	<code>\<rightharpoon></code>	\rightharpoonright
<code>\<leftharpoon></code>	\leftarrow	<code>\<rightharpoon></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>	\upharpoonright	<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\! $	<code>\<rparr></code>	$\! \rangle$
<code>\<lbrakk></code>	\llbracket	<code>\<rbrakk></code>	\rrbracket
<code>\<lbrace></code>	$\{$	<code>\<rbrace></code>	$\}$
<code>\<guillemotleft></code>	\llcorner	<code>\<guillemotright></code>	\lrcorner
<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>	\vDash
<code>\<tturnstile></code>	\vDash	<code>\<TTturnstile></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>	\lesapprox	<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>\<sqsubseteq></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>	\oplus	<code>\<Uplus></code>	\bigoplus
<code>\<noteq></code>	\neq	<code>\<sim></code>	\sim
<code>\<doteq></code>	\doteq	<code>\<simeq></code>	\simeq
<code>\<approx></code>	\approx	<code>\<asymp></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>	$\bar{}$
<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>\<rhhd></code>	\triangleleft
<code>\<unlhd></code>	\triangleleft	<code>\<unrhhd></code>	\triangleleft
<code>\<triangleleft></code>	\triangleleft	<code>\<triangleright></code>	\triangleright
<code>\<triangle></code>	\triangle	<code>\<triangleq></code>	\triangleq
<code>\<oplus></code>	\oplus	<code>\<Oplus></code>	\oplus
<code>\<otimes></code>	\otimes	<code>\<Otimes></code>	\otimes
<code>\<odot></code>	\odot	<code>\<Odot></code>	\odot
<code>\<ominus></code>	\ominus	<code>\<oslash></code>	\oslash
<code>\<dots></code>	\dots	<code>\<cdots></code>	\dots
<code>\<Sum></code>	\sum	<code>\<Prod></code>	\prod
<code>\<Coprod></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>\<inverse></code>	$^{-1}$	<code>\<onequarter></code>	$\frac{1}{4}$
<code>\<onehalf></code>	$\frac{1}{2}$	<code>\<threequarters></code>	$\frac{3}{4}$
<code>\<ordfeminine></code>	$^{\text{a}}$	<code>\<ordmasculine></code>	$^{\text{o}}$
<code>\<section></code>	\S	<code>\<paragraph></code>	\P
<code>\<exclamdown></code>	$\!;$	<code>\<questiondown></code>	$\!?$
<code>\<euro></code>	$\text{\text{€}}$	<code>\<pounds></code>	$\text{\text{£}}$
<code>\<yen></code>	$\text{\text{¥}}$	<code>\<cent></code>	$\text{\text{¢}}$
<code>\<currency></code>	$\text{\text{¤}}$	<code>\<degree></code>	$^{\circ}$
<code>\<hyphen></code>	$-$	<code>\<amalg></code>	\amalg

<code>\<mho></code>	Ω	<code>\<lozenge></code>	◇
<code>\<wp></code>	∅	<code>\<wrong></code>	∫
<code>\<struct></code>	◇	<code>\<acute></code>	´
<code>\<index></code>	1	<code>\<dieresis></code>	¨
<code>\<cedilla></code>	¸	<code>\<hungarumlaut></code>	˘
<code>\<module></code>	⟨ module ⟩	<code>\<bind></code>	⋈
<code>\<then></code>	⋈	<code>\<some></code>	€
<code>\<open></code>	⟨	<code>\<close></code>	⟩
<code>\<newline></code>			

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,

or any of its “improper” variants *erule*, *drule*, *frule*. 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 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 @{context} 1</code>	<code>simp</code>
<code>ALLGOALS (asm_full_simp_tac @{context})</code>	<code>simp_all</code>
<code>simp_tac @{context} 1</code>	<code>simp (no_asm)</code>
<code>asm_simp_tac @{context} 1</code>	<code>simp (no_asm_simp)</code>
<code>full_simp_tac @{context} 1</code>	<code>simp (no_asm_use)</code>
<code>asm_lr_simp_tac @{context} 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. 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.

```

stac a 1          subst a
hyp_subst_tac 1  hypsubst
split_all_tac 1  simp (no_asm_simp) only: split_tupled_all
                  ≈  simp only: split_tupled_all
                  <<  clarify

```

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:

```

tac1 THEN tac2      meth1, meth2
tac1 ORELSE tac2    meth1 | meth2
TRY tac              meth?
REPEAT1 tac         meth+
REPEAT tac          (meth+)?
EVERY [tac1, ...]   meth1, ...
FIRST [tac1, ...]  meth1 | ...

```

CHANGED (see [50]) 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 [50]) 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: subgoals are solved in a different order.

Iterated resolution, such as REPEAT (FIRSTGOAL (resolve_tac ...)), is usually better expressed using the *intro* and *elim* methods of Isar (see §9.4).

Bibliography

- [1] P. Andrews. *An Introduction to Mathematical Logic and Type Theory: to Truth through Proof*. Computer Science and Applied Mathematics. Academic Press, 1986.
- [2] D. Aspinall. Proof General. <http://proofgeneral.inf.ed.ac.uk/>.
- [3] 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.
- [4] C. Ballarin. Locales: A module system for mathematical theories. *Journal of Automated Reasoning*, 52(2):123–153, 2014.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] J. C. Blanchette. *Hammering Away: A User’s Guide to Sledgehammer for Isabelle/HOL*. <http://isabelle.in.tum.de/doc/sledgehammer.pdf>.
- [9] J. C. Blanchette. *Picking Nits: A User’s Guide to Nitpick for Isabelle/HOL*. <http://isabelle.in.tum.de/doc/nitpick.pdf>.
- [10] J. C. Blanchette, M. Desharnais, L. Panny, A. Popescu, and D. Traytel. *Defining (Co)datatypes in Isabelle/HOL*. <http://isabelle.in.tum.de/doc/datatypes.pdf>.
- [11] R. S. Boyer and J. S. Moore. *A Computational Logic Handbook*. Academic Press, 1988.

- [12] A. Chaieb. *Automated methods for formal proofs in simple arithmetics and algebra*. PhD thesis, Technische Universität München, 2008.
<http://www4.in.tum.de/~chaieb/pubs/pdf/diss.pdf>.
- [13] A. Chaieb and M. Wenzel. Context aware calculation and deduction — ring equalities via Gröbner Bases in Isabelle. In M. Kauers, M. Kerber, R. Miner, and W. Windsteiger, editors, *Towards Mechanized Mathematical Assistants (CALCULEMUS 2007)*, volume 4573. Springer-Verlag, 2007.
- [14] A. Church. A formulation of the simple theory of types. *Journal of Symbolic Logic*, 5:56–68, 1940.
- [15] K. Futatsugi, J. Goguen, J.-P. Jouannaud, and J. Meseguer. Principles of OBJ2. In *Symposium on Principles of Programming Languages*, pages 52–66, 1985.
- [16] G. Gentzen. Untersuchungen über das logische Schließen. *Math. Zeitschrift*, 1935.
- [17] M. J. C. Gordon and T. F. Melham, editors. *Introduction to HOL: A Theorem Proving Environment for Higher Order Logic*. Cambridge University Press, 1993.
- [18] F. Haftmann. *Code generation from Isabelle theories*.
<http://isabelle.in.tum.de/doc/codegen.pdf>.
- [19] F. Haftmann. *Haskell-style type classes with Isabelle/Isar*.
<http://isabelle.in.tum.de/doc/classes.pdf>.
- [20] B. Huffman and O. Kunčar. Lifting and Transfer: A Modular Design for Quotients in Isabelle/HOL. In *Certified Programs and Proofs (CPP 2013)*, volume 8307 of *Lecture Notes in Computer Science*. Springer-Verlag, 2013.
- [21] A. Krauss. *Defining Recursive Functions in Isabelle/HOL*.
<http://isabelle.in.tum.de/doc/functions.pdf>.
- [22] A. Krauss. *Automating Recursive Definitions and Termination Proofs in Higher-Order Logic*. PhD thesis, Institut für Informatik, Technische Universität München, Germany, 2009.
- [23] X. Leroy et al. *The Objective Caml system – Documentation and user’s manual*. <http://caml.inria.fr/pub/docs/manual-ocaml/>.
- [24] D. W. Loveland. *Automated Theorem Proving: A Logical Basis*. North-Holland Publishing Co., 1978.
- [25] U. Martin and T. Nipkow. Ordered rewriting and confluence. In M. E. Stickel, editor, *10th International Conference on Automated Deduction*, LNAI 449, pages 366–380. Springer, 1990.

- [26] D. Miller. A logic programming language with lambda-abstraction, function variables, and simple unification. *Journal of Logic and Computation*, 1(4), 1991.
- [27] R. Milner, M. Tofte, and R. Harper. *The Definition of Standard ML*. MIT Press, 1990.
- [28] 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.
- [29] T. Nipkow. Functional unification of higher-order patterns. In M. Vardi, editor, *Eighth Annual Symposium on Logic in Computer Science*, pages 64–74. IEEE Computer Society Press, 1993.
- [30] 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.
- [31] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle's Logics: HOL*. <http://isabelle.in.tum.de/doc/logics-HOL.pdf>.
- [32] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*. Springer, 2002. LNCS 2283.
- [33] 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.
- [34] T. Nipkow and C. Prehofer. Type reconstruction for type classes. *Journal of Functional Programming*, 5(2):201–224, 1995.
- [35] M. Odersky and al. An overview of the scala programming language. Technical Report IC/2004/64, EPFL Lausanne, Switzerland, 2004.
- [36] D. C. Oppen. Pretty printing. *ACM Transactions on Programming Languages and Systems*, 2(4), 1980.
- [37] L. C. Paulson. *Isabelle's Logics: FOL and ZF*. <http://isabelle.in.tum.de/doc/logics-ZF.pdf>.
- [38] L. C. Paulson. Natural deduction as higher-order resolution. *Journal of Logic Programming*, 3:237–258, 1986.
- [39] L. C. Paulson. The foundation of a generic theorem prover. *Journal of Automated Reasoning*, 5(3):363–397, 1989.

- [40] L. C. Paulson. Isabelle: The next 700 theorem provers. In P. Odifreddi, editor, *Logic and Computer Science*, pages 361–386. Academic Press, 1990.
- [41] L. C. Paulson. *ML for the Working Programmer*. Cambridge University Press, 2nd edition, 1996.
- [42] F. J. Pelletier. Seventy-five problems for testing automatic theorem provers. *Journal of Automated Reasoning*, 2:191–216, 1986. Errata, JAR 4 (1988), 235–236 and JAR 18 (1997), 135.
- [43] 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/>.
- [44] A. Pitts. The HOL logic. In M. J. C. Gordon and T. F. Melham, editors, *Introduction to HOL: A Theorem Proving Environment for Higher Order Logic*, pages 191–232. Cambridge University Press, 1993.
- [45] P. Rudnicki. An overview of the MIZAR project. In *1992 Workshop on Types for Proofs and Programs*. Chalmers University of Technology, Bastad, 1992.
- [46] C. Runciman, M. Naylor, and F. Lindblad. Smallcheck and Lazy Smallcheck: Automatic exhaustive testing for small values. In *Proceedings of the First ACM SIGPLAN Symposium on Haskell (Haskell 2008)*, pages 37–48. ACM, 2008.
- [47] P. Schroeder-Heister. A natural extension of natural deduction. *Journal of Symbolic Logic*, 49(4), 1984.
- [48] D. Traytel, S. Berghofer, and T. Nipkow. Extending Hindley-Milner Type Inference with Coercive Structural Subtyping. In H. Yang, editor, *APLAS 2011*, volume 7078 of *Lecture Notes in Computer Science*, pages 89–104, 2011.
- [49] A. Trybulec. Some features of the Mizar language. Presented at a workshop in Turin, Italy, 1993.
- [50] M. Wenzel. *The Isabelle/Isar Implementation*. <http://isabelle.in.tum.de/doc/implementation.pdf>.
- [51] M. Wenzel. *Isabelle/jEdit*. <http://isabelle.in.tum.de/doc/jedit.pdf>.
- [52] 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.

- [53] M. Wenzel. Isar — a generic interpretative approach to readable formal proof documents. In Y. Bertot, G. Dowek, A. Hirschowitz, C. Paulin, and L. Theys, editors, *Theorem Proving in Higher Order Logics: TPHOLs '99*, volume 1690 of *Lecture Notes in Computer Science*. Springer-Verlag, 1999.
- [54] 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>.
- [55] 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>.
- [56] M. Wenzel. Isabelle/jEdit — a Prover IDE within the PIDE framework. In J. Jeuring et al., editors, *Conference on Intelligent Computer Mathematics (CICM 2012)*, volume 7362 of *LNAI*. Springer, 2012.
- [57] M. Wenzel and S. Berghofer. *The Isabelle System Manual*.
<http://isabelle.in.tum.de/doc/system.pdf>.
- [58] M. Wenzel and L. C. Paulson. Isabelle/Isar. In F. Wiedijk, editor, *The Seventeen Provers of the World*, LNAI 3600. 2006.
- [59] F. Wiedijk. Mizar: An impression. Unpublished paper, 1999.
<http://www.cs.kun.nl/~freek/mizar/mizarintro.ps.gz>.
- [60] 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), **122**
- . (command), **120**
- .. (command), **120**
- ?thesis (variable), 118
- _ (fact), 114
- { (command), **109**
- } (command), **109**

- abbrev (antiquotation), **65**
- abbreviation (command), **82**, 228
- abs_def (attribute), **182**
- addafter (ML infix), **220**
- addbefore (ML infix), **220**
- addloop (ML infix), **205**
- addSafter (ML infix), **220**
- addSbefore (ML infix), **220**
- addSolver (ML infix), **203**
- addss (ML), **220**
- addSSolver (ML infix), **203**
- addSss (ML), **220**
- addSWrapper (ML infix), **220**
- addWrapper (ML infix), **220**
- adhoc_overloading (command), **255**
- algebra (HOL attribute), **267**
- algebra (HOL method), **267**
- also (command), **130**
- altstring (syntax), **51**, 52, 60
- and (keyword), 58, 111
- antiquotation (syntax), **66**
- any (inner syntax), **156**, 158
- apply (command), 114, 115, **126**
- apply_end (command), **126**
- aprop (inner syntax), **156**, 158
- args (syntax), **59**
- arith (HOL attribute), **265**
- arith (HOL method), **265**
- arith_split (HOL attribute), **265**
- arity (syntax), **55**
- assms (fact), 115
- assume (command), **110**
- assumes (element), **88**
- assumption (inference), **30**
- assumption (method), **122**
- atom (syntax), **59**
- atomize (attribute), **221**
- atomize (method), **221**
- attribute_setup (command), **99**
- attributes (syntax), **59**
- auto (method), **214**
- axiomatization (command), **84**, 103
- axmdecl (syntax), **60**

- back (command), **126**
- best (method), **214**
- bestsimp (method), **214**
- binder (keyword), **151**
- blast (method), **214**
- break (antiquotation option), **71**
- build (tool), 62
- bundle (command), **81**
- by (command), **120**

- calculation (fact), 130
- cartouche (inner syntax), **154**
- cartouche (syntax), **51**, 52, 60
- case (command), **132**
- case_conclusion (attribute), **132**
- case_names (attribute), **132**
- case_tac (HOL method), **275**
- cases (attribute), **140**

- cases (method), 118, 134, **135**
- chapter (command), **62**
- clamod (syntax), **215**
- clarify (method), **218**
- clarify_step (method), **219**
- clarsimp (method), **218**
- clasimpmod (syntax), **216**
- class (antiquotation), **65**
- class (command), **95**
- class_deps (command), **95**
- class_name (inner syntax), **158**
- class_syntax (ML antiquotation), **169**
- classdecl (syntax), **55**
- cleaning (HOL method), **251**
- code (HOL attribute), **278**
- code_abbrev (HOL attribute), **278**
- code_datatype (HOL command), **278**
- code_deps (HOL command), **278**
- code_identifier (HOL command), **278**
- code_monad (HOL command), **278**
- code_post (HOL attribute), **278**
- code_pred (HOL command), **278**
- code_printing (HOL command), **278**
- code_reflect (HOL command), **278**
- code_reserved (HOL command), **278**
- code_thms (HOL command), **278**
- code_unfold (HOL attribute), **278**
- coercion (HOL attribute), **264**
- coercion_enabled (HOL attribute), **264**
- coercion_map (HOL attribute), **264**
- coherent (HOL method), **268**
- coinduct (attribute), **140**
- coinduct (method), **135**
- coinductive (HOL command), **227**
- coinductive_set (HOL command), **227**
- comment (syntax), **54**
- cong (attribute), **193**
- const (antiquotation), **65**
- const_syntax (ML antiquotation), **169**
- constrains (element), **88**
- consts (command), **103**
- consumes (attribute), **132**
- context (command), **79**, 135
- context_elem (syntax), **89**
- contradiction (method), **213**
- corollary (command), **115**
- cut_tac (method), **185**
- datatype (HOL command), 232, **241**
- declaration (command), **85**
- declare (command), **85**
- deepen (method), **214**
- def (command), **110**
- default_sort (command), **102**
- defer (command), **126**
- defines (element), **88**
- definition (command), **82**
- defn (attribute), **82**
- defs (command), **103**
- delloop (ML infix), **205**
- delSWrapper (ML infix), **220**
- delWrapper (ML infix), **220**
- descending (HOL method), **251**
- descending_setup (HOL method), **251**
- dest (attribute), **211**
- dest (Pure attribute), **122**
- discharge (inference), **33**
- display (antiquotation option), **70**
- display_drafts (command), **76**
- done (command), **126**
- drule (method), **180**
- drule_tac (method), **185**
- elim (attribute), **211**
- elim (method), **180**

- elim (Pure attribute), **122**
- elim_format (Pure attribute), **182**
- elim_resolution (inference), **30**
- end (global command), **77**
- end (local command), **79**, 96
- erule (method), **180**
- erule_tac (method), **185**
- eta_contract (antiquotation option),
70
- eta_contract (attribute), **145**, 169
- expansion (inference), **33**
- export_code (HOL command), **278**

- fact (method), 60, **122**
- fail (method), **180**
- fast (method), **214**
- fastforce (method), **214**
- file (antiquotation), **65**
- finally (command), **130**
- find_consts (command), **175**
- find_theorems (command), **175**
- find_unused_assms (HOL com-
mand), **271**
- finish (inference), **29**
- fix (command), **110**
- fixes (element), **88**
- fixes (syntax), **84**
- float (syntax), **51**
- float_const (inner syntax), **154**
- float_token (inner syntax), **154**
- fold (method), **180**
- folded (attribute), **182**
- for (keyword), **105**
- force (method), **214**
- from (command), **113**
- frule (method), **180**
- frule_tac (method), **185**
- full_prf (antiquotation), **65**
- full_prf (command), **142**
- fun (HOL command), **231**
- fun_cases (HOL command), **231**

- function (HOL command), **231**
- functor (HOL command), **250**

- goal_spec (syntax), **120**
- goals (antiquotation), **65**
- goals_limit (antiquotation option),
71
- goals_limit (attribute), **145**
- guess (command), **128**

- have (command), **115**
- header (command), **62**
- help (command), **50**
- hence (command), **115**
- hide_class (command), **106**
- hide_const (command), **106**
- hide_fact (command), **106**
- hide_type (command), **106**
- hypsubst (method), **183**

- id (inner syntax), **154**
- ident (syntax), **51**, 154
- idt (inner syntax), **157**, 158
- idts (inner syntax), **157**, 158
- iff (attribute), **211**
- in (keyword), **80**
- include (command), **81**
- includes (keyword), **81**
- includes (syntax), 79, **81**, 116
- including (command), **81**
- ind_cases (HOL method), **275**
- indent (antiquotation option), **71**
- index (inner syntax), **157**, 158
- induct (attribute), **140**
- induct (method), 115, 134, **135**
- induct_simp (attribute), **138**
- induct_tac (HOL method), **275**
- induction (method), **135**
- induction_schema (HOL method),
236
- inductive (HOL command), **227**

- inductive_cases (HOL command), **275**
- inductive_set (HOL command), **227**
- infix (keyword), **151**
- infixl (keyword), **151**
- infixr (keyword), **151**
- init (inference), **29**
- injection (HOL method), **251**
- insert (method), **180**
- inst (syntax), **56**
- inst_step (method), **219**
- instance (command), **95**
- instantiation (command), **95**
- insts (syntax), **56**
- int (syntax), **54**
- interpret (command), **91**
- interpretation (command), **91**
- intro (attribute), **211**
- intro (method), **180**
- intro (Pure attribute), **122**
- intro_classes (method), **95**
- intro_locales (method), **88**
- iprover (HOL method), **266**
- is (keyword), **112**

- judgment (command), **221**

- keywords (keyword), **78**

- lemma (antiquotation), **65**
- lemma (command), **115**
- lemmas (command), **105**
- let (command), **112**
- lexicographic_order (HOL method), **236**
- lift_definition (HOL command), **260–263**
- lifting (HOL method), **251**
- lifting_forget (HOL command), **260**
- lifting_restore (HOL attribute), **260**
- lifting_setup (HOL method), **251**
- lifting_update (HOL command), **260**
- local_setup (command), **99**
- locale (command), **88**
- locale (syntax), **89**
- locale_deps (command), **88**
- locale_expr (syntax), **87**
- logic (inner syntax), **156, 158**
- longid (inner syntax), **154**
- longident (syntax), **51, 154**

- margin (antiquotation option), **71**
- meson (HOL method), **266**
- method (syntax), **119**
- method_setup (command), **127**
- metis (HOL method), **266**
- mixfix (syntax), **149**
- mkroot (tool), **62**
- ML (antiquotation), **65**
- ML (command), **99**
- ML_command (command), **99**
- ML_exception_trace (attribute), **100**
- ML_file (command), **99**
- ML_functor (antiquotation), **65**
- ML_op (antiquotation), **65**
- ML_prf (command), **99**
- ML_print_depth (attribute), **100**
- ML_source_trace (attribute), **100**
- ML_structure (antiquotation), **65**
- ML_type (antiquotation), **65**
- ML_val (command), **99**
- mode (antiquotation option), **71**
- modes (syntax), **143**
- mono (HOL attribute), **227**
- moreover (command), **130**

- name (syntax), **53**
- nameref (syntax), **53**
- names_long (antiquotation option), **70**

- names_long (attribute), **145**
- names_short (antiquotation option), **70**
- names_short (attribute), **145**
- names_unique (antiquotation option), **70**
- names_unique (attribute), **145**
- nat (syntax), **51**, 51, 154
- next (command), **109**
- nitpick (HOL command), **271**
- nitpick_params (HOL command), **271**
- no_adhoc_overloading (command), **255**
- no_notation (command), **152**
- no_syntax (command), **164**
- no_translations (command), **164**
- no_type_notation (command), **152**
- no_vars (attribute), **182**
- nonterminal (command), **164**
- notation (command), **152**
- note (command), **113**
- notepad (command), **108**
- notes (element), **88**
- nothing (fact), 114
- num_const (inner syntax), **154**
- num_token (inner syntax), **154**

- obtain (command), **128**
- obtains (element), **115**, 118
- OF (attribute), **122**
- of (attribute), **122**
- oops (command), **110**
- oracle (command), **106**
- output (keyword), 166
- overloading (command), **98**

- params (attribute), **132**
- parname (syntax), **53**
- parse_ast_translation (command), **169**
- parse_translation (command), **169**
- partial_function (HOL command), **238**
- partial_function_mono (HOL attribute), **238**
- partiality_descending (HOL method), **251**
- partiality_descending_setup (HOL method), **251**
- pat_completeness (HOL method), **236**
- prefer (command), **126**
- prems (fact), 112
- presume (command), **110**
- prf (antiquotation), **65**
- prf (command), **142**
- primrec (HOL command), **231**
- print_abbrevs (command), **82**
- print_ast_translation (command), **169**
- print_attributes (command), **175**
- print_bundles (command), **81**
- print_cases (command), **132**
- print_claset (command), **211**
- print_classes (command), **95**
- print_codeproc (HOL command), **278**
- print_codesetup (HOL command), **278**
- print_commands (command), **50**
- print_defn_rules (command), **82**
- print_dependencies (command), **91**
- print_facts (command), **175**
- print_induct_rules (command), **140**
- print_inductives (command), **227**
- print_interps (command), **91**
- print_locale (command), **88**
- print_locales (command), **88**
- print_methods (command), **175**
- Print_Mode.with_modes (ML), **147**
- print_mode_value (ML), **147**

- print_options (command), **179**
- print_quot_maps (HOL command), **260**
- print_quotconsts (HOL command), **251**
- print_quotients (HOL command), **260**
- print_quotientsQ3 (HOL command), **251**
- print_quotmapsQ3 (HOL command), **251**
- print_rules (command), **122**
- print_simpset (command), **193**
- print_state (command), **142**
- print_statement (command), **115**
- print_syntax (command), **160**, 169, 171
- print_term_bindings (command), **175**
- print_theorems (command), **175**
- print_theory (command), **175**
- print_trans_rules (command), **130**
- print_translation (command), **169**
- proof
 - default, 122
 - fake, 122
 - terminal, 122
 - trivial, 122
- proof (command), 114, 115, **120**, 120, 124
- prop (antiquotation), **65**
- prop (command), **142**
- prop (inner syntax), **156**, 158
- prop (syntax), **56**
- prop_pat (syntax), **57**
- props (syntax), **58**
- ptrn (inner syntax), **157**, 158
- ptrns (inner syntax), **157**, 158
- qed (command), **120**, 120
- quickcheck (HOL command), **271**
- quickcheck_generator (HOL command), **271**
- quickcheck_params (HOL command), **271**
- quot_del (HOL attribute), **260**
- quot_lifted (HOL attribute), **251**
- quot_map (HOL attribute), **260**
- quot_preserve (HOL attribute), **251**
- quot_respect (HOL attribute), **251**
- quot_thm (HOL attribute), **251**
- quotes (antiquotation option), **71**
- quotient_definition (HOL command), **251**
- quotient_type (HOL command), **251**
- rail (antiquotation), **72**
- raw_tactic (method), **185**
- real (syntax), **54**
- recdef (HOL command), **239**
- recdef_cong (HOL attribute), **241**
- recdef_simp (HOL attribute), **241**
- recdef_tc (HOL command), **239**
- recdef_wf (HOL attribute), **241**
- record (HOL command), **244**
- regularize (HOL method), **251**
- relation (HOL method), **236**
- relator_distr (HOL attribute), **260**
- relator_domain (HOL attribute), **257**, **259**
- relator_eq (HOL attribute), **257**
- relator_eq_onp (HOL attribute), **260**
- relator_mono (HOL attribute), **260**
- rename_tac (method), **185**
- rep_datatype (HOL command), 232, **241**
- resolution (inference), **30**
- rotate_tac (method), **185**
- rotated (attribute), **182**
- rule (attribute), **211**
- rule (method), 121, **213**

- rule (Pure attribute), **122**
- rule (Pure method), 114, 121, **122**, 124, 302
- rule_format (attribute), **221**
- rule_tac (method), **185**
- rulify (attribute), **221**

- safe (method), **218**
- safe_step (method), **219**
- schematic_corollary (command), **115**
- schematic_lemma (command), **115**
- schematic_theorem (command), **115**
- sect (command), **62**
- section (command), **62**
- setloop (ML infix), **205**
- setSolver (ML infix), **203**
- setSSolver (ML infix), **203**
- setup (command), **99**
- setup_lifting (HOL command), **260**
- show (command), 111, **115**, 120
- show_abbrevs (antiquotation option), **70**
- show_abbrevs (attribute), **145**
- show_brackets (attribute), **145**
- show_consts (attribute), **145**
- show_hyps (attribute), **145**
- show_main_goal (attribute), **145**
- show_markup (attribute), **145**
- show_question_marks (attribute), **145**
- show_sorts (antiquotation option), **70**
- show_sorts (attribute), **145**
- show_structs (antiquotation option), **70**
- show_tags (attribute), **145**
- show_types (antiquotation option), **70**
- show_types (attribute), **145**
- show_variants (attribute), **255**

- shows (element), **115**
- simp (attribute), **193**
- simp (method), **188**
- simp_all (method), **188**
- simp_break (attribute), **198**
- simp_debug (attribute), **198**
- simp_depth_limit (attribute), **188**
- simp_trace (attribute), **198**
- simp_trace_depth_limit (attribute), **198**
- simp_trace_new (attribute), **198**
- simplified (attribute), **206**
- Simplifier.mk_solver (ML), **203**
- Simplifier.premis_of (ML), **202**
- Simplifier.set_subgoal (ML), **202**
- Simplifier.set_termless (ML), **197**
- simpmod (syntax), **189**
- simproc_setup (command), **200**
- size_change (HOL method), **236**
- sledgehammer (HOL command), **269**
- sledgehammer_params (HOL command), **269**
- slow (method), **214**
- slow_step (method), **219**
- slowsimp (method), **214**
- SML_file (command), **99**
- solve_direct (HOL command), **269**
- solver (ML type), **203**
- sorry (command), 110, **120**
- sort (inner syntax), **157**, 159
- sort (syntax), **55**
- source (antiquotation option), **71**
- specification (HOL command), **254**
- split (attribute), **193**
- split (method), **183**, 190
- split_format (HOL attribute), **257**
- Splitter.add_split (ML), **205**
- Splitter.del_split (ML), **205**
- step (method), **219**
- str_token (inner syntax), **154**
- string (syntax), **51**, 52

- string_token (inner syntax), **154**
- structure (keyword), 90, 158
- subclass (command), **95**
- subgoal_tac (method), **185**
- subgoals (antiquotation), **65**
- sublocale (command), **91**
- subset (command), **62**
- subsection (command), **62**
- subst (method), **183**
- subsubset (command), **62**
- subsubsection (command), **62**
- succeed (method), **180**
- swapped (attribute), **211**
- symident (syntax), **51**
- syntax (command), **164**
- syntax_ambiguity_limit (attribute), **161**
- syntax_ambiguity_warning (attribute), **161**
- syntax_ast_stats (attribute), **164**
- syntax_ast_trace (attribute), **164**
- syntax_const (ML antiquotation), **169**
- syntax_declaration (command), **85**
- tactic (method), **185**
- tagged (attribute), **182**
- tags (syntax), **71**
- target (syntax), **79**
- term (antiquotation), **65**
- term (command), **142**
- term (syntax), **56**
- term abbreviations, 113
- term_pat (syntax), **57**
- term_type (antiquotation), **65**
- termination (HOL command), **231**
- text (antiquotation), **65**
- text (command), **62**
- text (syntax), **54**
- text_raw (command), **62**
- THEN (attribute), **182**
- then (command), **113**, 115
- theorem (command), **115**
- theorems (command), **105**
- theory (antiquotation), **65**
- theory (command), **77**
- thesis (variable), 113
- thin_tac (method), **185**
- this (fact), 108, 113
- this (method), **122**
- this (variable), 113
- thm (antiquotation), **65**
- thm (command), **142**
- thm_deps (command), **175**
- thmdecl (syntax), **60**
- thmdef (syntax), **60**
- thmref (syntax), **61**
- thmrefs (syntax), **61**
- thus (command), **115**
- tid (inner syntax), **154**
- transfer (HOL method), **257**
- transfer' (HOL method), **257**
- Transfer.transferred (HOL attribute), **257**
- transfer_domain_rule (HOL attribute), **257**
- transfer_prover (HOL method), **257**
- transfer_rule (HOL attribute), **257**
- translations (command), **164**
- try (HOL command), **269**
- try0 (HOL command), **269**
- tvar (inner syntax), **154**
- txt (command), **62**
- txt_raw (command), **62**
- typ (antiquotation), **65**
- typ (command), **142**
- type (antiquotation), **65**
- type (inner syntax), **157**, 158
- type (syntax), **55**
- type_name (inner syntax), **157**
- type_notation (command), **152**
- type_synonym (command), **103**, 248

- type_syntax (ML antiquotation), xnum_token (inner syntax), **154**
169
- typed_print_translation (command), **169**
- typeddecl (command), **103**, 103
- typedef (command), **103**, 242
- typedef (HOL command), **247**
- typedefree (syntax), **51**, 154
- typeof (antiquotation), **65**
- typespec (syntax), **56**
- typespec_sorts (syntax), **57**
- typevar (syntax), **51**, 52, 154

- ultimately (command), **130**
- unfold (method), **180**
- unfold_locales (method), **88**
- unfolded (attribute), **182**
- unfolding (command), **113**
- unify_search_bound (attribute),
223
- unify_trace_bound (attribute), **223**
- unify_trace_simp (attribute), **223**
- unify_trace_types (attribute), **223**
- untagged (attribute), **182**
- untransferred (HOL attribute), **257**
- unused_thms (command), **175**
- url (antiquotation), **65**
- using (command), **113**

- value (HOL command), 68, **271**
- values (HOL command), **271**
- var (inner syntax), **154**
- var (syntax), **51**, 52, 154
- vars (syntax), **58**
- verbatim (syntax), **51**, 52

- weak-discharge (inference), **33**
- where (attribute), **122**
- with (command), **113**
- wrapper (ML type), **220**
- write (command), **152**