Interactive Formal Verification

Lawrence C Paulson
Computer Laboratory
University of Cambridge

This lecture course introduces interactive formal proof using Isabelle. The lecture notes consist of copies of the slides, some of which have brief remarks attached. Isabelle documentation can be found on the Internet at the URL http://www.cl.cam.ac.uk/research/hvg/Isabelle/documentation.html. The most important single manual is the *Tutorial on Isabelle/HOL*. Reading the *Tutorial* is an excellent way of learning Isabelle in depth. However, the *Tutorial* is very long and a little outdated; although its details remain correct, it presents a style of proof that has become increasingly obsolete with the advent of structured proofs and ever greater automation. These lecture notes take a very different approach and refer you to specific sections of the *Tutorial* that are particularly appropriate.

Tobias Nipkow has just written a new tutorial entitled *Programming and Proving in Isabelle/HOL*. It is much shorter than the original *Tutorial*, and much more up-to-date. If you would like to read a tutorial from cover to cover, this is the one to read.

The other tutorials listed on the documentation page are mainly for advanced users.

Interactive Formal Verification
1 : Introduction

Lawrence C Paulson
Computer Laboratory
University of Cambridge

Motivation

- Complex systems almost inevitably contain bugs.
- Debugging suffers from diminishing returns. Many critical bugs are *never fixed*!
- Critical systems (avionics, ...) are required to meet a standard of 10^{-9} failures per hour. Testing to such a standard is infeasible.

> "Program testing can be used to show the presence of bugs, but never to show their absence!"
>
> — Edsger W. Dijkstra

What is Interactive Proof?

- Work in a logical formalism
  - with precise definitions of concepts
  - and a formal deductive system
- Construct hierarchies of definitions and proofs
  - libraries of formal mathematics
  - specifications of components and properties

Interactive Theorem Provers

- Based on higher-order logic
  - Isabelle, HOL (many versions), PVS
- Based on constructive type theory
  - Coq, Twelf, Agda, ...
- Based on first-order logic with recursion
  - ACL2

Here are some useful web links:

```
Isabelle: http://www.cl.cam.ac.uk/research/hvg/Isabelle/
HOL4: http://hol.sourceforge.net/
HOL Light: http://www.cl.cam.ac.uk/~jrh13/hol-light/
PVS: http://pvs.csl.sri.com/
Coq: http://coq.inria.fr/
ACL2: http://www.cs.utexas.edu/users/moore/acl2/
```

The LCF Architecture

- A small kernel implements the logic and can generate theorems.
- All specification methods and automatic proof procedures expand to full proofs.
- Unsoundness is less likely with this architecture
- ... but the implementation is more complicated, and performance can suffer.
- Used in Isabelle, HOL, Coq but not PVS or ACL2.

Theorem Provers: Key Features

- Logical formalism (higher-order, type theory etc.)
- Control issues:
  - User interface / Proof language
  - Automation
- Libraries of formalised mathematics
- Tools: typesetting, library search,…

Isabelle

- Isabelle is a generic interactive theorem prover, developed by Lawrence Paulson (Cambridge) and Tobias Nipkow (Munich). First release in 1986.
- Integrated tool support for
  - Automated provers
  - Counter-example finding
  - Code generation from logical terms
  - LaTeX document generation

Higher-Order Logic

- First-order logic extended with functions and sets
- Polymorphic types, including a type of truth values
- No distinction between terms and formulas
- ML-style functional programming

"HOL = functional programming + logic"

Basic Syntax of Formulas

formulas *A*, *B*, ... can be written as

(*A*)	t = u	~*A*
A & *B*	*A* | *B*	*A* --> *B*
A <-> *B*	ALL *x*. *A*	EX *x*. *A*

(Among many others)

Isabelle also supports symbols such as

≤ ≥ ≠ ∧ ∨ ⟶ ↔ ∀ ∃

See the *Tutorial*, section 1.3: “Types, terms and formulae”

Some Syntactic Conventions

In $\forall x.\ A \wedge B$, the quantifier spans the entire formula

Parentheses are **required** in $A \wedge (\forall x\ y.\ B)$

Binary logical connectives associate to the right: $A \to B \to C$ is the same as $A \to (B \to C)$

$\neg A \wedge B = C \vee D$ is the same as $((\neg A) \wedge (B = C)) \vee D$

See the *Tutorial*, section 1.3: “Types, terms and formulae”

Basic Syntax of Terms

- The typed λ-calculus:
  - constants, c
  - variables, x and *flexible* variables, ?x
  - abstractions λx. t
  - function applications t u
- Numerous infix operators and binding operators for arithmetic, set theory, etc.

See the *Tutorial*, section 1.3: “Types, terms and formulae”

Types

- Every term has a type; Isabelle infers the types of terms automatically. We write $t :: \tau$
- Types can be *polymorphic*, with a system of type classes (inspired by the Haskell language) that allows sophisticated overloading.
- A formula is simply a term of type `bool`.
- There are types of ordered pairs and functions.
- Other important types are those of the natural numbers (`nat`) and integers (`int`).

Product Types for Pairs

- (x_1, x_2) has type $\tau_1 * \tau_2$ provided $x_i :: \tau_i$
- $(x_1, \ldots, x_{n\text{-}1}, x_n)$ abbreviates $(x_1, \ldots, (x_{n\text{-}1}, x_n))$
- Extensible record types can also be defined.

Function Types

- Infix operators are curried functions
  - `+ :: nat => nat => nat`
  - `& :: bool => bool => bool`
  - Curried function notation: $\lambda x\ y.\ t$
- Function arguments can be paired
  - Example: `nat*nat => nat`
  - Paired function notation: $\lambda(x,y).\ t$

Arithmetic Types

- `nat`: the natural numbers (nonnegative integers)
  - inductively defined: `0`, `Suc` *n*
  - operators include `+ - * div mod`
  - relations include `< ≤ dvd` (divisibility)
- `int`: the integers, with `+ - * div mod` …
- `rat`, `real`: `+ - * / sin cos ln` …
- arithmetic constants and laws for these types

Only integer constants are available. Note *that* traditional notation for floating point numbers would be inappropriate, but rational numbers can be expressed.

HOL as a Functional Language

recursive data type of lists

```
datatype 'a list = Nil | Cons 'a "'a list"

fun app :: "'a list => 'a list => 'a list"  where
  "app Nil ys = ys"
| "app (Cons x xs) ys = Cons x (app xs ys)"

fun rev  where
  "rev Nil = Nil"
| "rev (Cons x xs) = app (rev xs) (Cons x Nil)"
```

recursive functions
(types can be inferred)

Recursive data types can be defined as in ML, although with somewhat less generality. Recursive functions can also be declared, provided Isabelle can establish their termination; all functions in higher-order logic are total. Naturally terminating recursive definitions pose no difficulties for Isabelle. In complicated situations, it is possible to give a hint.

Proof by Induction

declaring a lemma

use it to simplify other formulas

```
lemma [simp]: "app xs Nil = xs"
  apply (induct xs)
  apply auto
done
```

two steps: *induction* followed by *automation*

end of proof

Example of a *Structured Proof*

- base case and inductive step can be proved explicitly
- Invaluable for proofs that need intricate manipulation of facts

```
lemma "app xs Nil = xs"
proof (induct xs)
  case Nil
  show "app Nil Nil = Nil"
    by auto
next
  case (Cons a xs)
  show "app (Cons a xs) Nil = Cons a xs"
    by auto
qed
```

Interactive Formal Verification 2: Isabelle Theories

Lawrence C Paulson
Computer Laboratory
University of Cambridge

Formal Theories

- Collections of specifications: types, constants, functions, sets and relations…
- even *axioms* occasionally, but it is safer to define explicit models satisfying desired properties.

Axiom systems are frequently inconsistent!

- Theories can specify mathematics, formal models or abstract implementations.

A Tiny Theory

name of the new theory

```
theory BT imports Main begin

datatype 'a bt =
    Lf
  | Br 'a  "'a bt"  "'a bt"

fun reflect :: "'a bt => 'a bt" where
  "reflect Lf = Lf"
| "reflect (Br a t1 t2) = Br a (reflect t2) (reflect t1)"
```

the theory it builds upon

declarations of types, constants, etc

```
lemma reflect_reflect_ident: "reflect (reflect t) = t"
  apply (induct t)
   apply auto
  done

end
```

proving a theorem

See the *Tutorial*, section 1.2 (Theories) and 2.1 (An Introductory Theory).

Notes on Theory Structure

- A theory can *import* any existing theories.
- Types, constants, etc., must be *declared before use*.
- The various declarations and proofs may otherwise appear in any order.
- Many declarations can be confined to *local scopes*.
- A finished theory can be imported by others.

Some Fancy Type Declarations

```
typedecl loc  --"an unspecified type of locations"

type_synonym val   = nat --"values"
type_synonym state = "loc => val"
type_synonym aexp  = "state => val"
type_synonym bexp  = "state => bool"  --"functions on states"

datatype
  com = SKIP
      | Assign loc aexp          ("_ :== _ " 60)
      | Semi   com com           ("_; _"  [60, 60] 10)
      | Cond   bexp com com      ("IF _ THEN _ ELSE _"  60)
      | While  bexp com          ("WHILE _ DO _"  60)
```

a new basic type

end-of-line comments

concrete syntax for commands

recursive type of commands

Notes on Type Declarations

- Type synonyms merely introduce *abbreviations*.
- Recursive data types are less general than in functional programming languages.
  - No recursion into the domain of a function.
  - Mutually recursive definitions can be tricky.
- Recursive types are equipped with proof methods for *induction* and *case analysis*.

See the tutorial, section 2.5.

Basic Constant Definitions

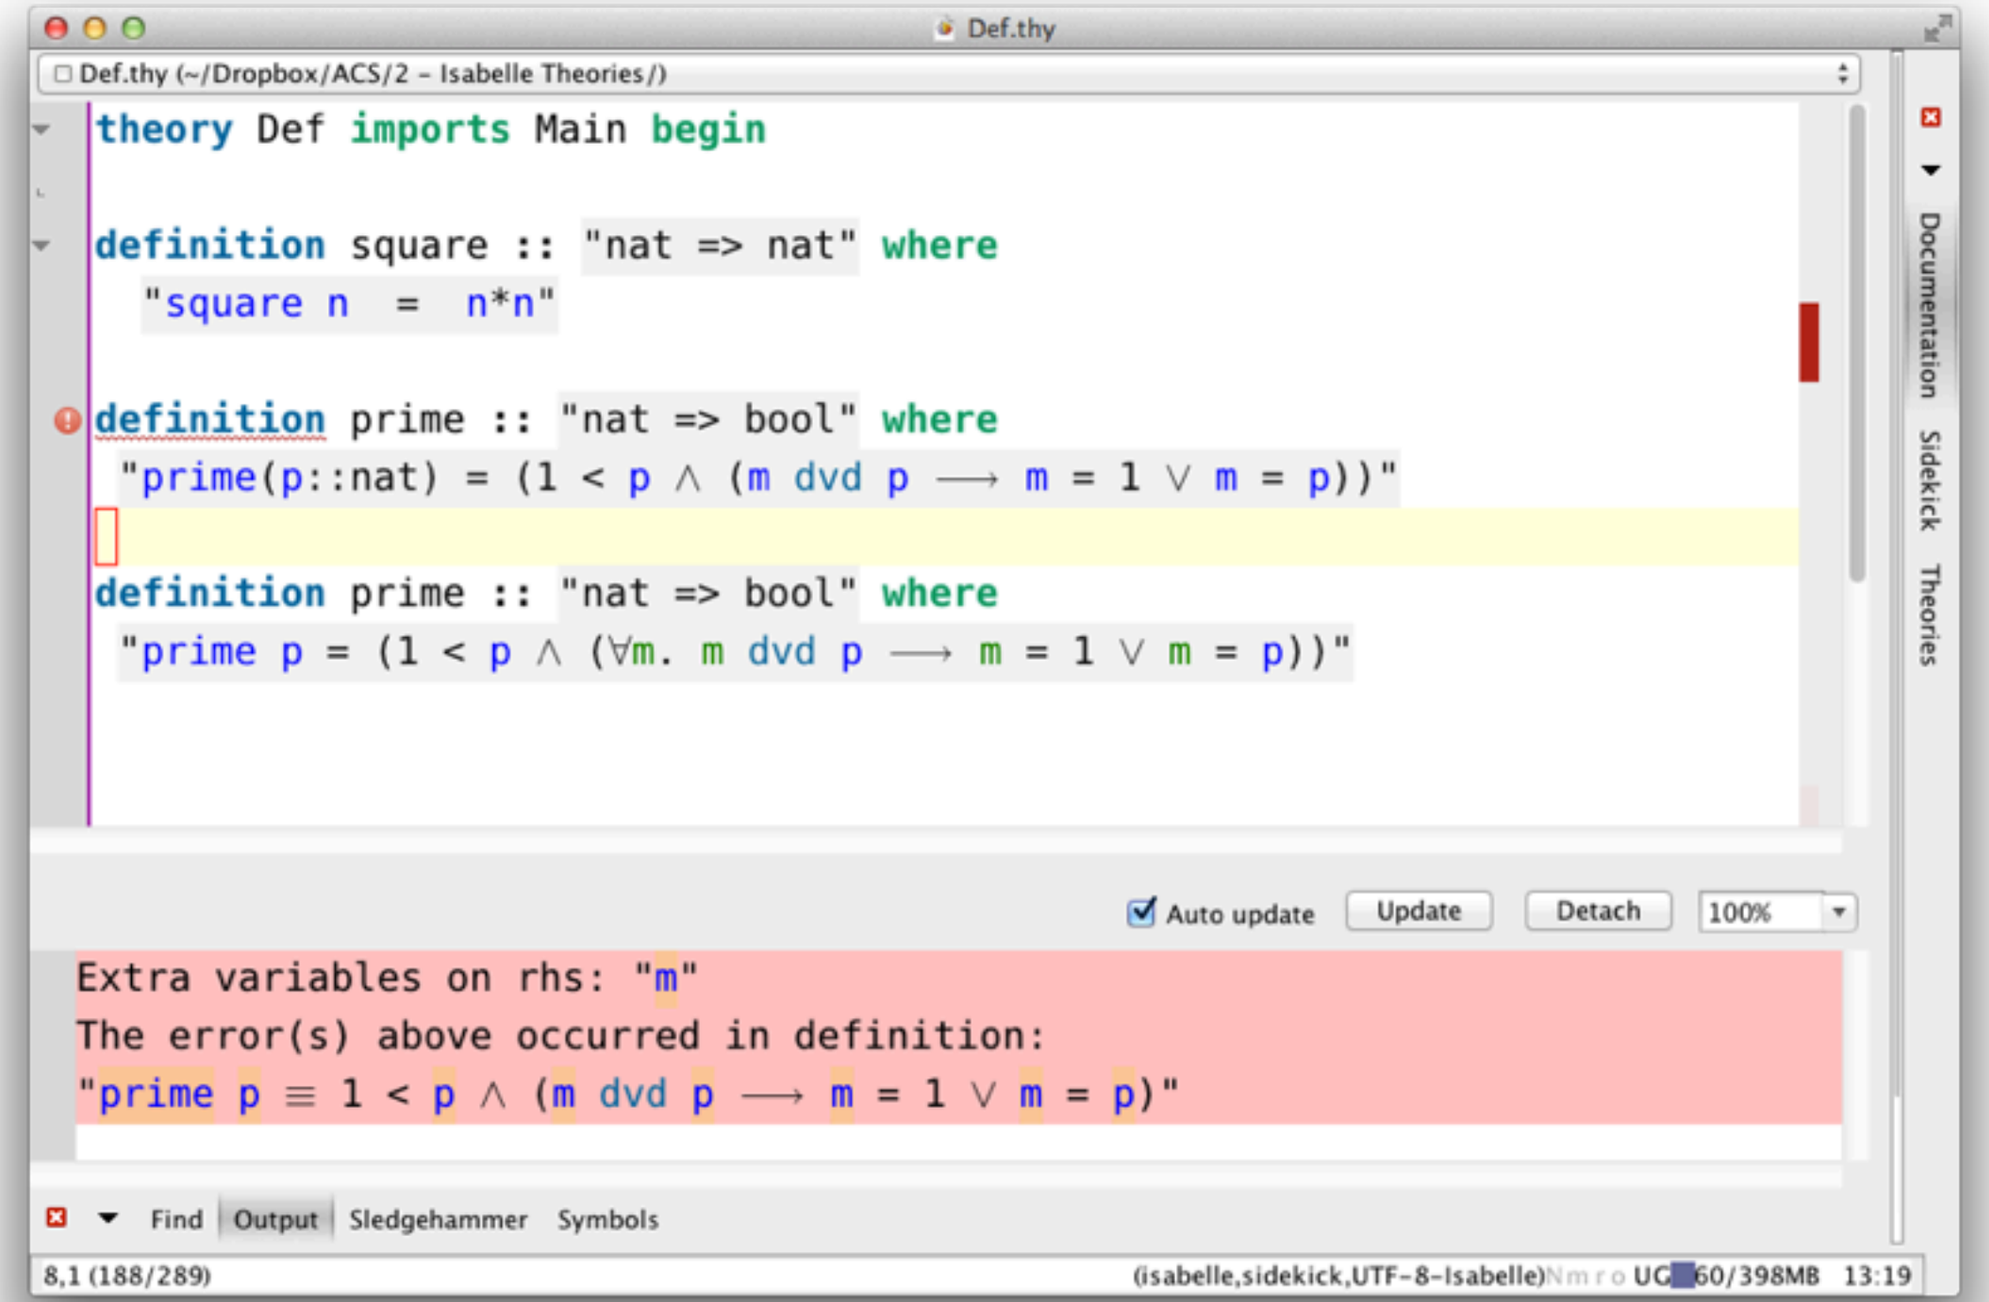

See the *Tutorial*, Section 2.7.2 Constant Definitions.

Notes on Constant Definitions

- Basic definitions are *not* recursive.
- Every variable on the right-hand side must also appear on the left.
- In proofs, definitions are *not* expanded by default!
  - Defining the constant *C* to denote *t* yields the theorem *C*`_def`, asserting *C*=*t*.
  - Abbreviations can be declared through a separate mechanism.

Lists in Isabelle

- We illustrate data types and functions using a reduced Isabelle theory (one without lists).
- The standard Isabelle environment has a *comprehensive list library*:
  - Functions # (cons), @ (append), `map`, `filter`, `nth`, `take`, `drop`, `takeWhile`, `dropWhile`, ...
  - Cases: (`case` xs `of` [] ⇒ [] | x#xs ⇒ ...)
  - Over 600 theorems!

List Induction Principle

To show φ(*xs*), it suffices to show the *base case* and *inductive step*:

- φ(Nil)
- φ(*xs*) ⇒ φ(Cons(*x*,*xs*))

The principle of case analysis is similar, expressing that any list has one of the forms Nil or Cons(*x*,*xs*) (for some *x* and *xs*).

Isabelle/jEdit

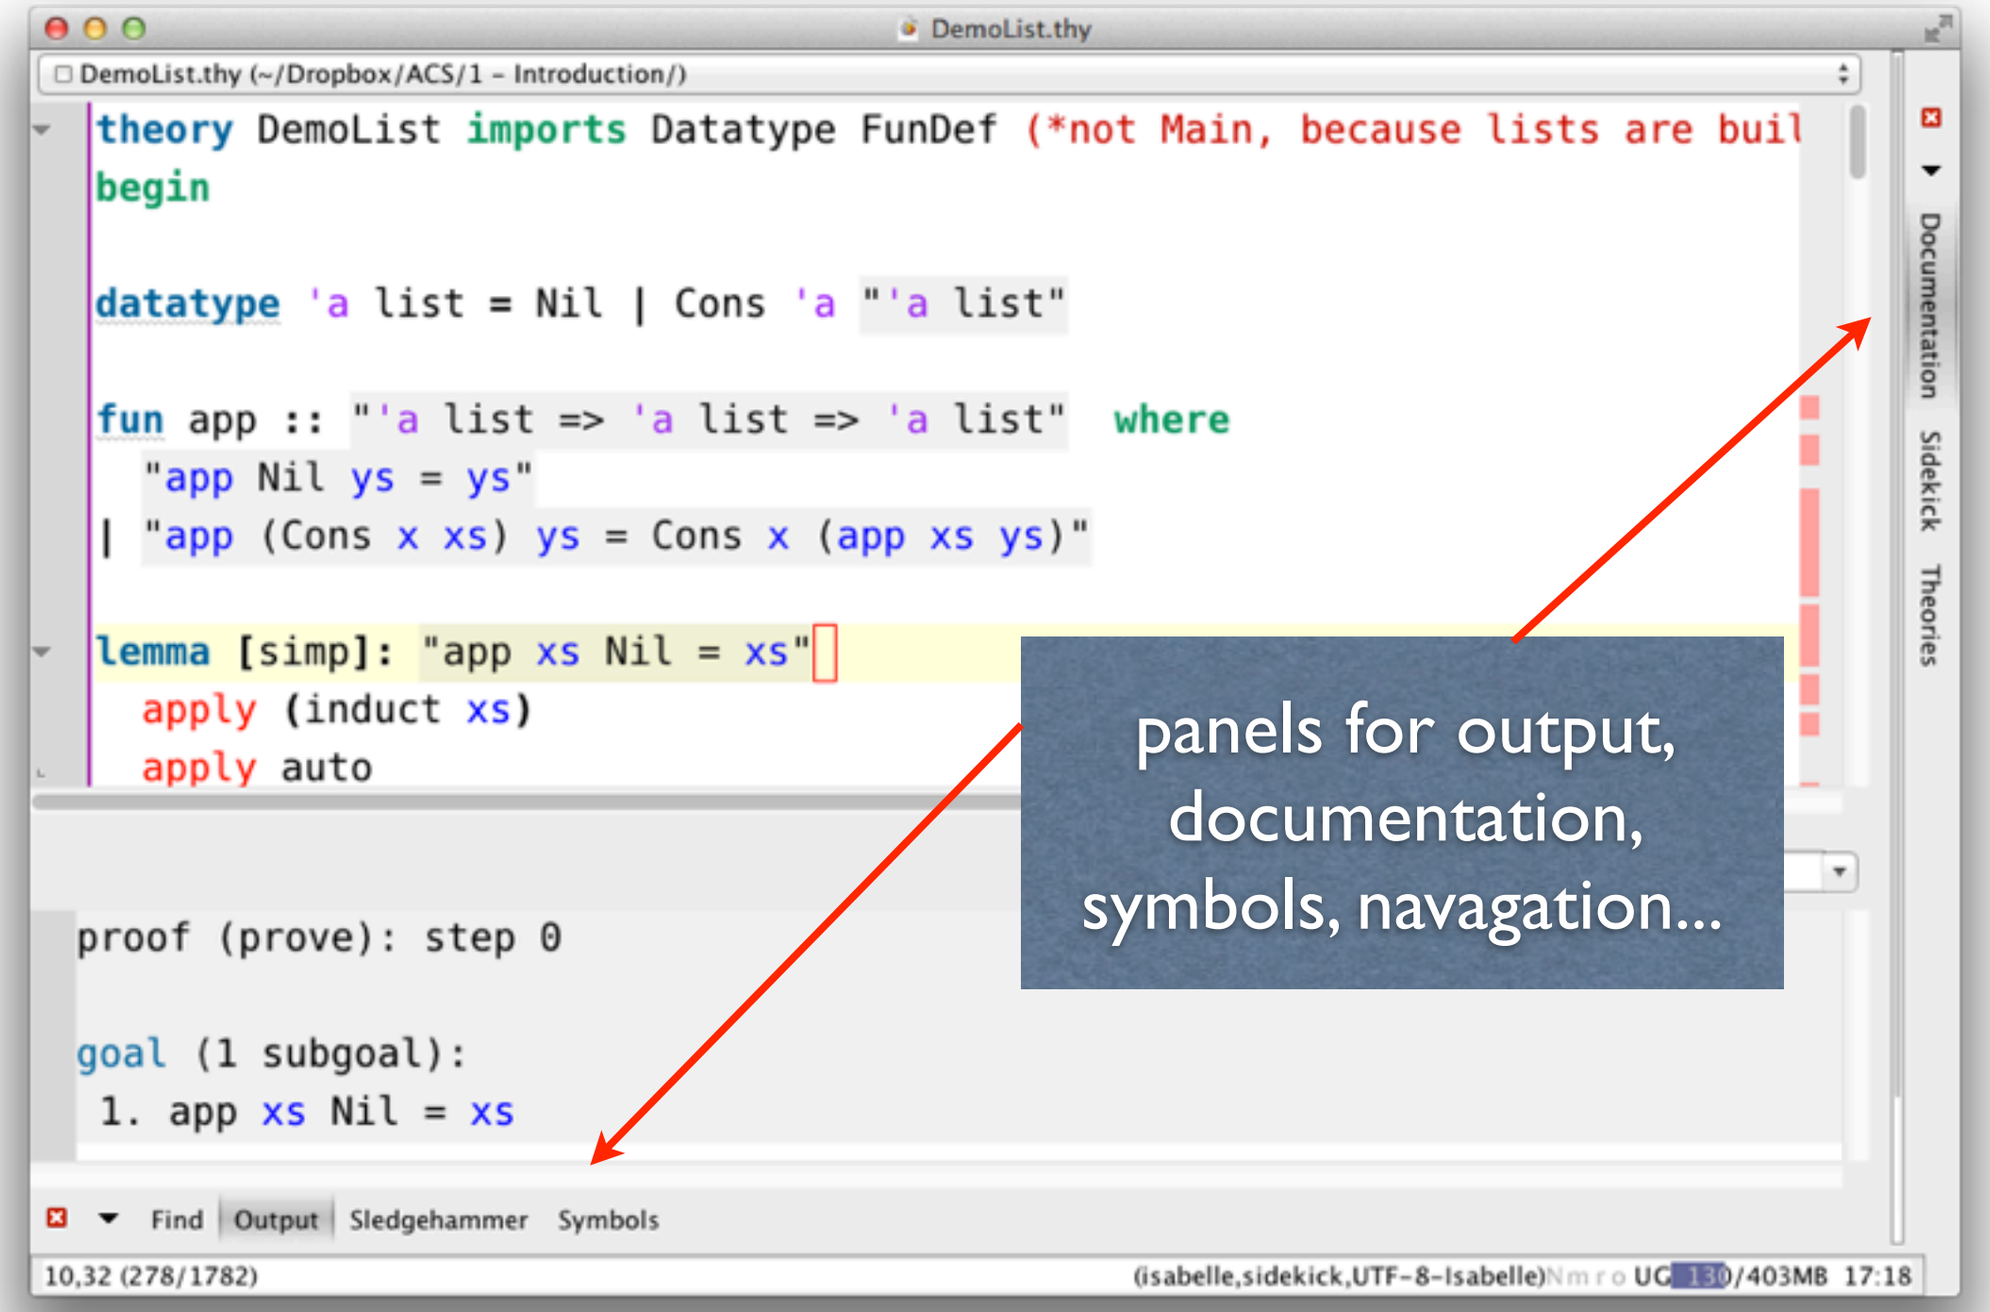

Isabelle's new user interface is the work of Makarius Wenzel. The entire proof document is processed as far as possible, errors and all.
It allows inspection of proof states, identifier declarations, etc.

All documentation is accessible from the sidebar.

Launch using the command "isabelle jedit FILENAME"

Proof General

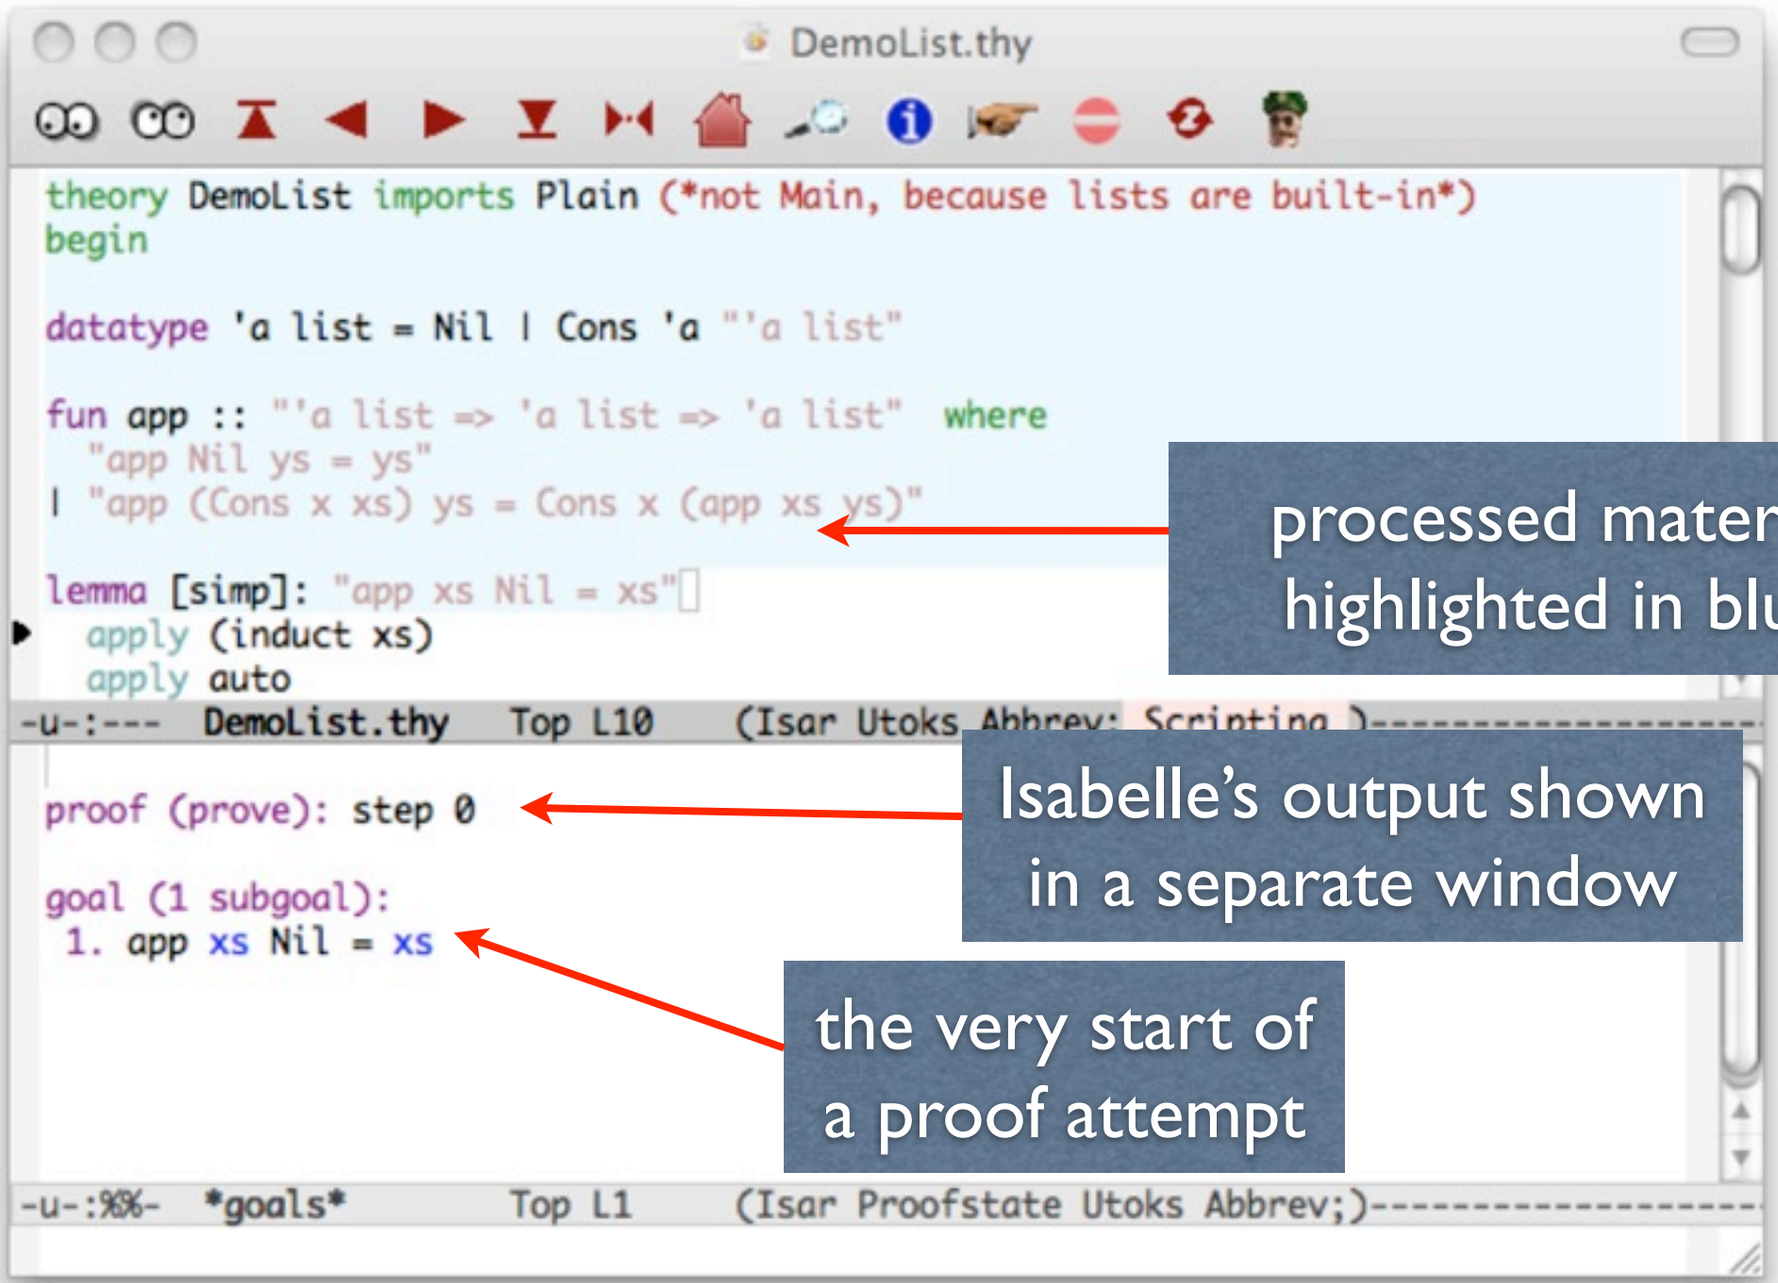

Isabelle's original user interface, Proof General, was developed by David Aspinall. It has a separate website: http://proofgeneral.inf.ed.ac.uk/

Proof General runs under Emacs, preferably version 23

Launch using the command "isabelle emacs FILENAME"

Proof by Induction

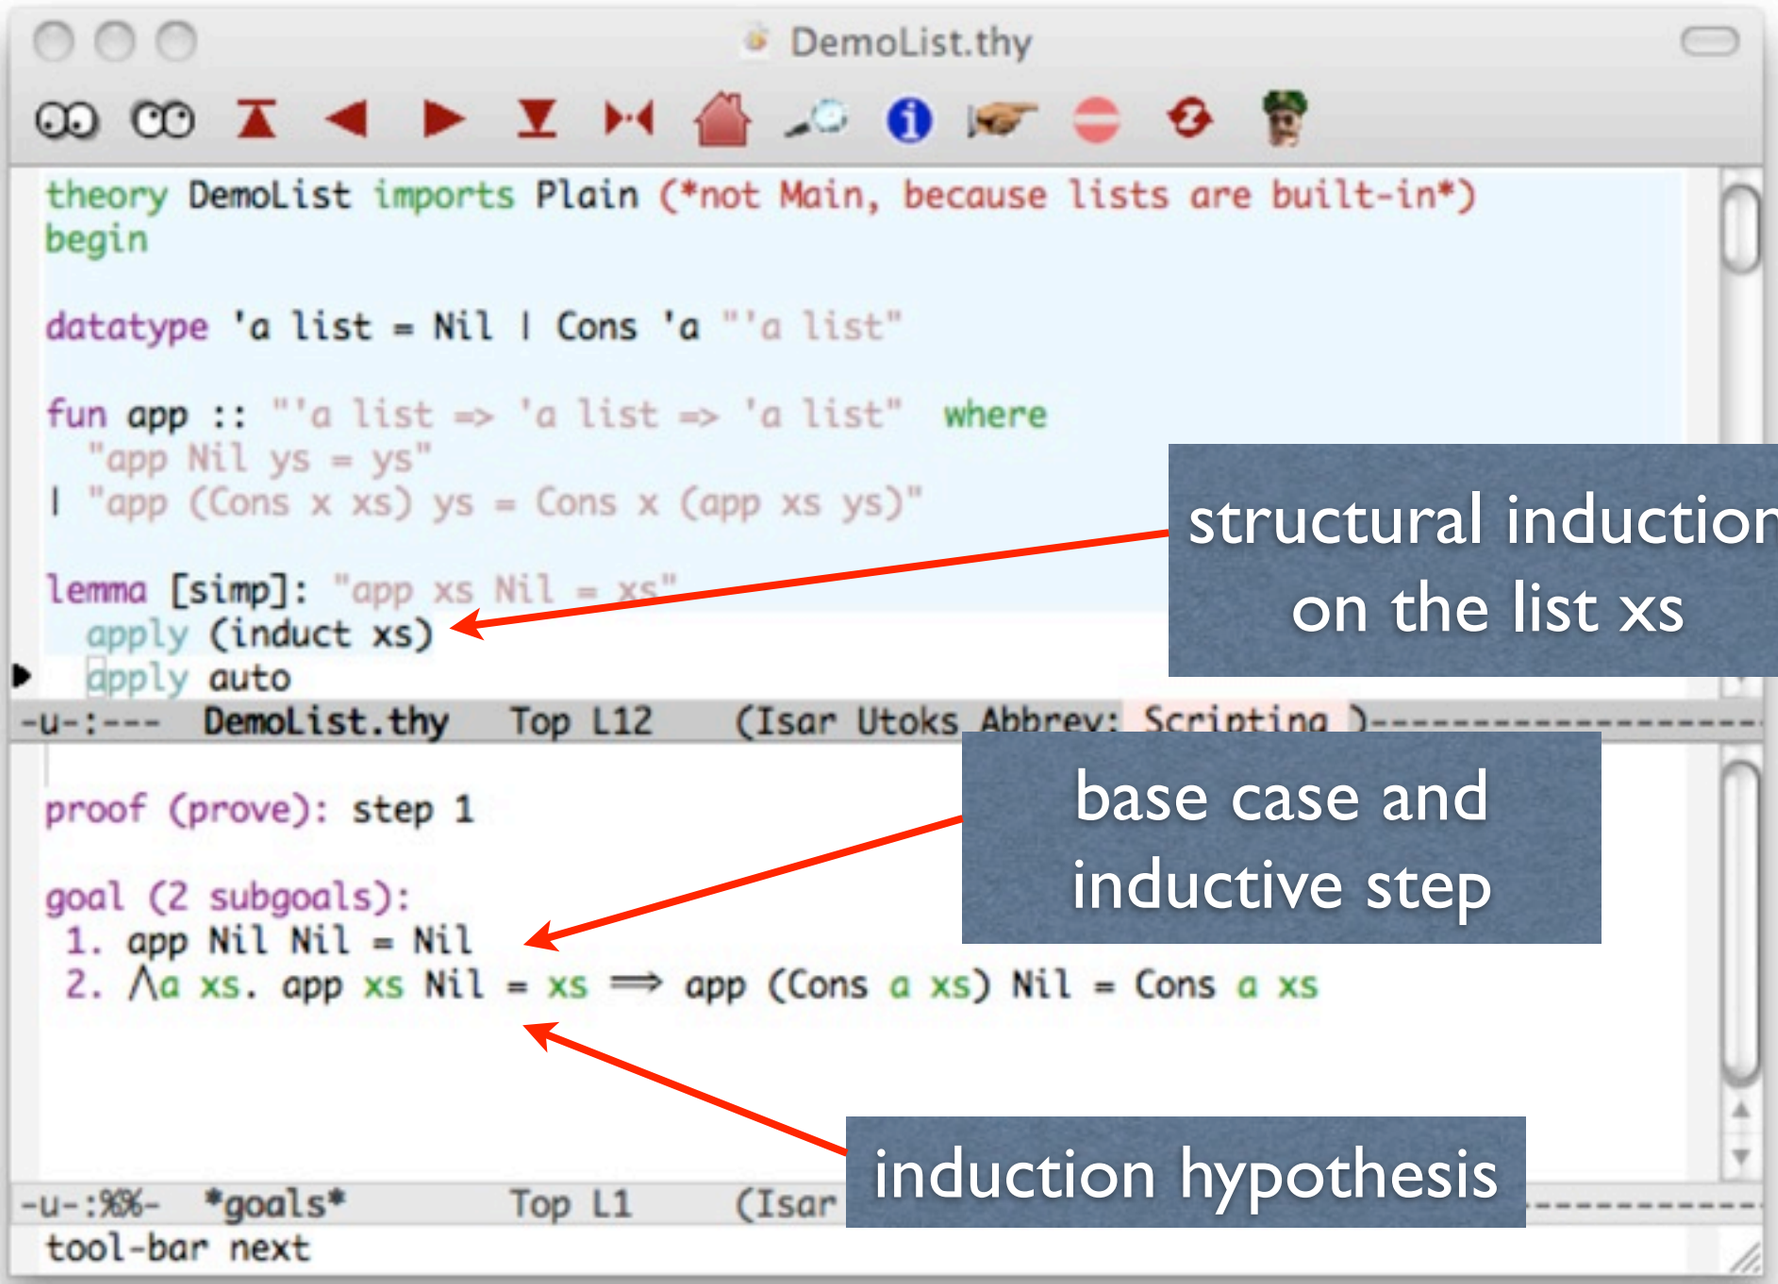

See the tutorial, section 2.3 (An Introductory Proof). For the moment, there is no important difference between induct_tac (used in the tutorial) and induct (used above). With both of these proof methods, you name an induction variable and it selects the corresponding structural induction rule, based on that variable's type. It then produces an instance of induction sufficient to prove the property in question.

Finishing a Proof

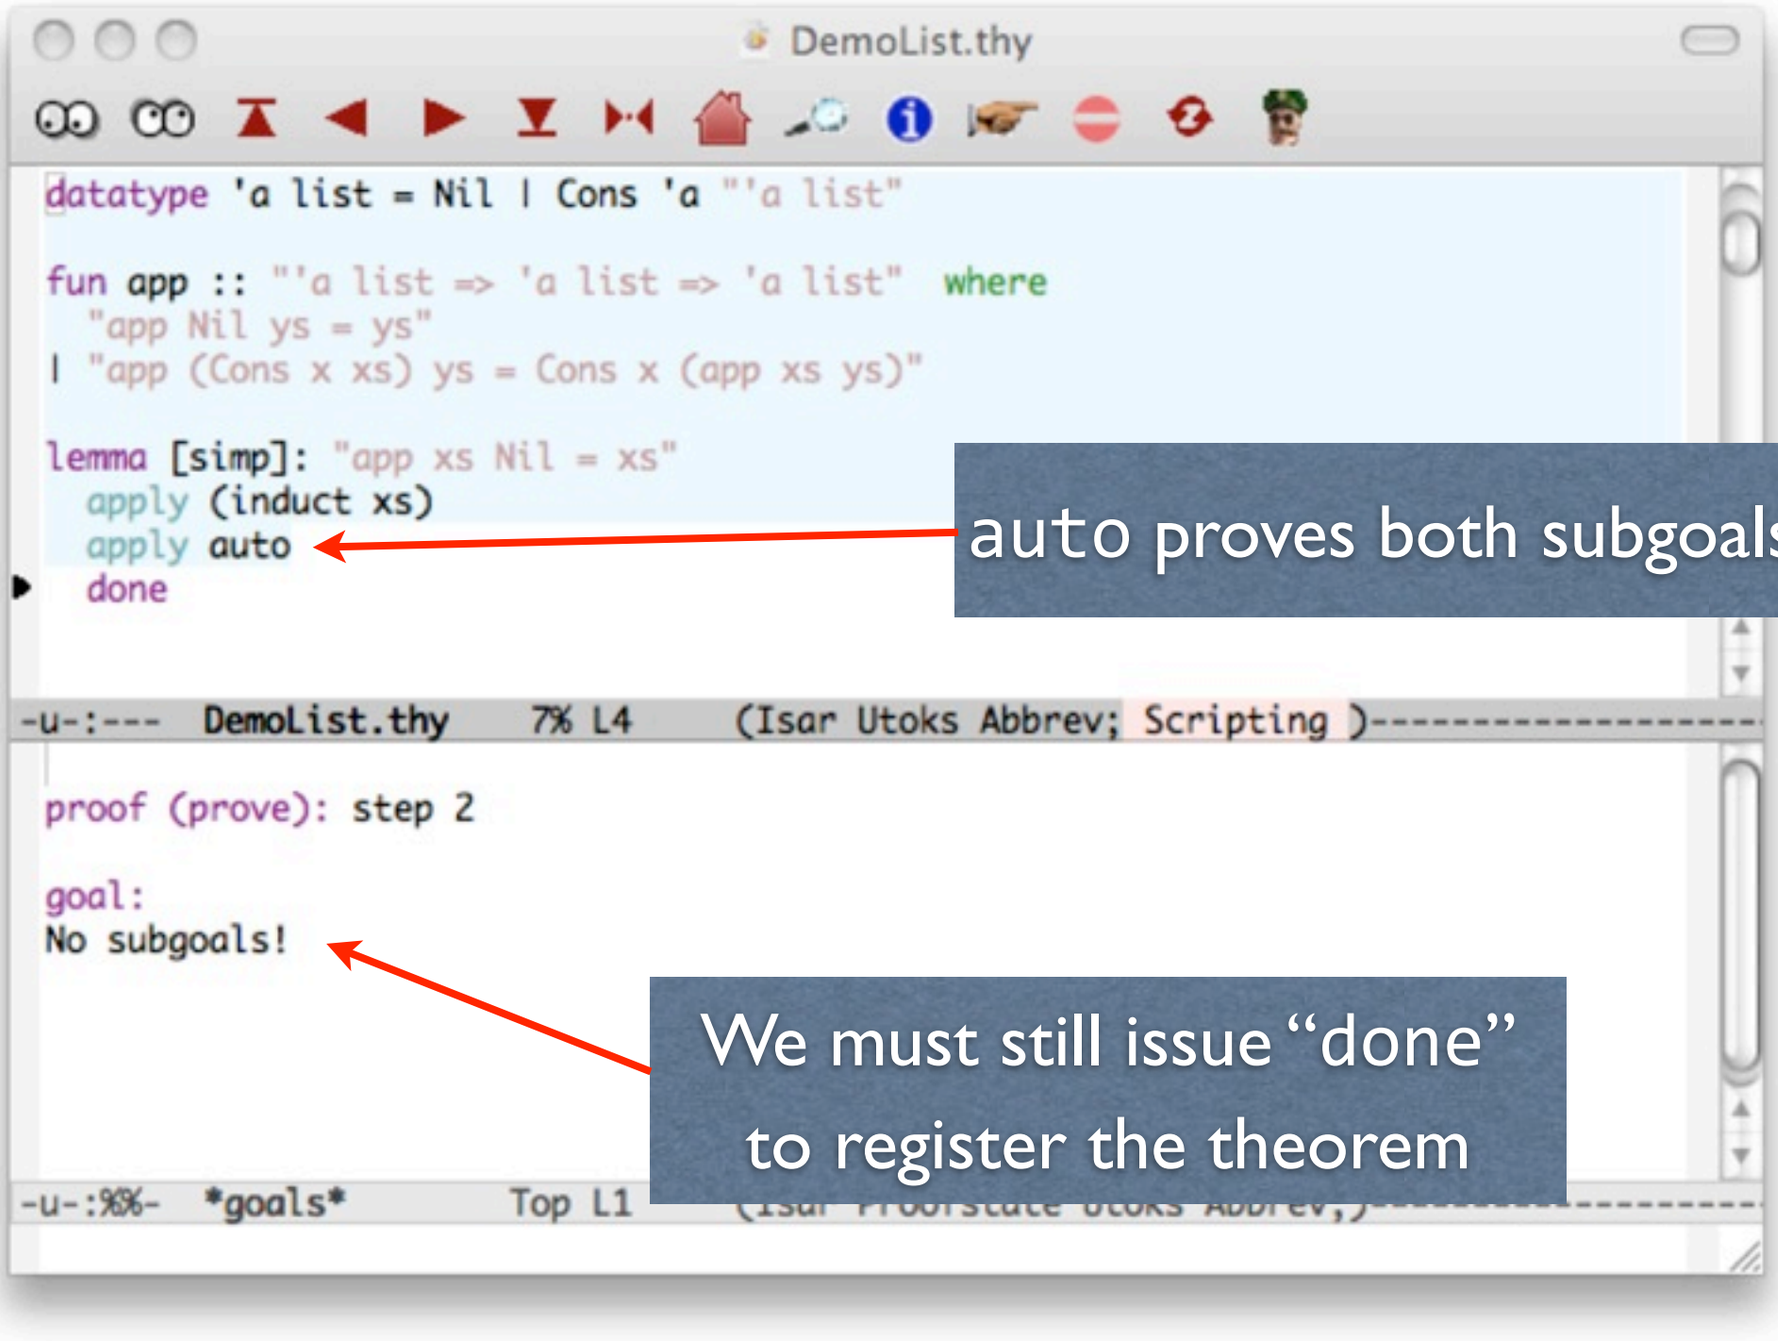

By default, Isabelle simplifies applications of recursive functions that match their defining recursion equations. This is quite different to the treatment of non-recursive definitions.

Another Proof Attempt

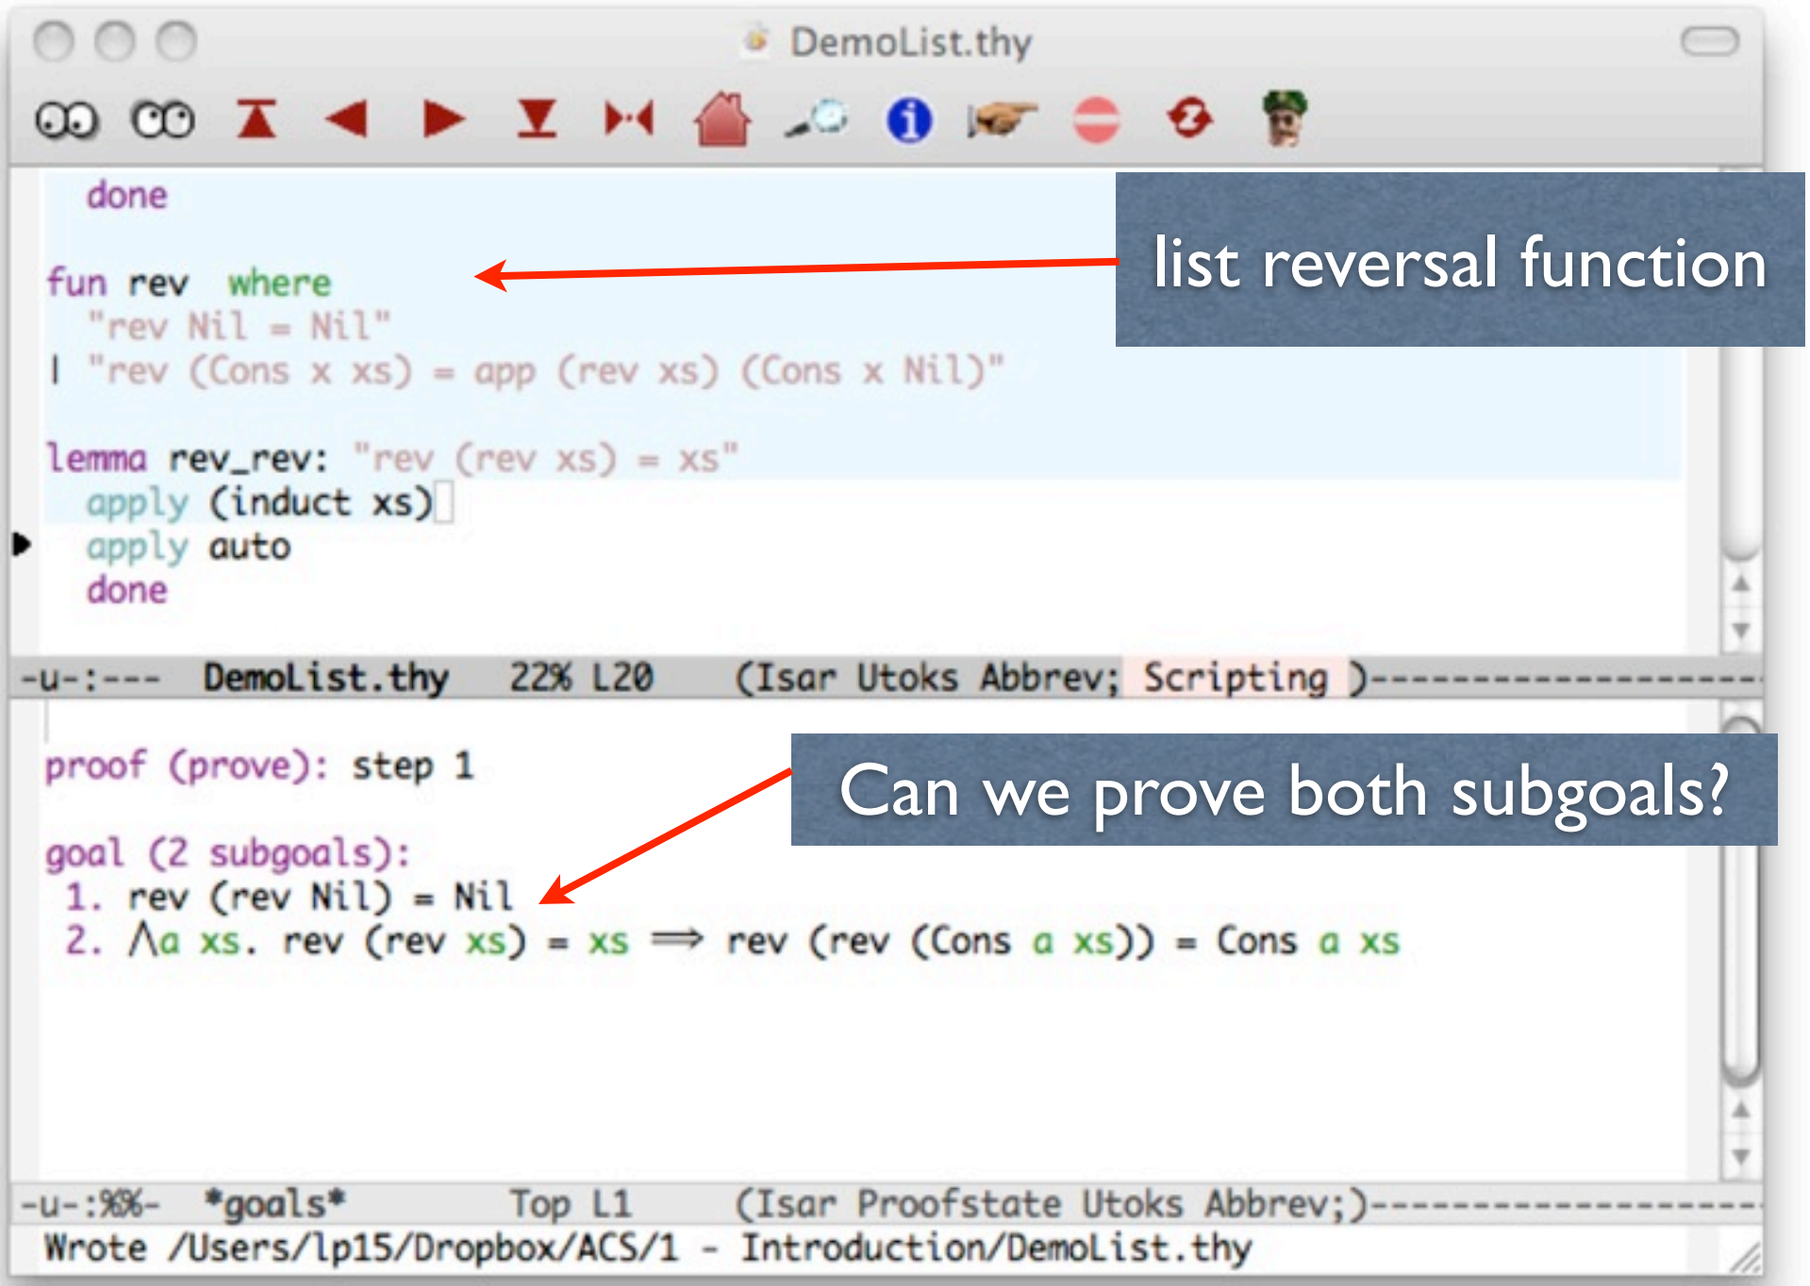

Stuck!

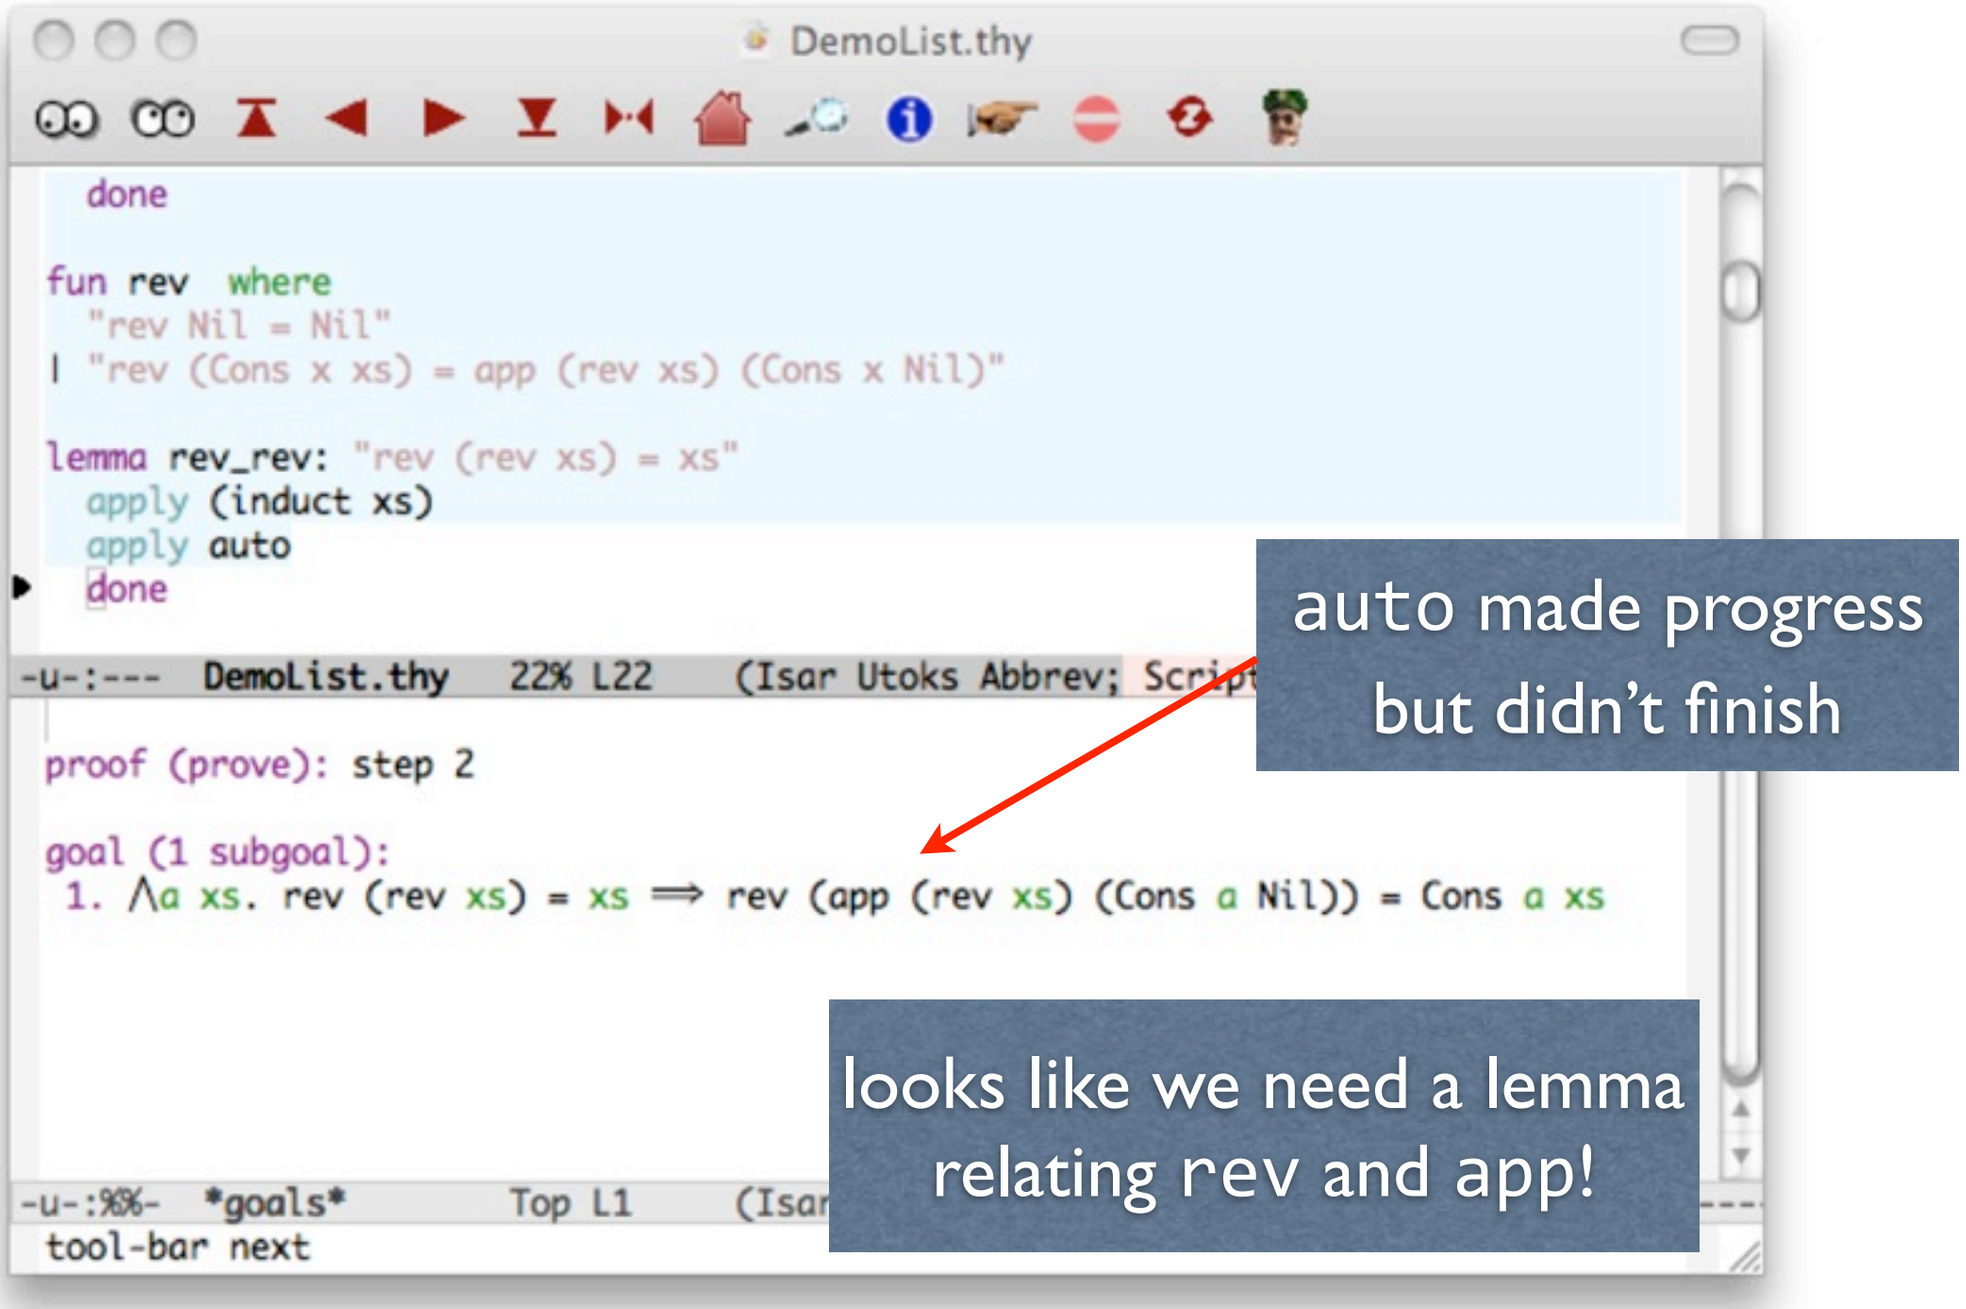
DemoList.thy
done
fun rev where
"rev Nil = Nil"
| "rev (Cons x xs) = app (rev xs) (Cons x Nil)"
lemma rev_rev: "rev (rev xs) = xs"
apply (induct xs)
apply auto
done
-u-:--- DemoList.thy 22% L22 (Isar Utoks Abbrev; Script
proof (prove): step 2
goal (1 subgoal):
1. ⋀a xs. rev (rev xs) = xs ⟹ rev (app (rev xs) (Cons a Nil)) = Cons a xs
auto made progress but didn't finish
looks like we need a lemma relating rev and app!
-u-:%%- *goals* Top L1 (Isar
tool-bar next

Stuck Again!

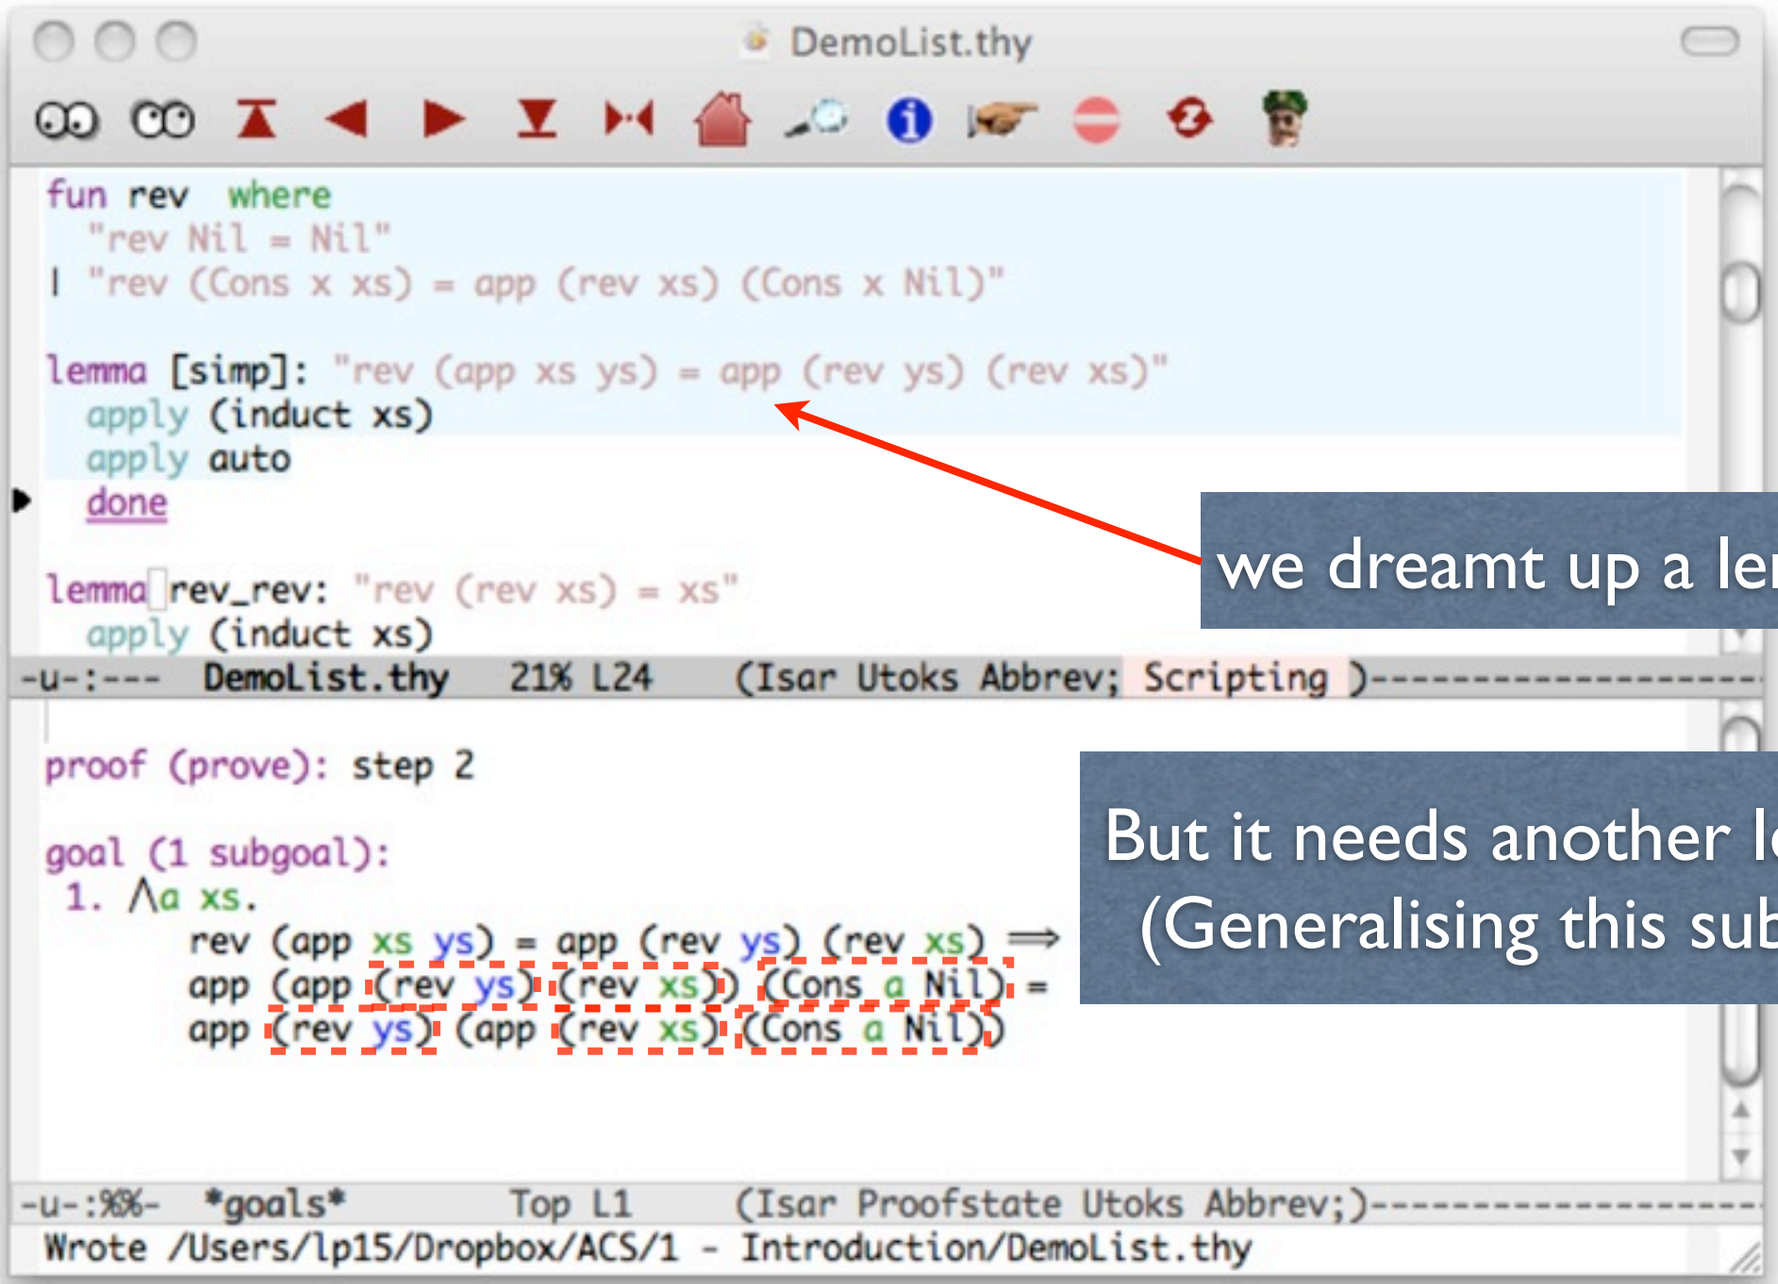

The subgoal that we cannot prove looks very complicated. But when we notice the repeated terms in it, we see that it is an instance of something simple and natural: the associativity of the function `app`. This fact does not involve the function `rev`! We see in this example how crucial it is to prove properties in the most abstract and general form.

The Final Piece of the Jigsaw

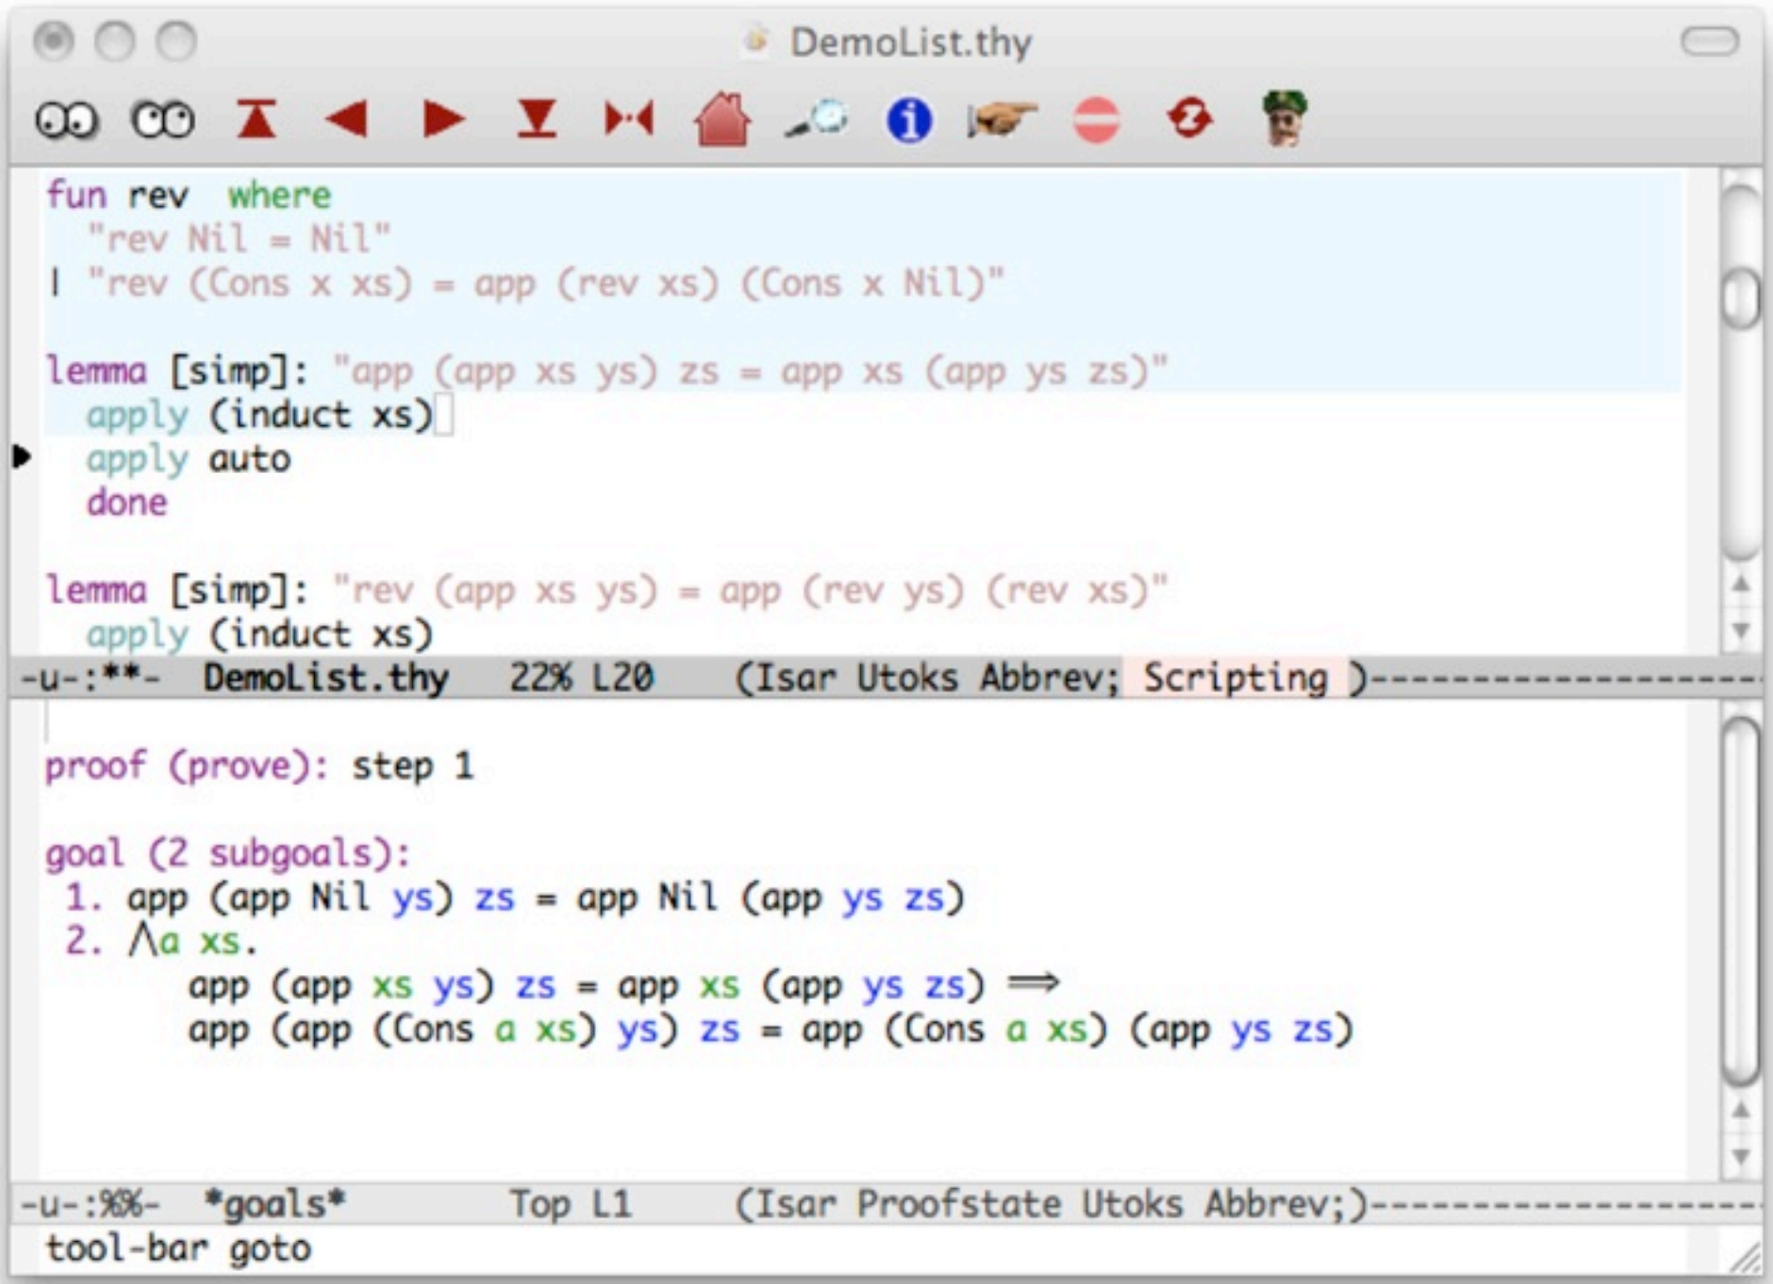

This proof of associativity will be successful, and with its help, the other lemmas are easily proved.

The Finished Proof

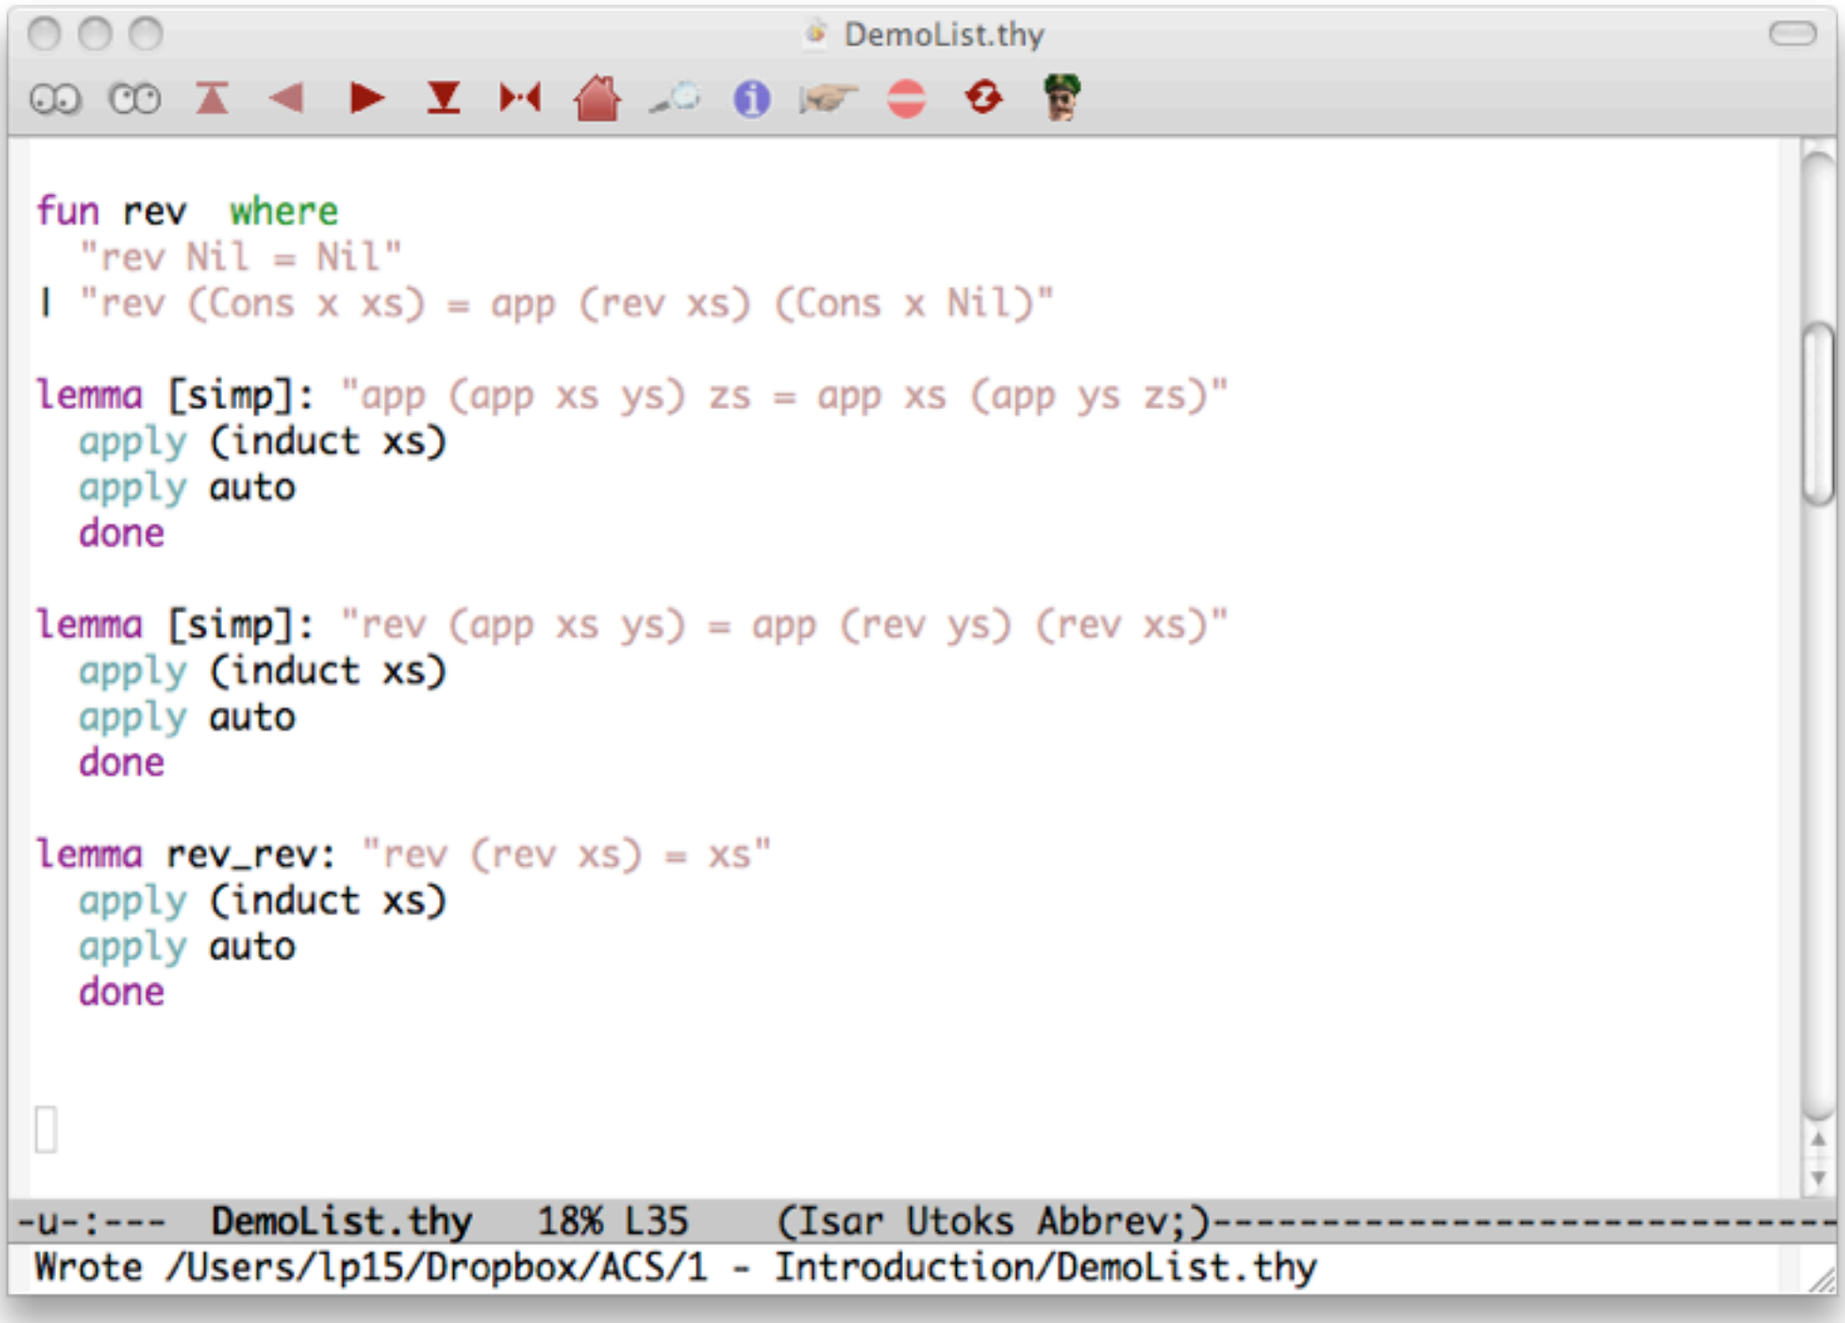

The lemmas must be proved in the correct order. Each is needed to prove the next.

It is actually more usable to give each lemma a name and to supply the relevant names to `auto`. The two lemmas proved above, especially the associativity of append, do not look like they would always be useful in simplification, so normally they would be proved without the `[simp]` attribute.

Interactive Formal Verification *3*: Elementary Proof

Lawrence C Paulson
Computer Laboratory
University of Cambridge

Elements of Interactive Proof

- Quite a few theorems can be proved by a combination of *induction* and *simplification*.
- Induction can be a straightforward *structural induction* rule derived from a type declaration, but other induction rules are quite specialised.
- Simplification typically refers to *rewriting* according to the definition of a recursive function…
- but it has many refinements, including automatic *case splitting*, simple logical reasoning and sophisticated *arithmetic* reasoning.

Goals and Subgoals

- We start with one subgoal: the statement to be proved.
- Proof *tactics* and *methods* typically replace a single subgoal by zero or more new subgoals.
  - But certain methods, notably `auto` and `simp_all`, operate on *all* outstanding subgoals.
- We finish when no subgoals remain. *The theorem is proved!*

See the Tutorial, 2.3 An Introductory Proof. The list of subgoals is always flat. Towards the end of this course, there is a brief introduction to structured proofs.

Structure of a Subgoal

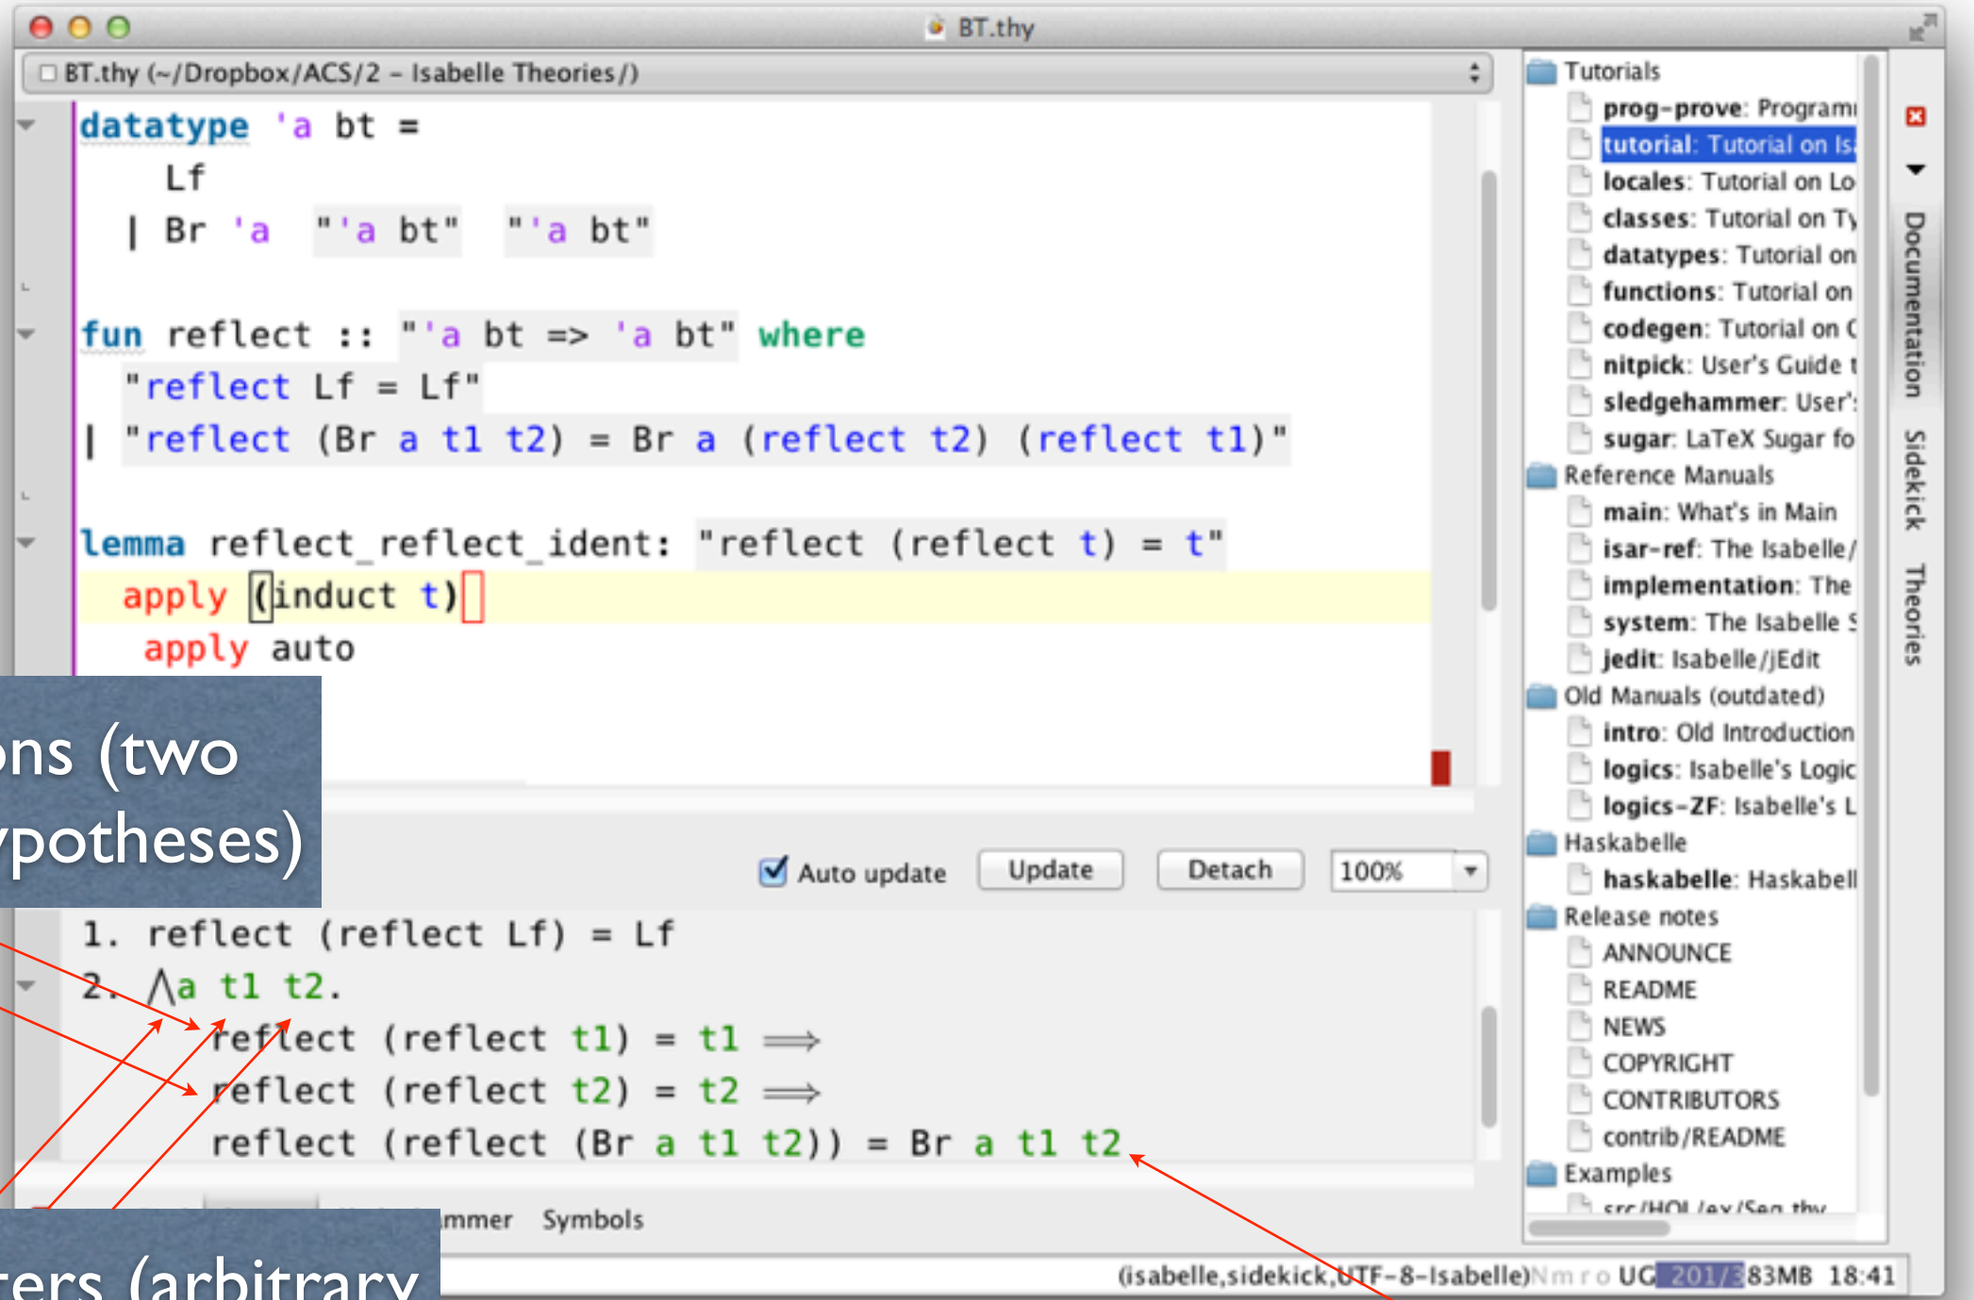

assumptions (two induction hypotheses)

parameters (arbitrary local variables)

conclusion

Proof by Rewriting

```
app (Cons x xs) ys  ⇒ Cons x (app xs ys)
   rev (Cons x xs)  ⇒ app (rev xs) (Cons x Nil)
   rev (app xs ys)  ⇒ app (rev ys) (rev xs)
app (app xs ys) zs  ⇒ app xs (app ys zs)
```

recursive defns

induction hyp

lemma

```
rev (app (Cons a xs) ys) = app (rev ys) (rev (Cons a xs))
```

```
rev (app (Cons a xs) ys) =
rev (Cons a (app xs ys)) =
app (rev (app xs ys)) (Cons a Nil) =
app (app (rev ys) (rev xs)) (Cons a Nil) =
app (rev ys) (app (rev xs) (Cons a Nil))
```

```
app (rev ys) (rev (Cons a xs)) =
app (rev ys) (app (rev xs) (Cons a Nil))
```

At each step, the highlighted term is rewritten to something else. Eventually, the left hand side and right hand side of the desired equation have become equal. (This equation is the induction step for our lemma, rev (app xs ys) = app (rev ys) (rev xs).)

Rewriting with Equivalences

```
(x dvd -y) = (x dvd y)
(a * b = 0) = (a = 0 ∨ b = 0)
(A - B ⊆ C) = (A ⊆ B ∪ C)
(a*c ≤ b*c) = ((0<c → a ≤ b) ∧ (c<0 → b ≤ a))
```

introduces a case split on the sign of c

- Logical equivalencies are just boolean equations.
- They lead to a clear and simple proof style.
- They can also be written with the syntax $P \leftrightarrow Q$.

Automatic Case Splitting

Simplification will replace

$P(\texttt{if}\ b\ \texttt{then}\ x\ \texttt{else}\ y)$

by

$$(b \to P(x)) \wedge (\neg b \to P(y))$$

- By default, this only happens when simplifying the conclusion. But assumptions can also be split.
- Other kinds of case splitting can be enabled.

See the *Tutorial*, 3.1.9 Automatic Case Splits

Conditional Rewrite Rules

```
xs ≠ [] ⇒ hd (xs @ ys) = hd xs
```

```
n ≤ m ⇒ (Suc m) - n = Suc (m - n)
```

```
[|a ≠ 0; b ≠ 0|] ⇒ b / (a*b) = 1 / a
```

- *First* match the left-hand side, *then* **recursively** prove the conditions by simplification.
- If successful, applying the resulting rewrite rule.

Termination Issues

- *Looping*: `f(x) = h(g(x))`, `g(x) = f(x+2)`
- *Looping*: `P(x) ⇒ x=0`
  - `simp` will try to use this rule to simplify its own precondition!
- `x+y = y+x` is actually okay!
  - *Permutative rewrite rules* are applied but only if they make the term "lexicographically smaller".

The Methods `simp` and `auto`

- `simp` performs *rewriting* (along with simple arithmetic simplification) on the *first* subgoal
- `auto` simplifies *all subgoals*, not just the first.
- `auto` also applies all obvious *logical steps*
  - Splitting conjunctive goals and disjunctive assumptions
  - Performing obvious quantifier removal

See the *Tutorial*, 3.1 Simplification. This section describes the options and possibilities thoroughly.

Variations on `simp` and `auto`

```
simp add: add_assoc

simp del: rev_rev (no_asm_simp)

simp (no_asm)

simp_all (no_asm_simp) add: … del: …

auto simp add: … del: …
```

using another rewrite rule

omitting a certain rule

not simplifying the assumptions

ignoring all assumptions

do `simp` for all subgoals

`auto` with options

Rules for Arithmetic

- An identifier can denote a *list* of lemmas.
- `add_ac` and `mult_ac`: associative/commutative properties of addition and multiplication
- `algebra_simps`: useful for multiplying out polynomials
- `field_simps`: useful for multiplying out the denominators when proving inequalities

  Example: `auto simp add: field_simps`

These identifiers denote lists of theorems that work together well as rewrite rules for performing various simplification tasks.

Basics of Proof by Induction

- State the desired theorem using "`lemma`", with its name and optionally `[simp]`
- Identify the *induction variable*
  - Its type should be some datatype (incl. `nat`)
    - It should appear as the *argument of a recursive function*.
- Complicating issues include unusual recursions and auxiliary variables.

Completing the Proof

- Apply "`induct`" with the chosen variable.
- The first subgoal will be the base case, and it should be trivial using "`simp`".
- Other subgoals will involve induction hypotheses and the proof of each may require several steps.
- Naturally, the first thing to try is "`auto`", but much more is possible.

Basics of Isabelle/jEdit

- Based on the jEdit text editor.
- Isabelle automatically processes the entire visible window, errors and all, using parallel threads.
- Identifiers and other elements can be *inspected* using hover-click.
- *Dockable panels* give access to the Isabelle output, theory structure, manuals, symbols, etc.

A View of Isabelle/jEdit

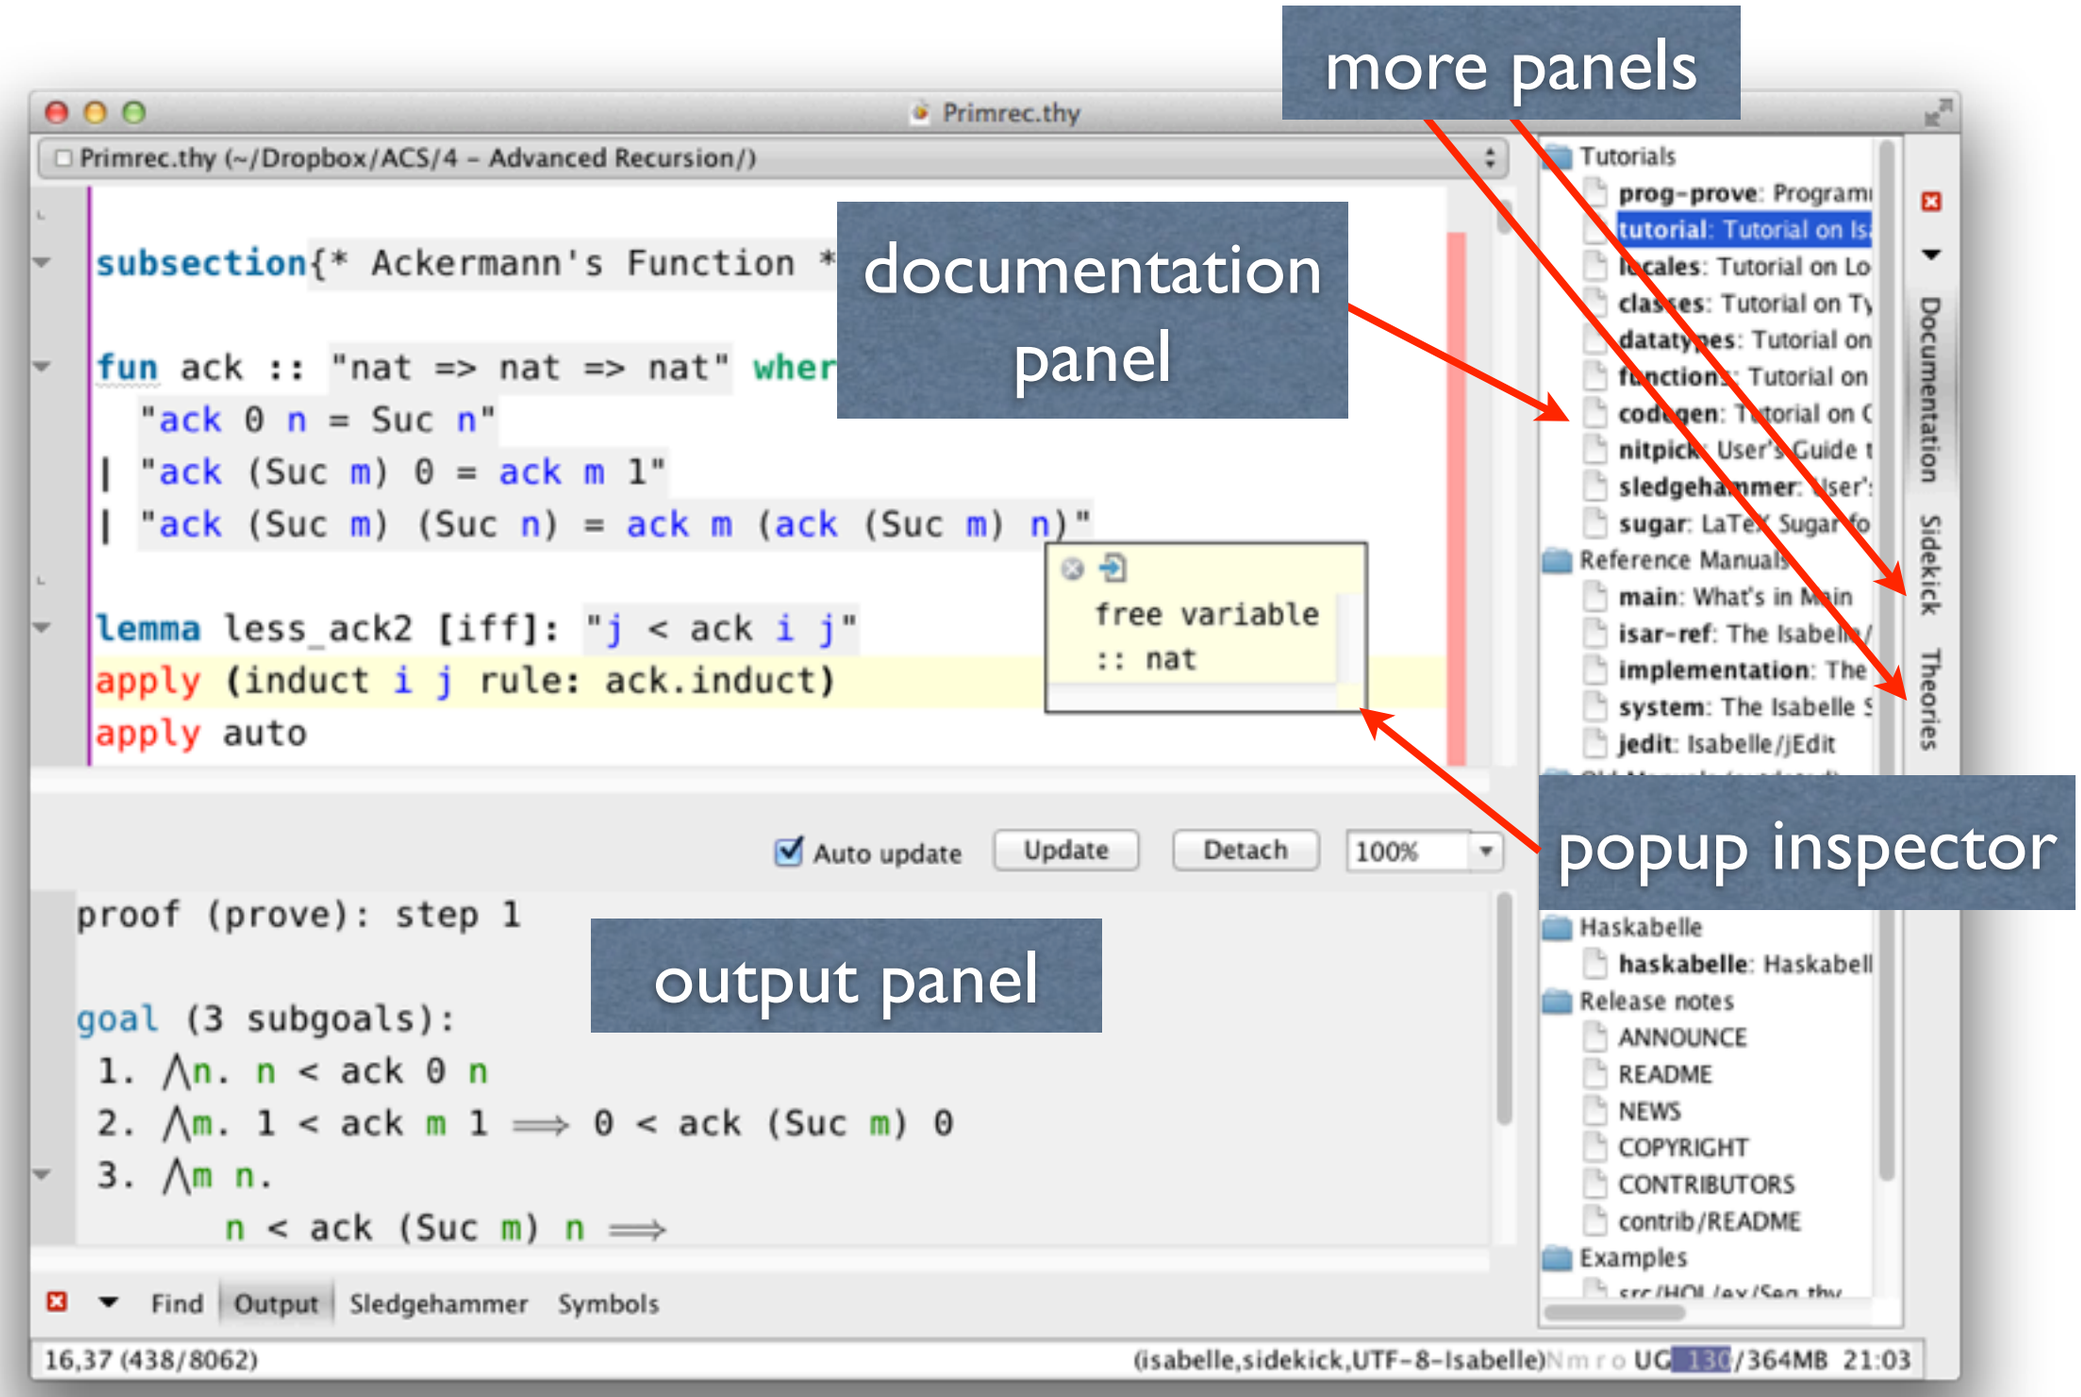

Interactive Formal Verification
4: Advanced Recursion, Induction and Simplification

Lawrence C Paulson
Computer Laboratory
University of Cambridge

Why does Induction Fail?

In a formal proof—like in a program—even trivial errors can be fatal. Everything must be set up *exactly* right…

- The statement being proved is too weak, so the induction hypothesis is too weak.
- You have chosen an inappropriate induction rule for this proof.
- Or maybe you just don't know how to make use of the induction hypotheses.

A Failing Proof by Induction

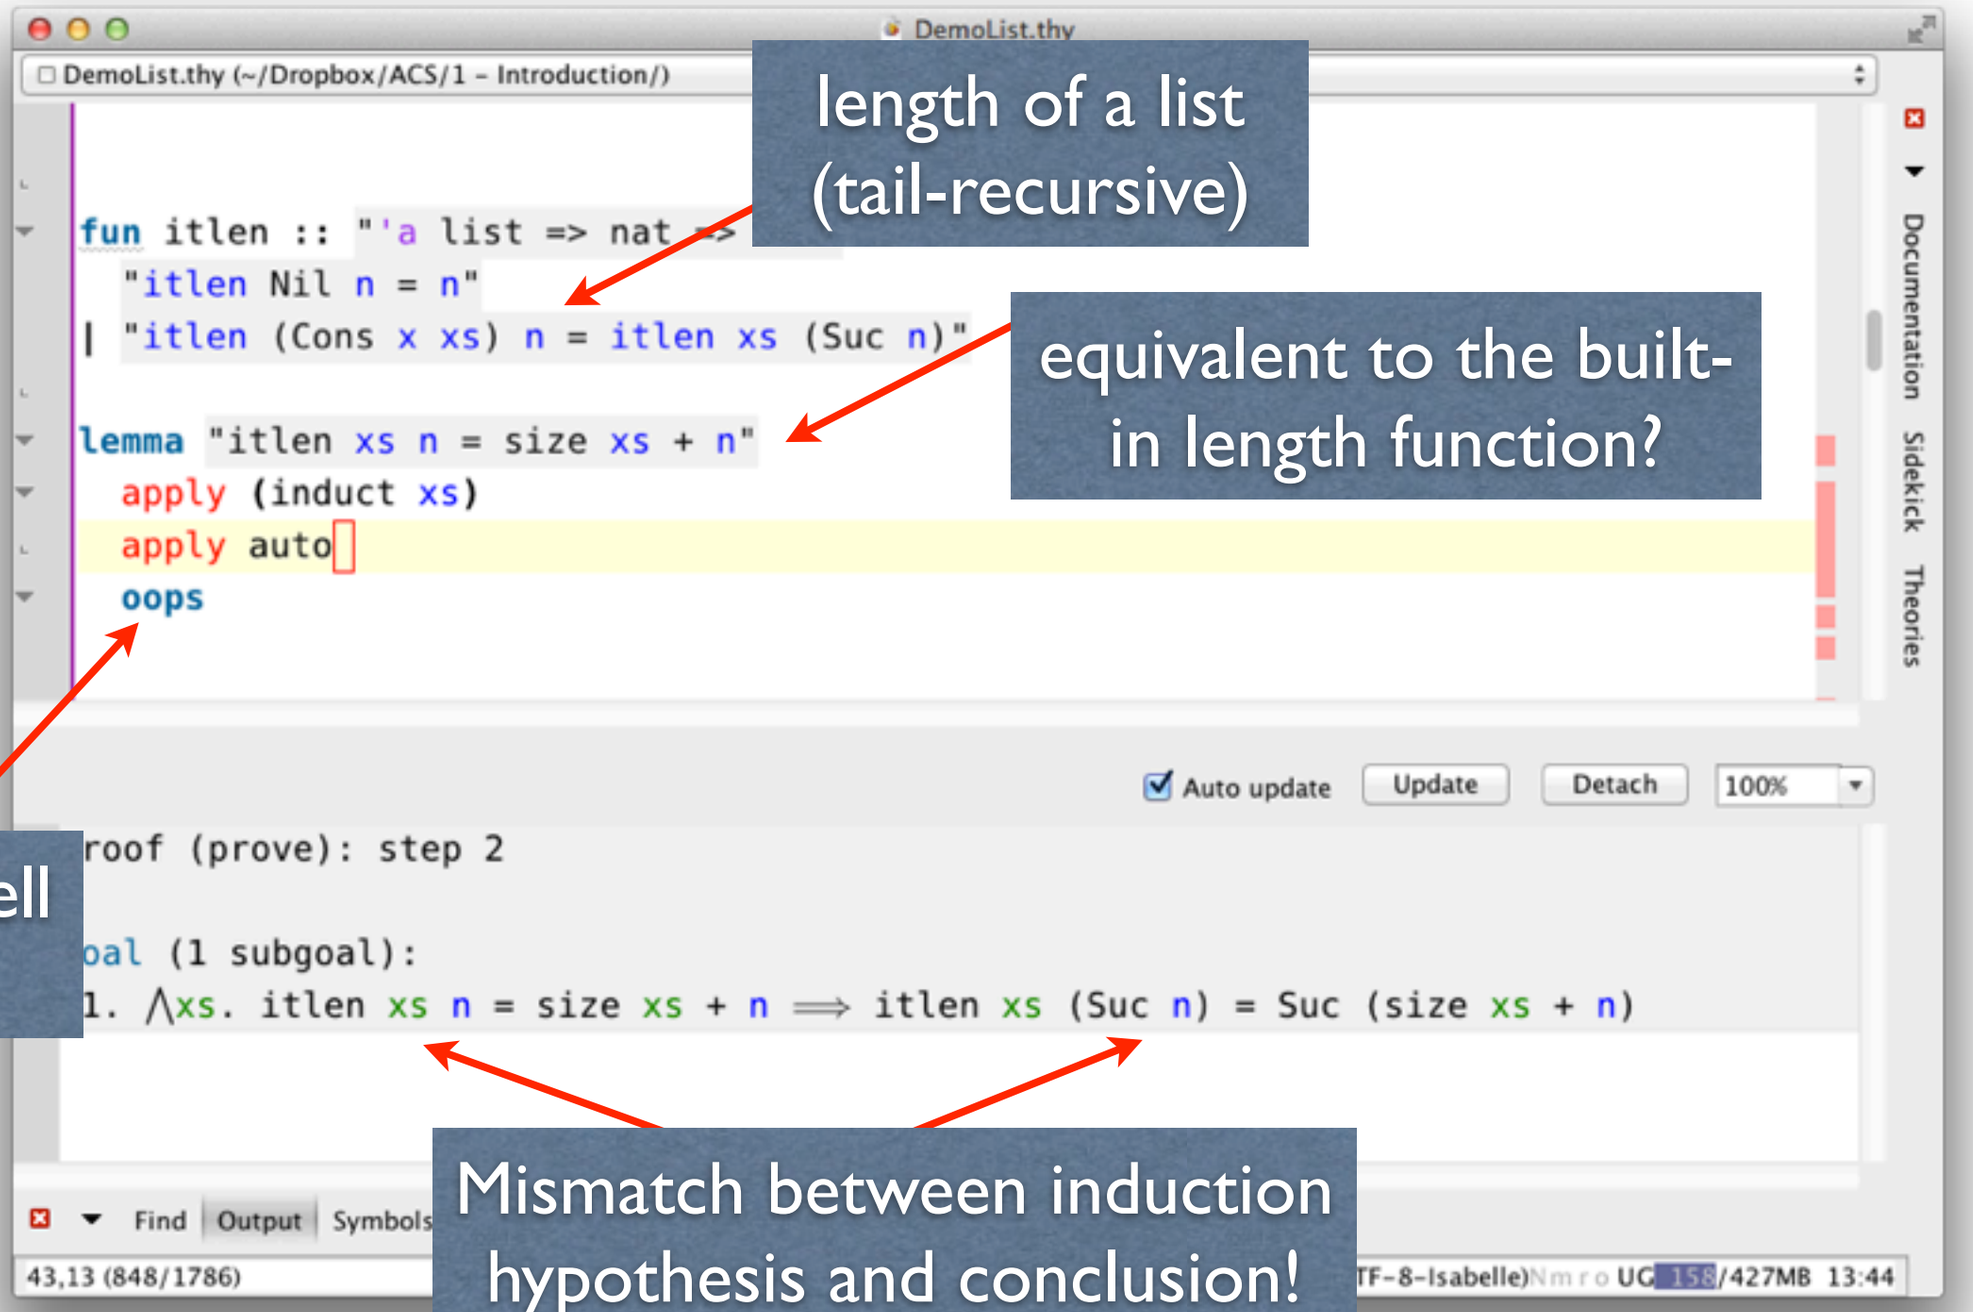

Generalising the Induction

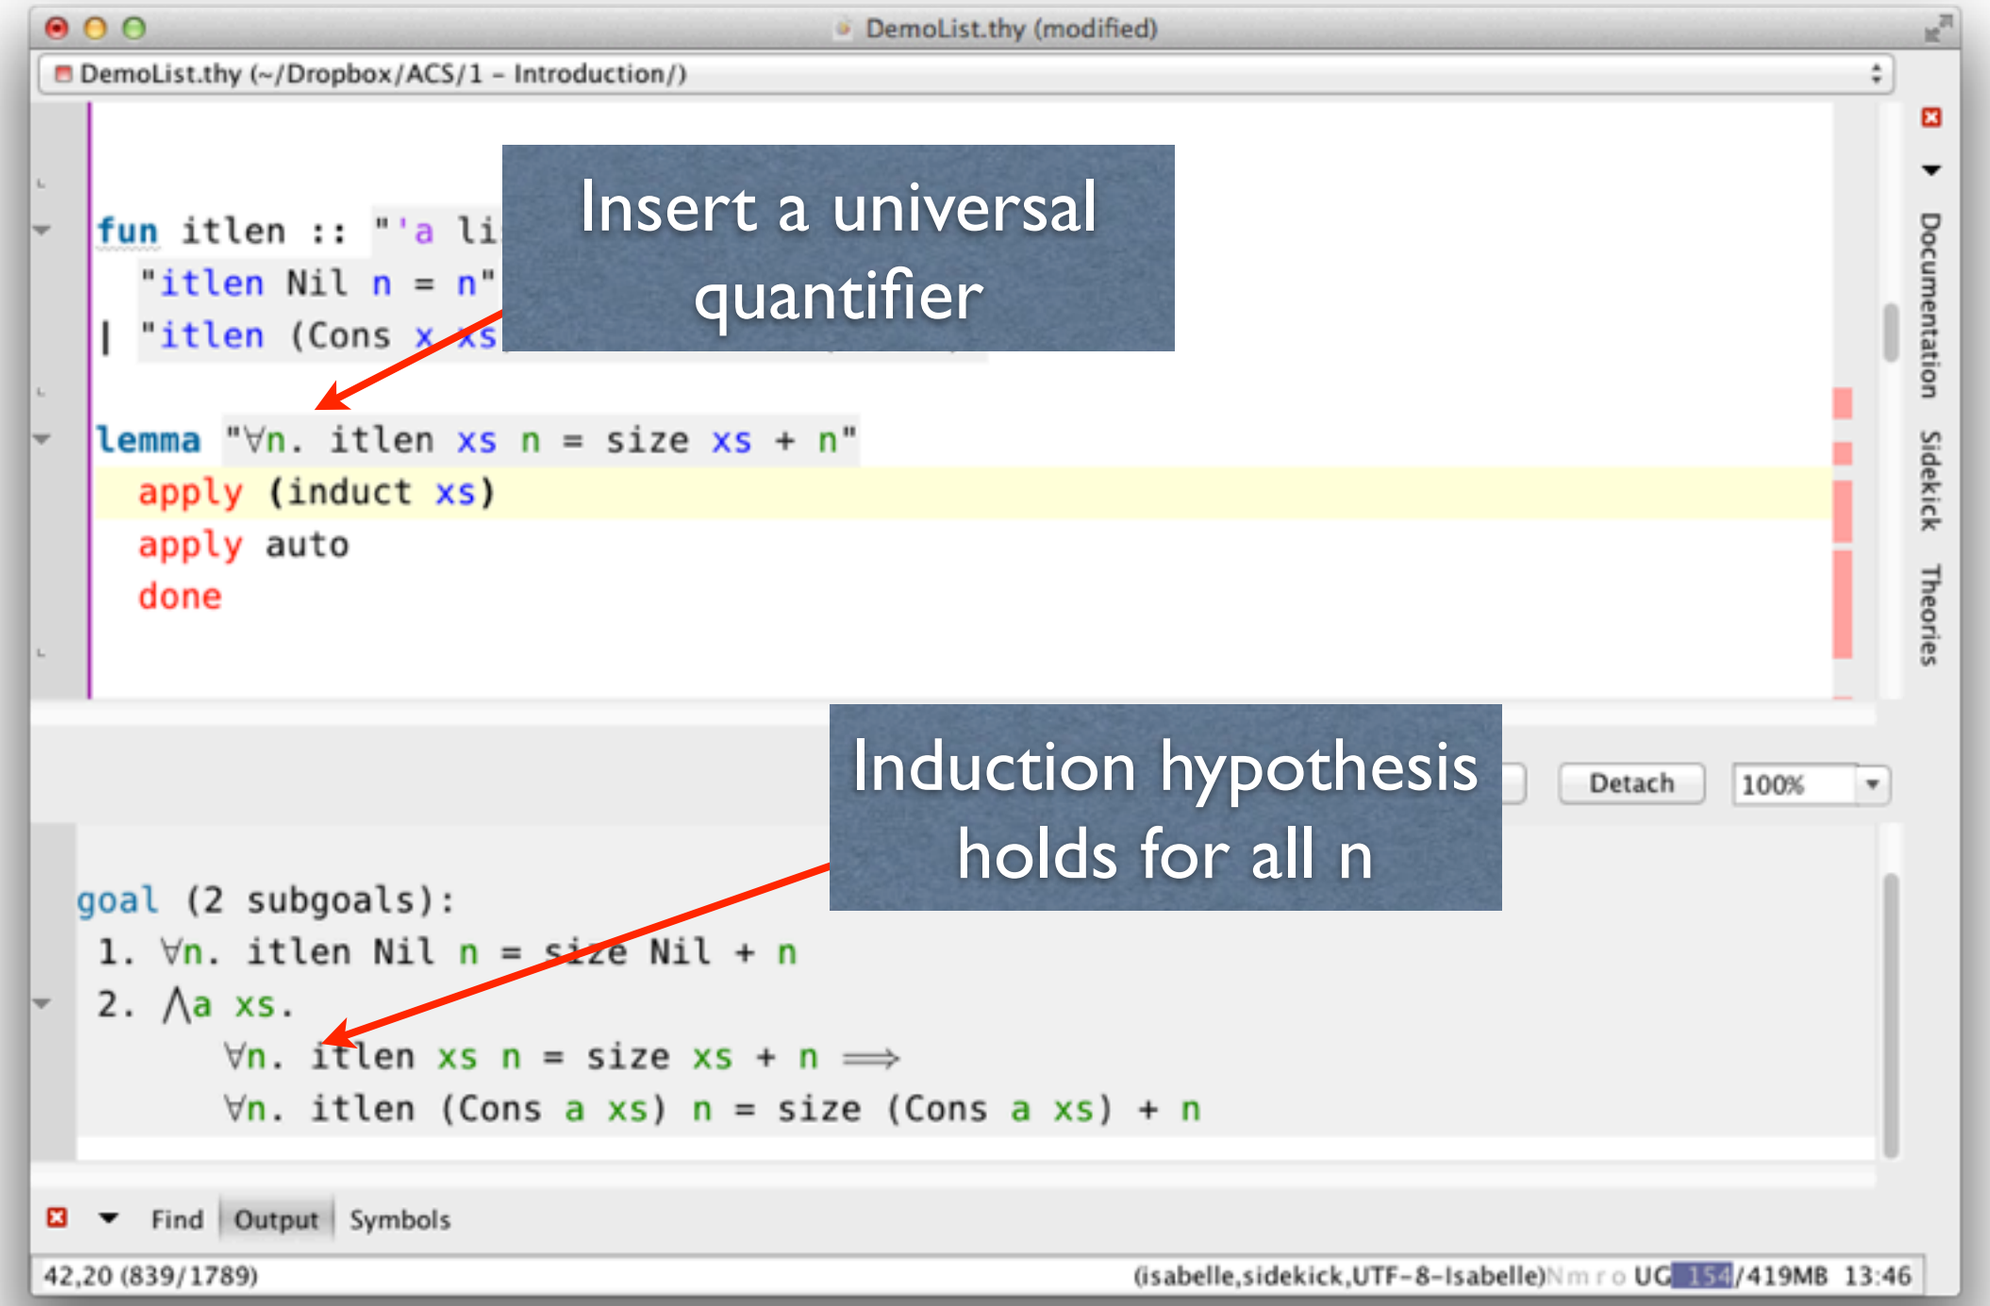

The need to generalise the induction formula in order to obtain a more general induction hypothesis Is well known from mathematics. Logically, note that the induction formula above has only one free variable: xs. The induction formula on the previous slide has two free variables: xs and n.

Generalising: A Better Way

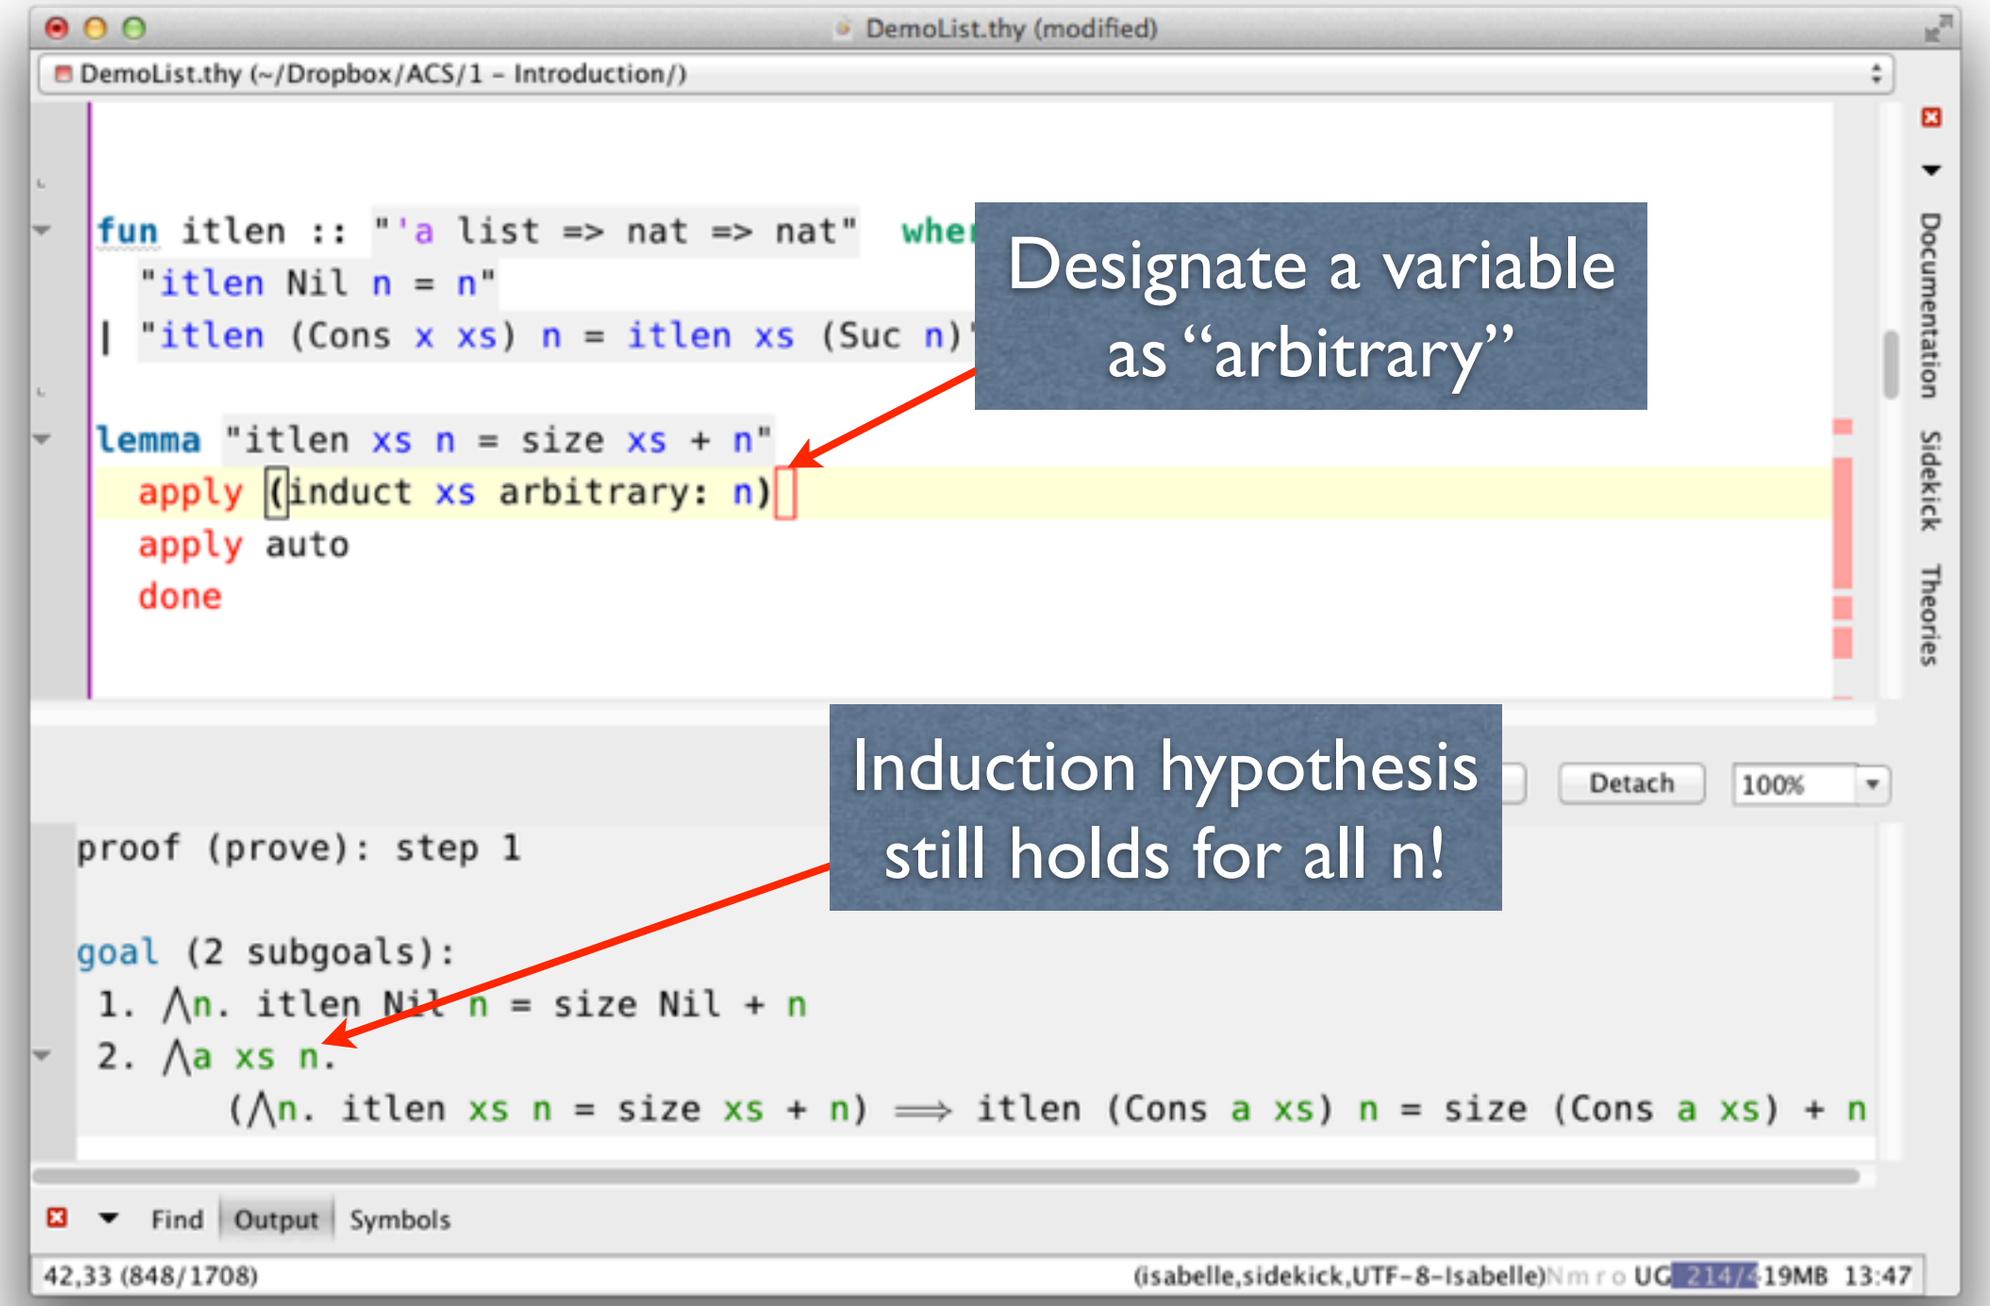

The approach described above is logically similar to the one on the previous slide, but it avoids the use of a universal quantifier (∀) in the theorem statement. Because Isabelle is a logical framework, it has meta-level versions of the universal quantifier and the implication symbol, and we generally avoid universal quantifiers in theorems. But it is important to remember that behind the convenience of the method illustrated here is a straightforward use of logic: we are still generalising induction formula. For more complicated examples, see the *Tutorial*, 9.2.1 **Massaging the Proposition**.

Unusual Recursions

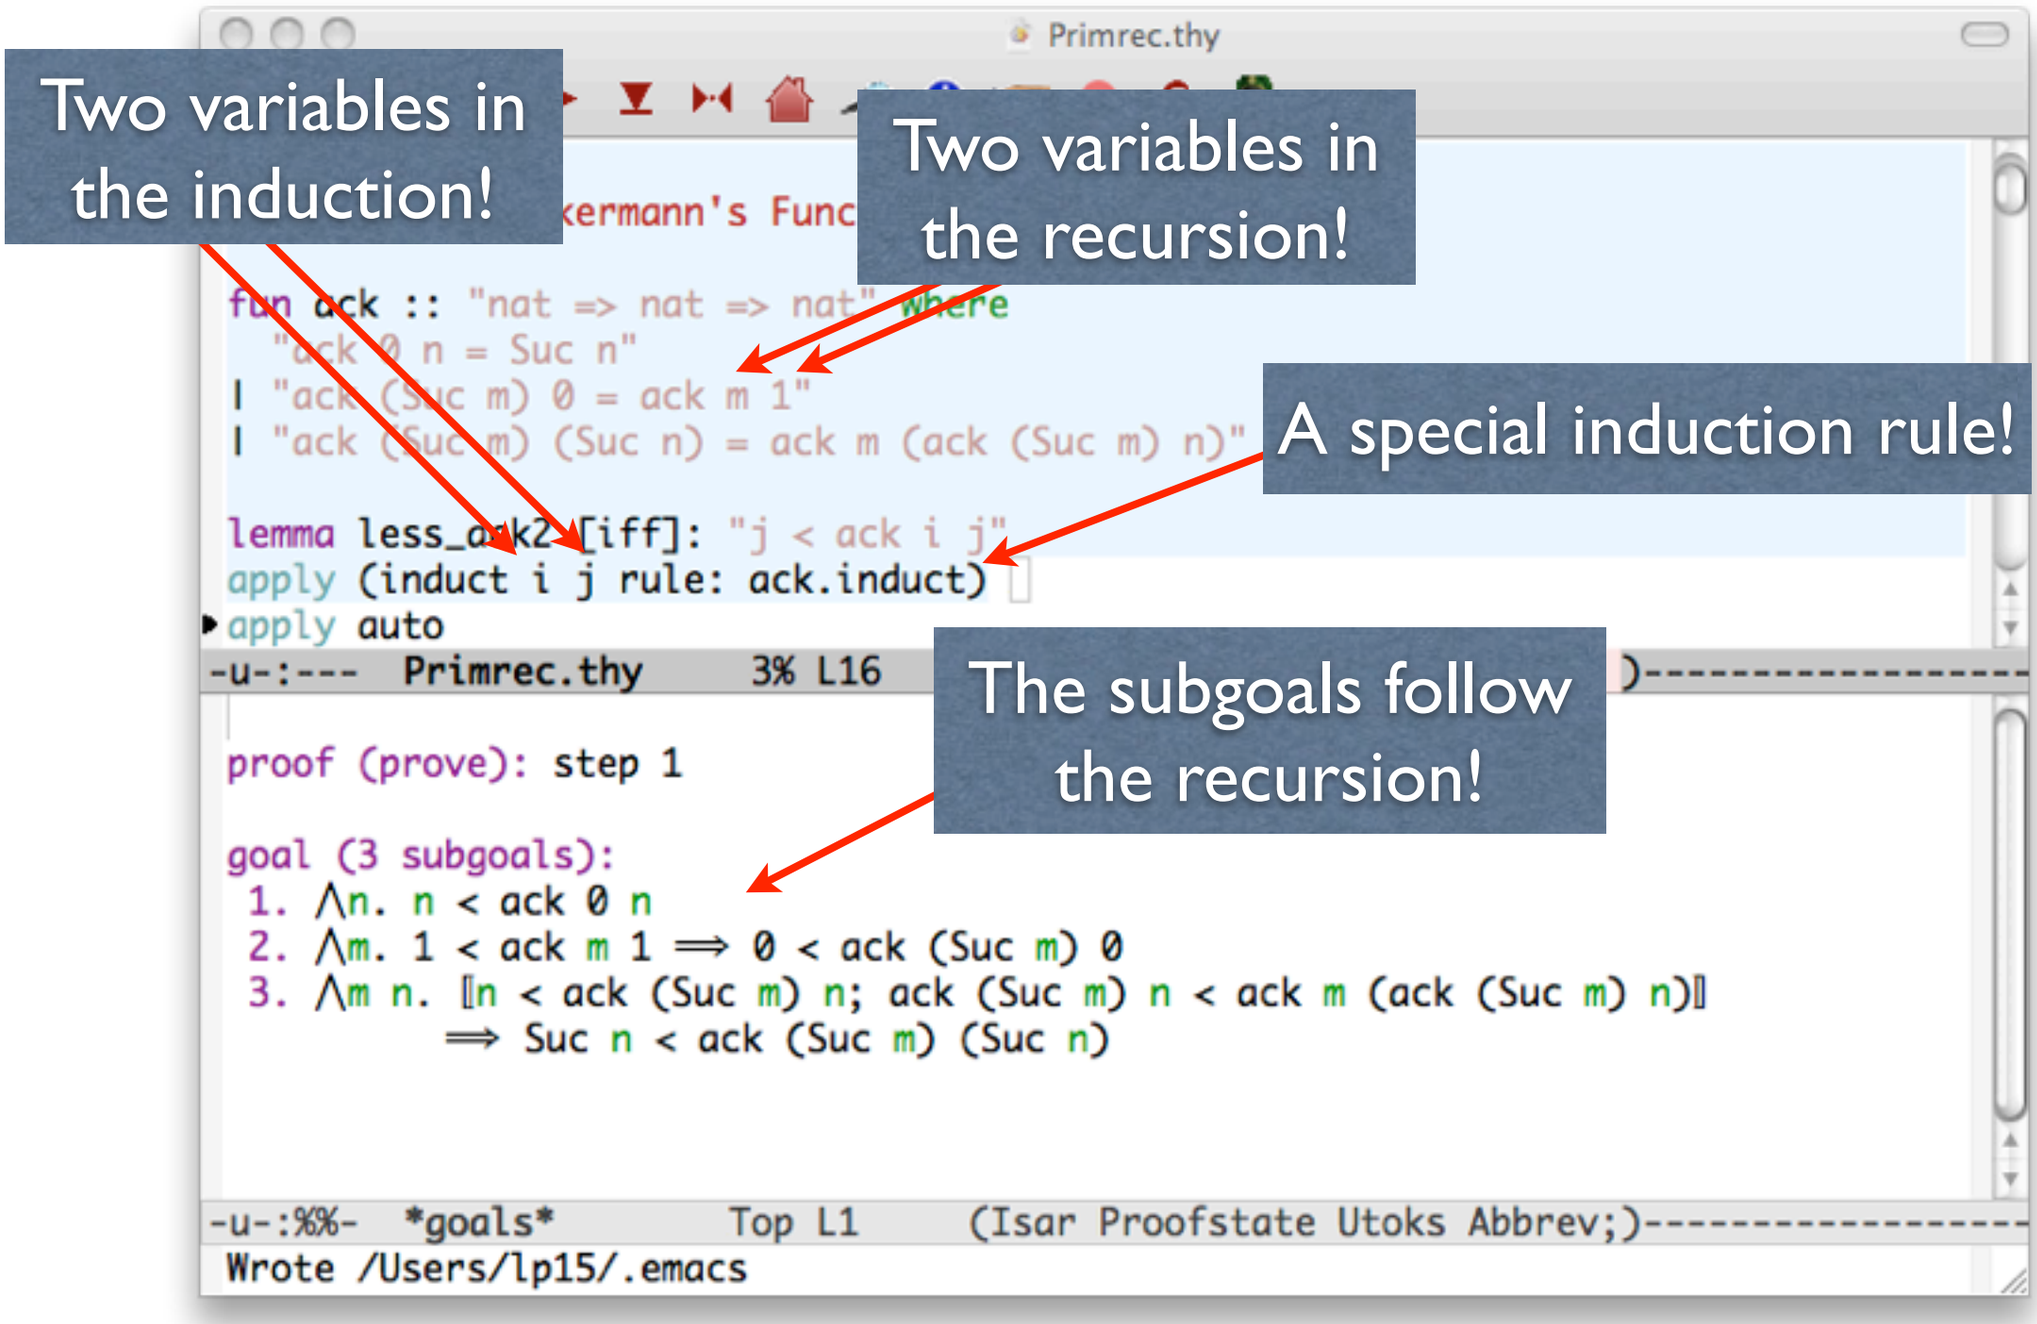

For full documentation, see *Defining Recursive Functions in Isabelle/HOL*, by Alexander Krauss.

Recursion: Key Points

- Recursion in one variable, following the structure of a datatype declaration, is called *primitive*.
- Recursion in multiple variables, terminating by size considerations, can be handled using `fun`.
  - `fun` produces a special induction rule.
  - `fun` can handle **nested recursion**.
  - `fun` also handles *pattern matching*, which it **completes**.

Isabelle provides the command `primrec` for primitive recursion as well. It is closely based on the internal derivation of recursion, and can handle function definitions involving certain complicated features (in particular, higher-order primitive recursion) where fun fails. See the *Tutorial*, **2.1 An Introductory Theory**. More difficult examples of `primrec` are covered in **3.3 Case Study: Compiling Expressions**.

Special Induction Rules

- They follow the function's recursion exactly.
- For Ackermann, they reduce $P\ x\ y$ to
  - $P\ 0\ n$, for arbitrary n
  - $P\ (Suc\ m)\ 0$ assuming $P\ m\ 1$, for arbitrary m
    - $P\ (Suc\ m)\ (Suc\ n)$ assuming $P\ (Suc\ m)\ n$ and $P\ m\ (ack\ (Suc\ m)\ n)$, for arbitrary m and n
- **Usually** they do what you want. Trial and error is tempting, but ultimately you will need to think!

The Ackermann example proves several lemmas using the special rule, but several others using ordinary mathematical induction!

Another Unusual Recursion

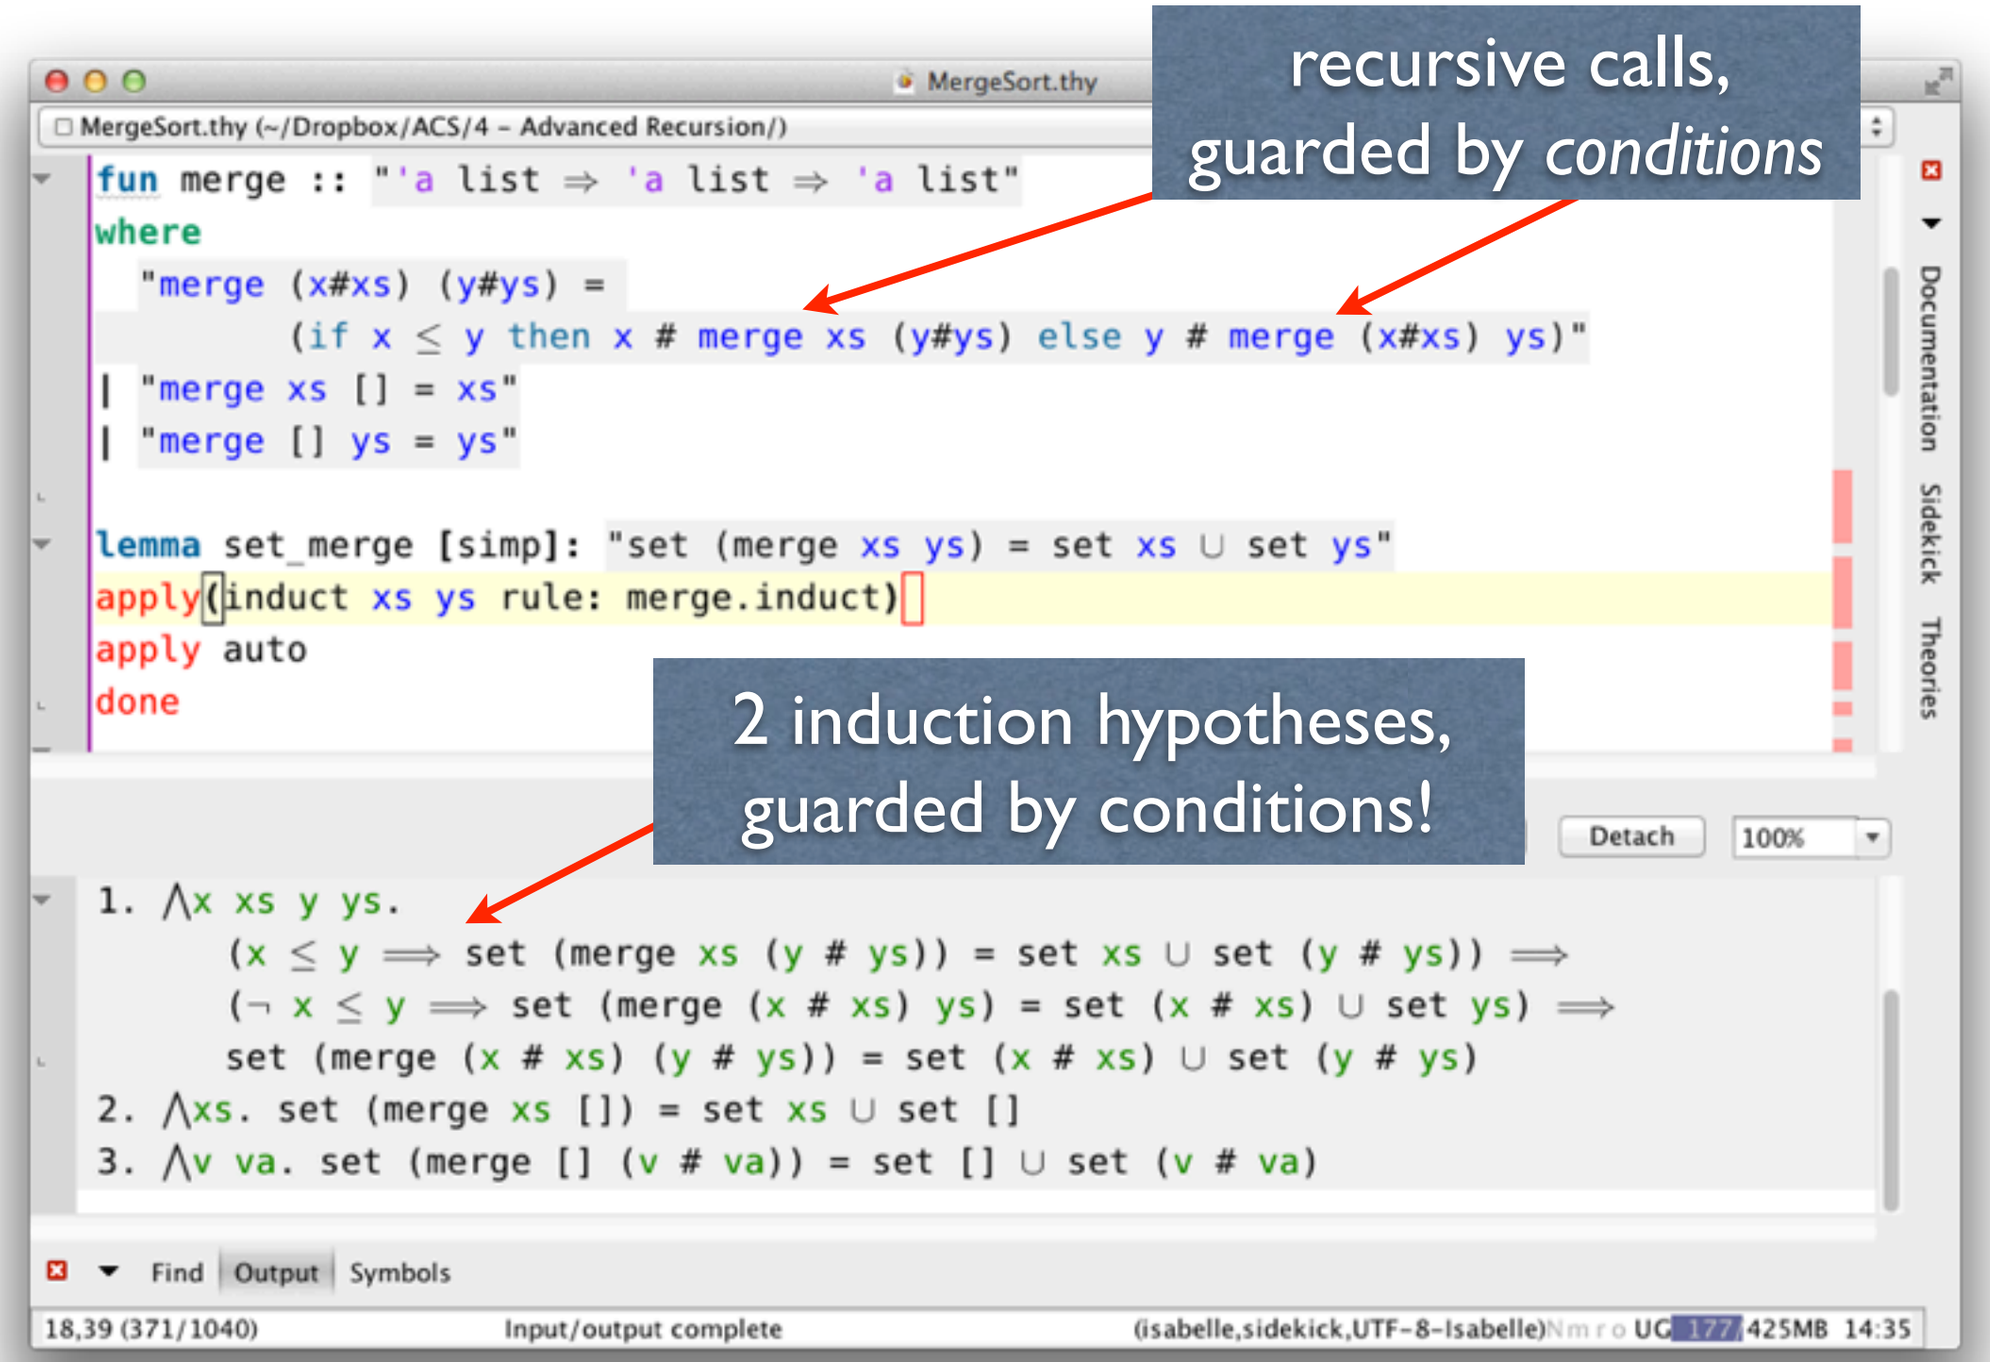

Again, see *Defining Recursive Functions in Isabelle/HOL*. Each induction hypothesis can only be used if the corresponding condition is provable.

Proof Outline

```
set (merge (x#xs) (y#ys)) = set (x # xs) ∪ set (y # ys)

      set (if x ≤ y then x # merge xs (y#ys)
            else y # merge (x#xs) ys)          =   ...
                              =
  (x ≤ y ⟶ set(x # merge xs (y#ys)) = ...) &
  (¬ x ≤ y ⟶ set(y # merge (x#xs) ys) = ...)
                              =
 (x ≤ y ⟶ {x} ∪ set(merge xs (y#ys)) = ...) &
 (¬ x ≤ y ⟶ {y} ∪ set(merge (x#xs) ys) = ...)
                              =
 (x ≤ y ⟶ {x} ∪ set xs ∪ set (y # ys) = ...) &
 (¬ x ≤ y ⟶ {y} ∪ set (x # xs) ∪ set ys = ...)
```

The first rewriting step in the proof unfolds the definition of `merge`. The second one is a case-split involving `if`. This step introduces a conjunction of implications, creating contexts that exactly match the induction hypotheses. But first, the definition of `set` (a function that maps a list to the finite set of its elements) must be unfolded. The last step highlighted above applies the induction hypotheses. The remaining steps, not shown, prove the equality between the set expressions just produced and the right-hand side of the original subgoal.

The Case Expression

- Similar to that found in the functional language ML.
- Automatically generated for every datatype.
- The simplifier can (upon request!) perform case-splits analogous to those for "`if`".
- Case splits in *assumptions* (not the conclusion) never happen unless requested.

Case-Splits for Lists

```
fun ordered :: "'a list => bool"
where
  "ordered [] = True"
| "ordered (x#l) =
    (case l of [] => True
             | Cons y xs => (x≤y & ordered (y#xs)))"
```

The definition shown on the slide describes the same function as the following one:

```
fun ordered :: "'a list => bool"
where
  "ordered [] = True"
| "ordered [x] = True"
| "ordered (x#y#xs) = (x \<le> y & ordered (y#xs))"
```

Case-Splitting in Action

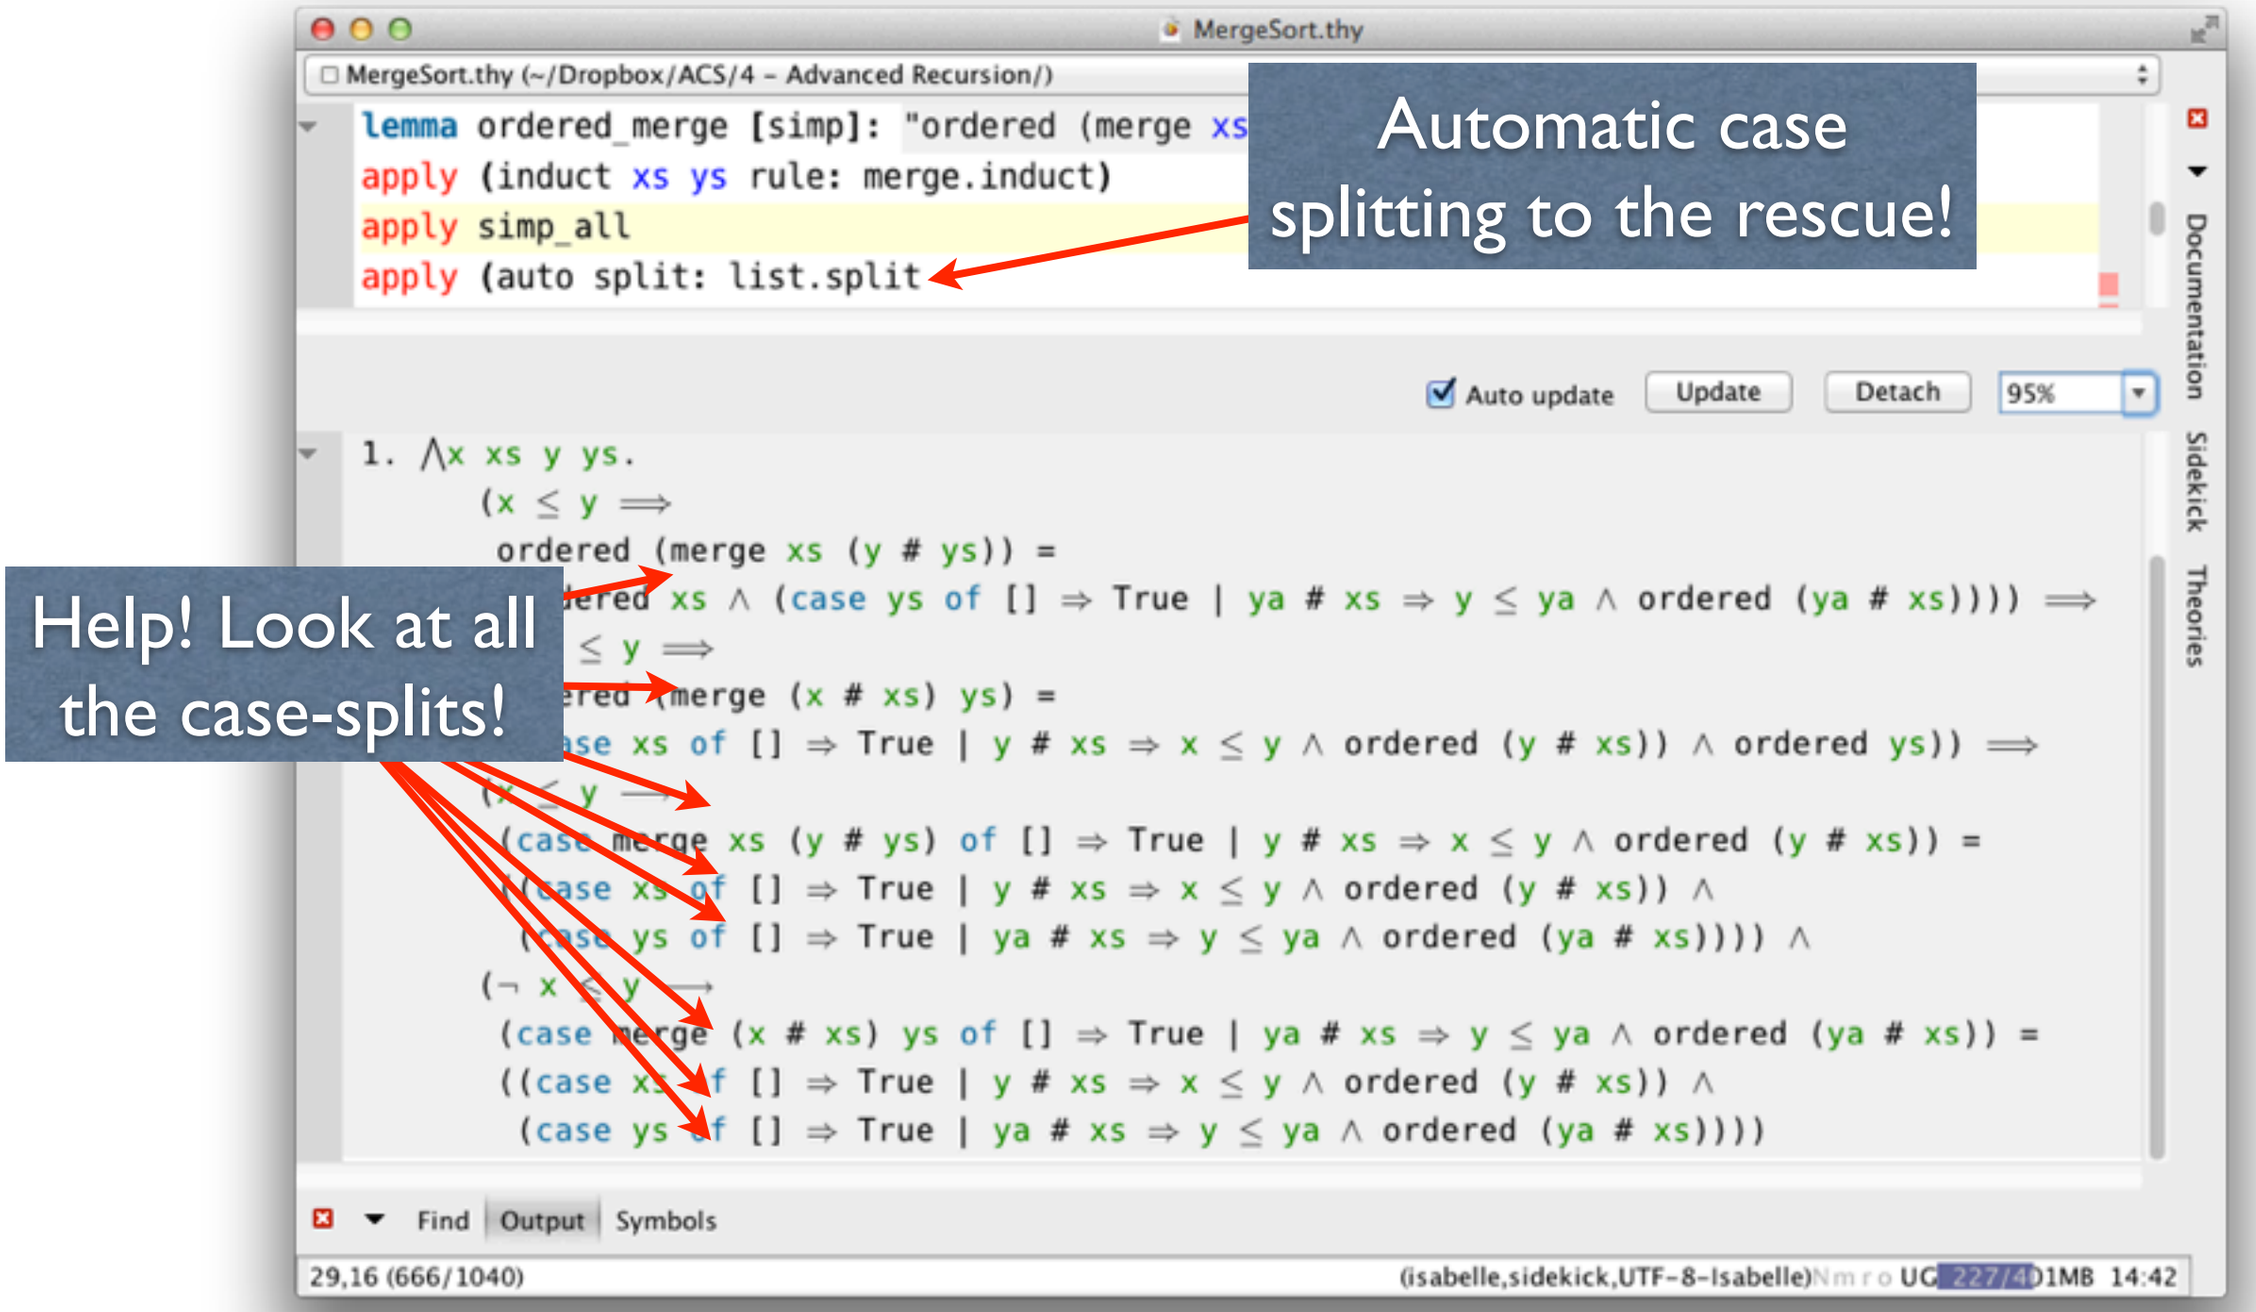

There isn't room to show the full subgoal, but the second part of the conjunction (beginning with ¬ x ≤ y) has a similar form to the first part, which is visible above.

Note that the last step used was `simp_all`, rather than `auto`. The latter would break up the subgoal according to its logical structure, leaving us with 14 separate subgoals! Simplification, on the other hand, seldom generates multiple subgoals. The one common situation where this can happen is indeed with case splitting, but in our example, case splitting completely proves the theorem.

Completing the Proof

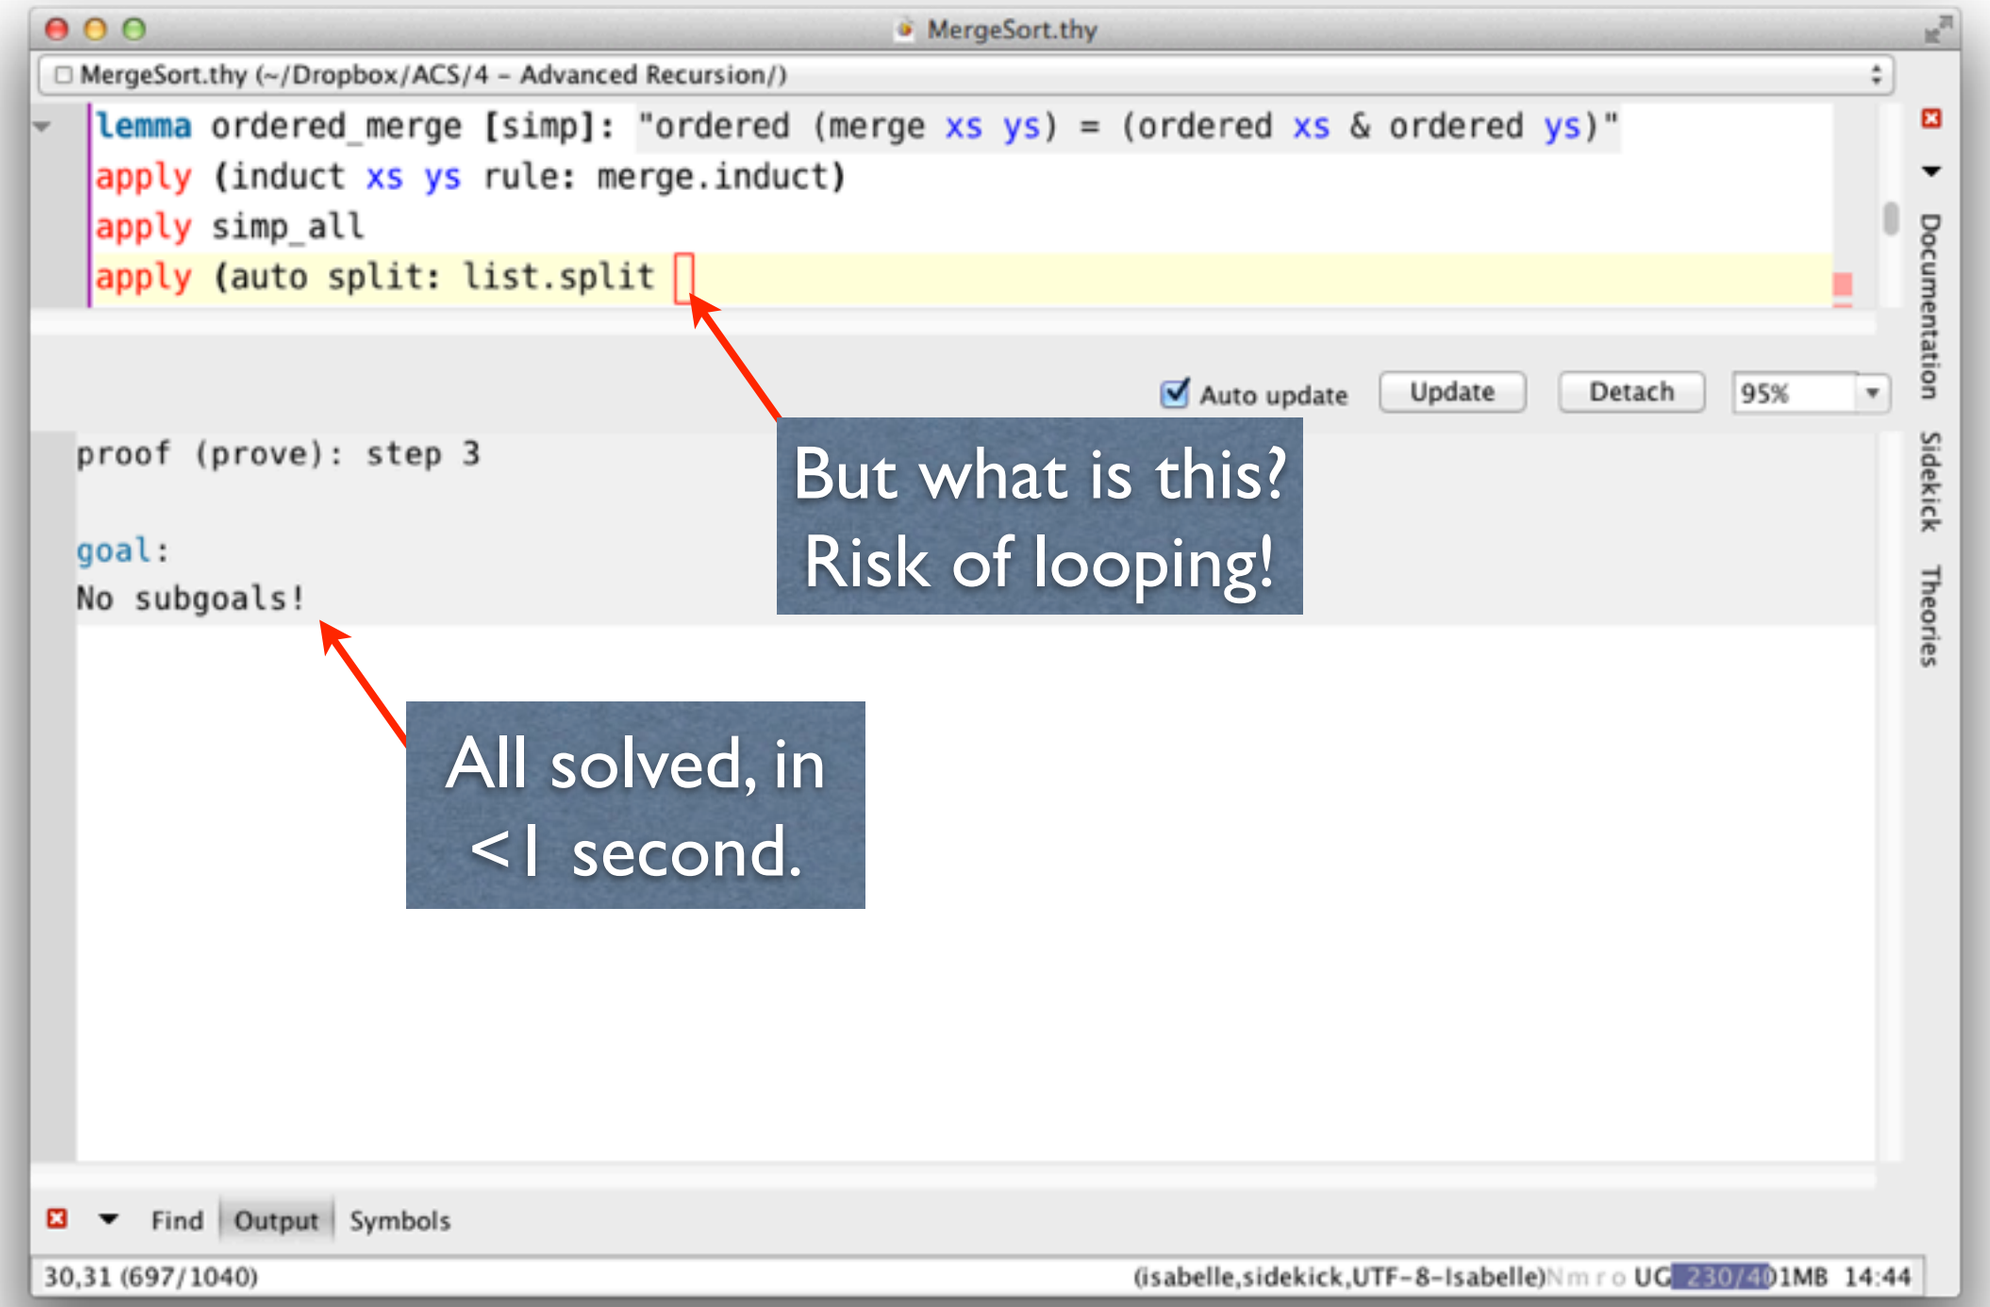

The identifier `ordered.simps` refers to the two equations that make up the definition of the function `ordered`. The suffix `(2)` selects the second of these. Now "`simp del: ordered.simps(2)`" tells `auto` to ignore this equation. Otherwise, the call will run forever.

Case Splitting for Lists

Simplification will replace

$$P\ (\texttt{case}\ xs\ \texttt{of}\ [\,]\ \texttt{=>}\ a\ \mid\ \texttt{Cons}\ x\ l\ \texttt{=>}\ b\ x\ l)$$

by

$$(xs = [\,] \longrightarrow P(a)) \wedge (\forall x\ l.\ xs = x \# l \longrightarrow P(b\ x\ l))$$

- It creates a case for each datatype constructor.
- Here it causes the simplifier to loop if combined with the second rewrite rule for ordered.

Specifically, a case split will create an instance where the list has the form a#l, and therefore ordered(a#l) will rewrite to another instance of case, *ad infinitum*.

Summary

- Many forms of recursion are available.
- The supplied induction rule often leads to simple proofs.
- The "case" operator can often be dealt with using automatic case splitting...
- but complex simplifications can run forever!

How to Trace the Simplifier

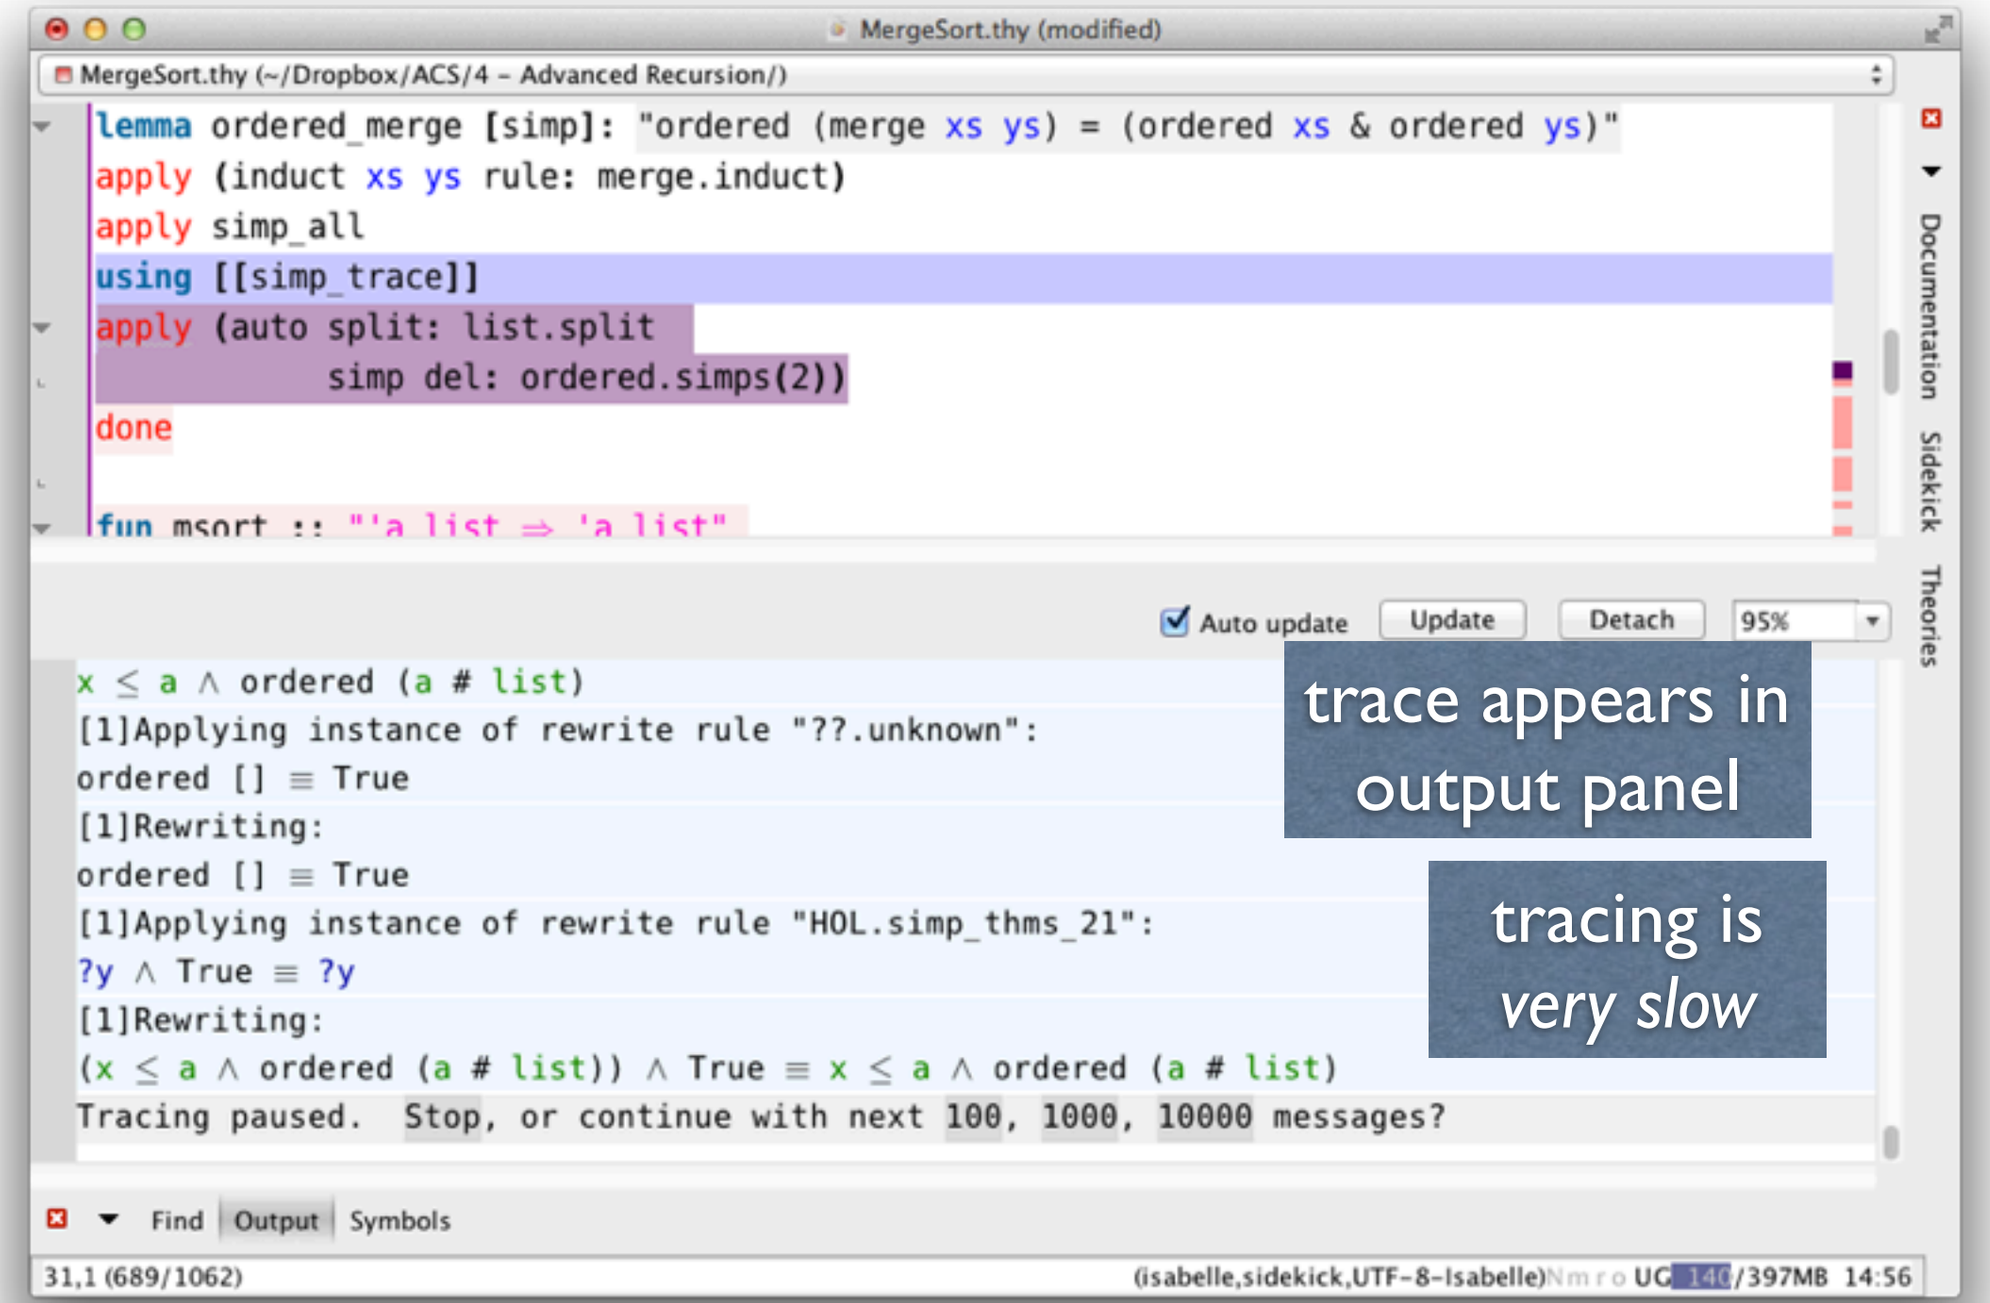

Interactive Formal Verification
5: Logic in Isabelle

Lawrence C Paulson
Computer Laboratory
University of Cambridge

Logical Frameworks

- A formalism to represent other formalisms
- Support for *natural deduction*
- A common basis for implementations
- Type theories are commonly used, but Isabelle uses a simple meta-logic whose main primitives are
  - $\Rightarrow$ (implication)
  - $\bigwedge$ (universal quantification)

Isabelle's Family of Logics

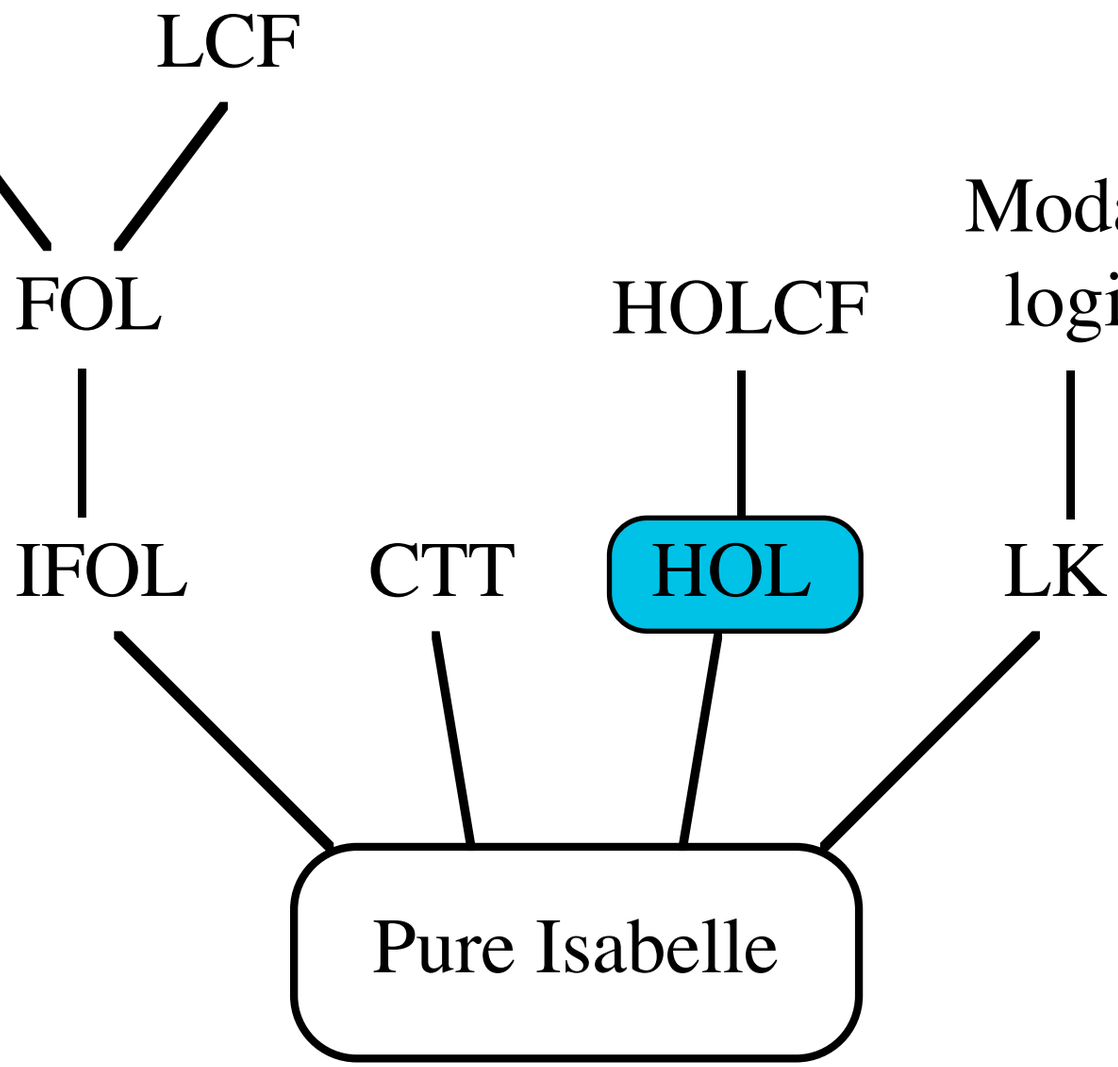

Natural Deduction Basics

- Proof is done using mainly *inference rules* rather than axioms.
- For each logical symbol, there are rules to *introduce* and *eliminate* it.
- Assumptions can be *introduced* and *discharged*.
- Contrast with *Hilbert-style* proof systems, where typically the main inference rule is *modus ponens*…
- and there are many cryptic axioms, each combining a number of logical symbols.

Natural Deduction in Isabelle

$$\frac{P \quad Q}{P \wedge Q}$$

```
P ⇒ (Q ⇒ P ∧ Q))
```

$$\frac{P \wedge Q}{P}$$

```
P ∧ Q ⇒ P
```

$$\frac{P \wedge Q}{Q}$$

```
P ∧ Q ⇒ Q
```

$$\frac{P \to Q \quad P}{Q}$$

```
P⟶Q ⇒ (P ⇒ Q)
```

See the *Tutorial*, Chapter 5: **The Rules of the Game**. The first of these is an *introduction* rule, `conjI` in Isabelle. The following three are *elimination* rules: `conjunct1`, `conjunct2`, and `mp`. Isabelle parlance, these three are actually *destruction* rules because they lack the general form of an elimination rule in natural deduction.

Meta-implication

- The symbol ⇒ (or ==>) expresses the relationship between premise and conclusion
- ... and between subgoal and goal.
- It is distinct from →, which is not part of Isabelle's underlying logical framework.
- `P⇒(Q⇒R)` is abbreviated as `⟦P;Q⟧ ⇒ R`

The distinction between meta- and object-connectives is a common source of confusion among students. This distinction is inherent in the use of a logical framework. There is no reason why an object-logic would have an implication symbol at all. Isabelle gives a special significance to ⇒, in particular for expressing the structure of inference rules, as shown on previous slide. This would be impossible if we make no distinction between ⇒ and →.

A Trivial Proof

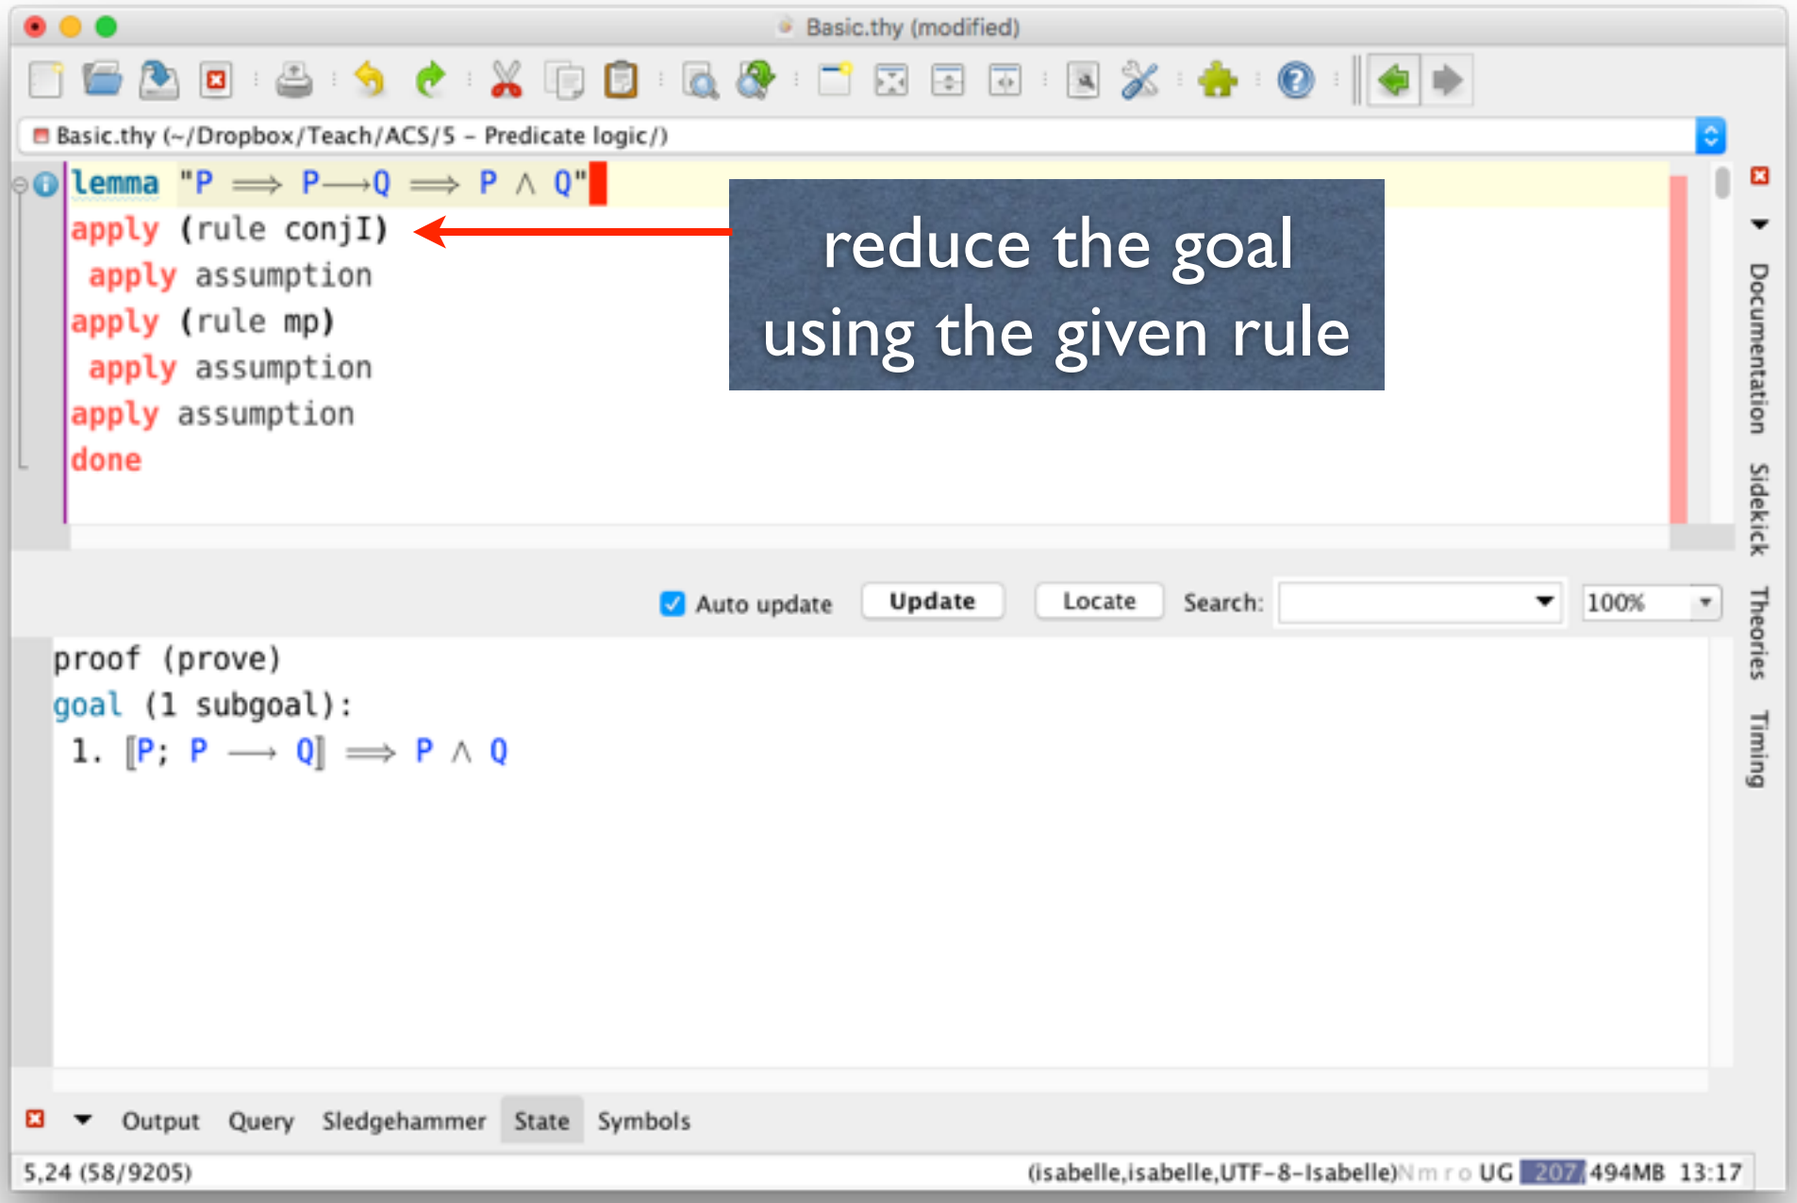

The method “`rule`” is one of the most primitive in Isabelle. It matches the conclusion of the supplied rule with that of the a subgoal, which is replaced by new subgoals: the corresponding instances of the rule’s premises. See the *Tutorial*, **5.7 Interlude: the Basic Methods for Rules**.

Normally, it applies to the *first* subgoal, though a specific goal number can be specified; many other proof methods follow the same convention.

Proof by Assumption

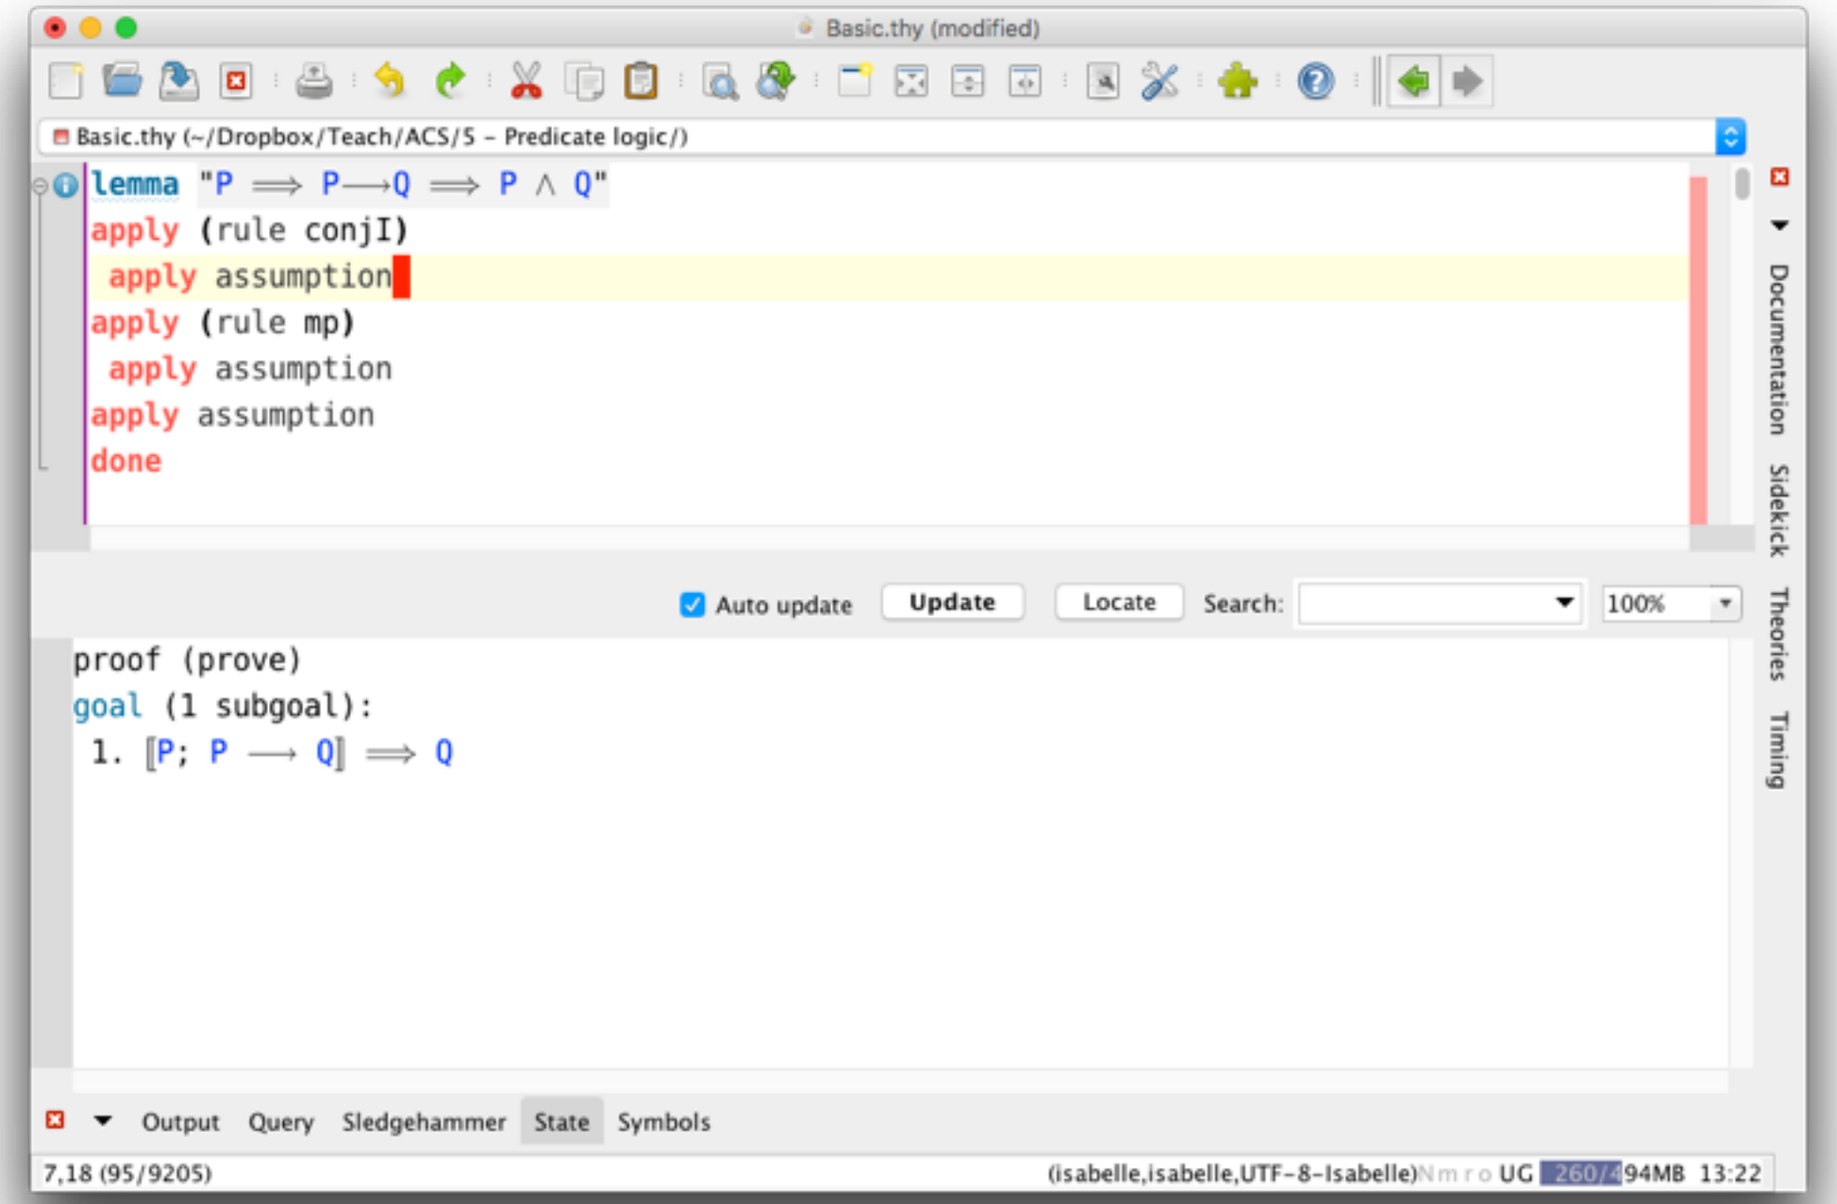

The method "`assumption`" proves (and deletes!) a subgoal if it can unify the subgoal's conclusion with one of its premises.

Unknowns in Subgoals

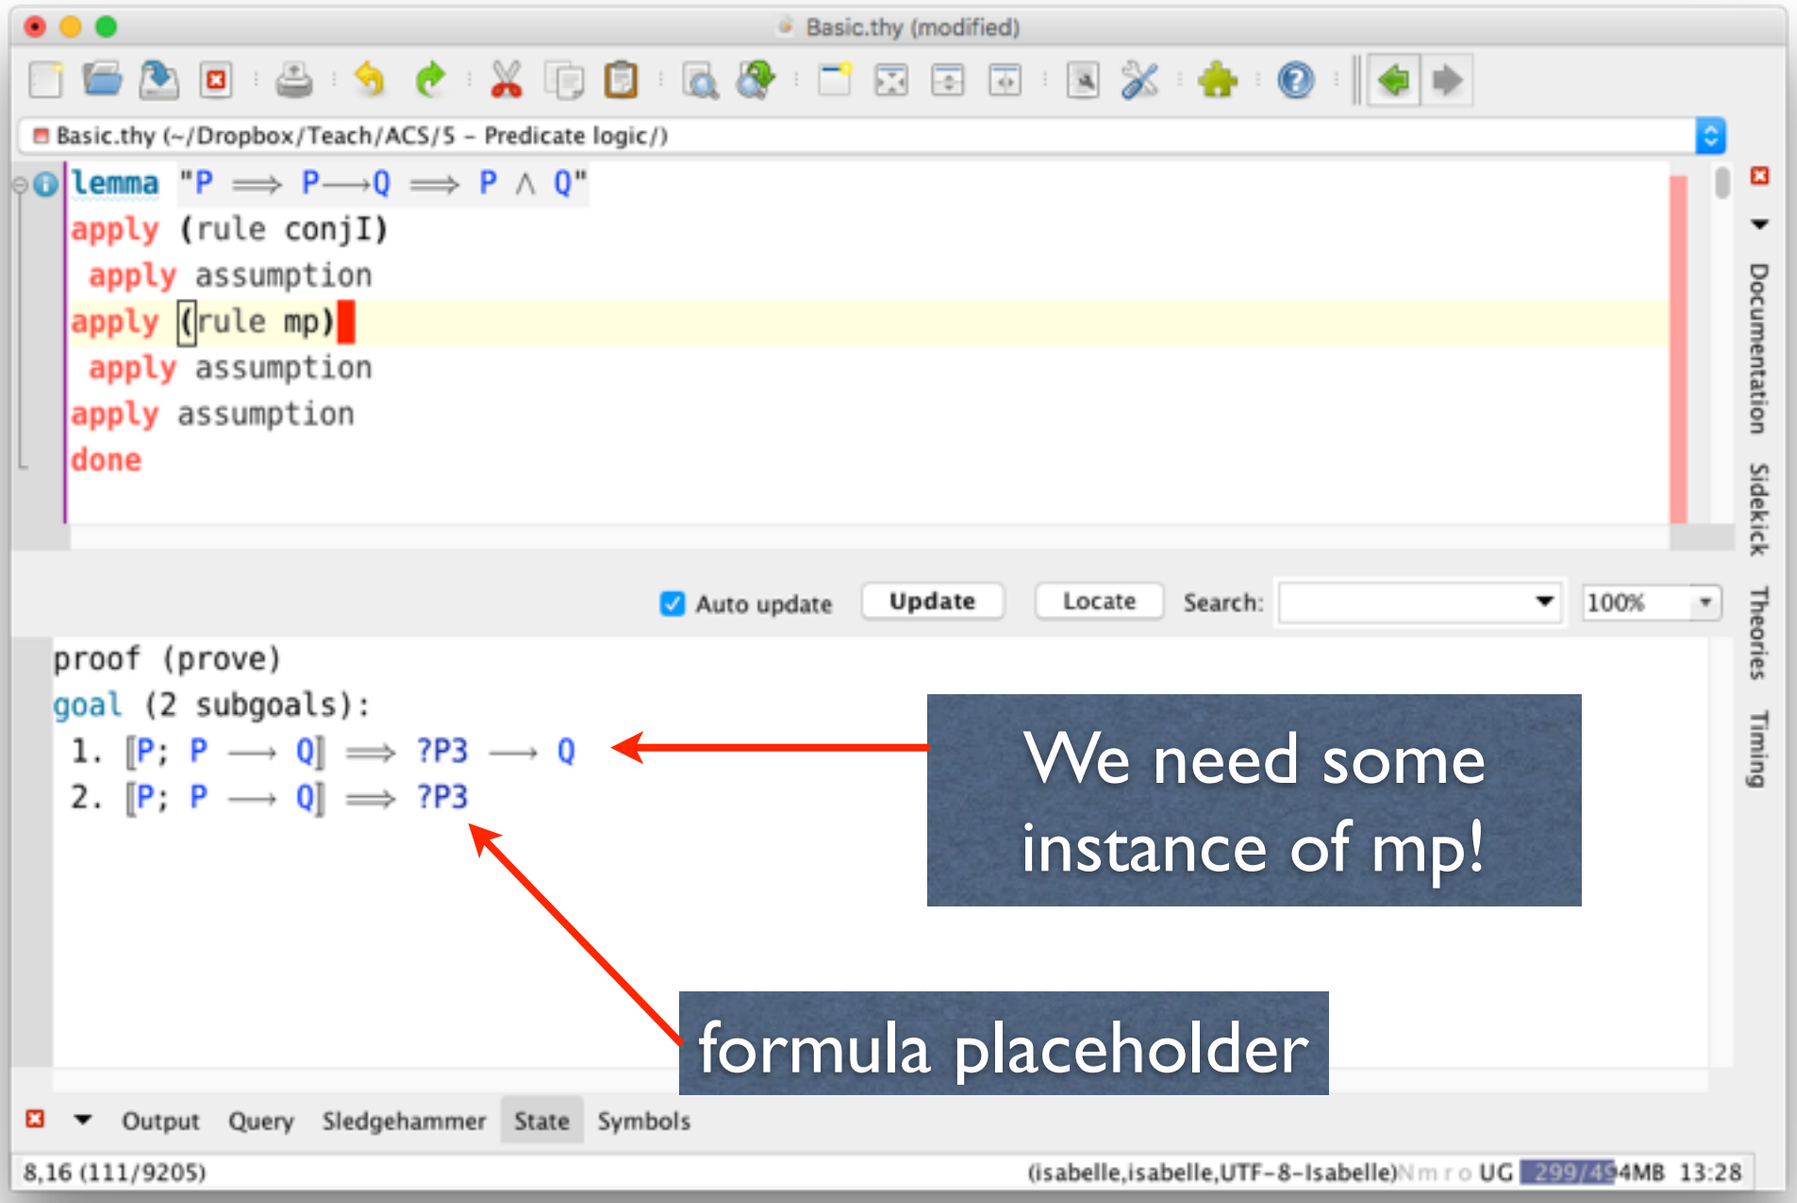

Isabelle includes a class of variables whose names begin with the ? character. They are called unknowns or schematic variables. Logically, they are no different from ordinary free variables, but Isabelle treats them differently: it allows them to be replaced by other expressions during unification. Isabelle rewrite rules and inference rules contain many such variables, but we normally suppress the question marks to make them easier to read. For example, the rule `conjI` is really `?P ==> (?Q ==> ?P & ?Q)`.

Unknowns and Unification

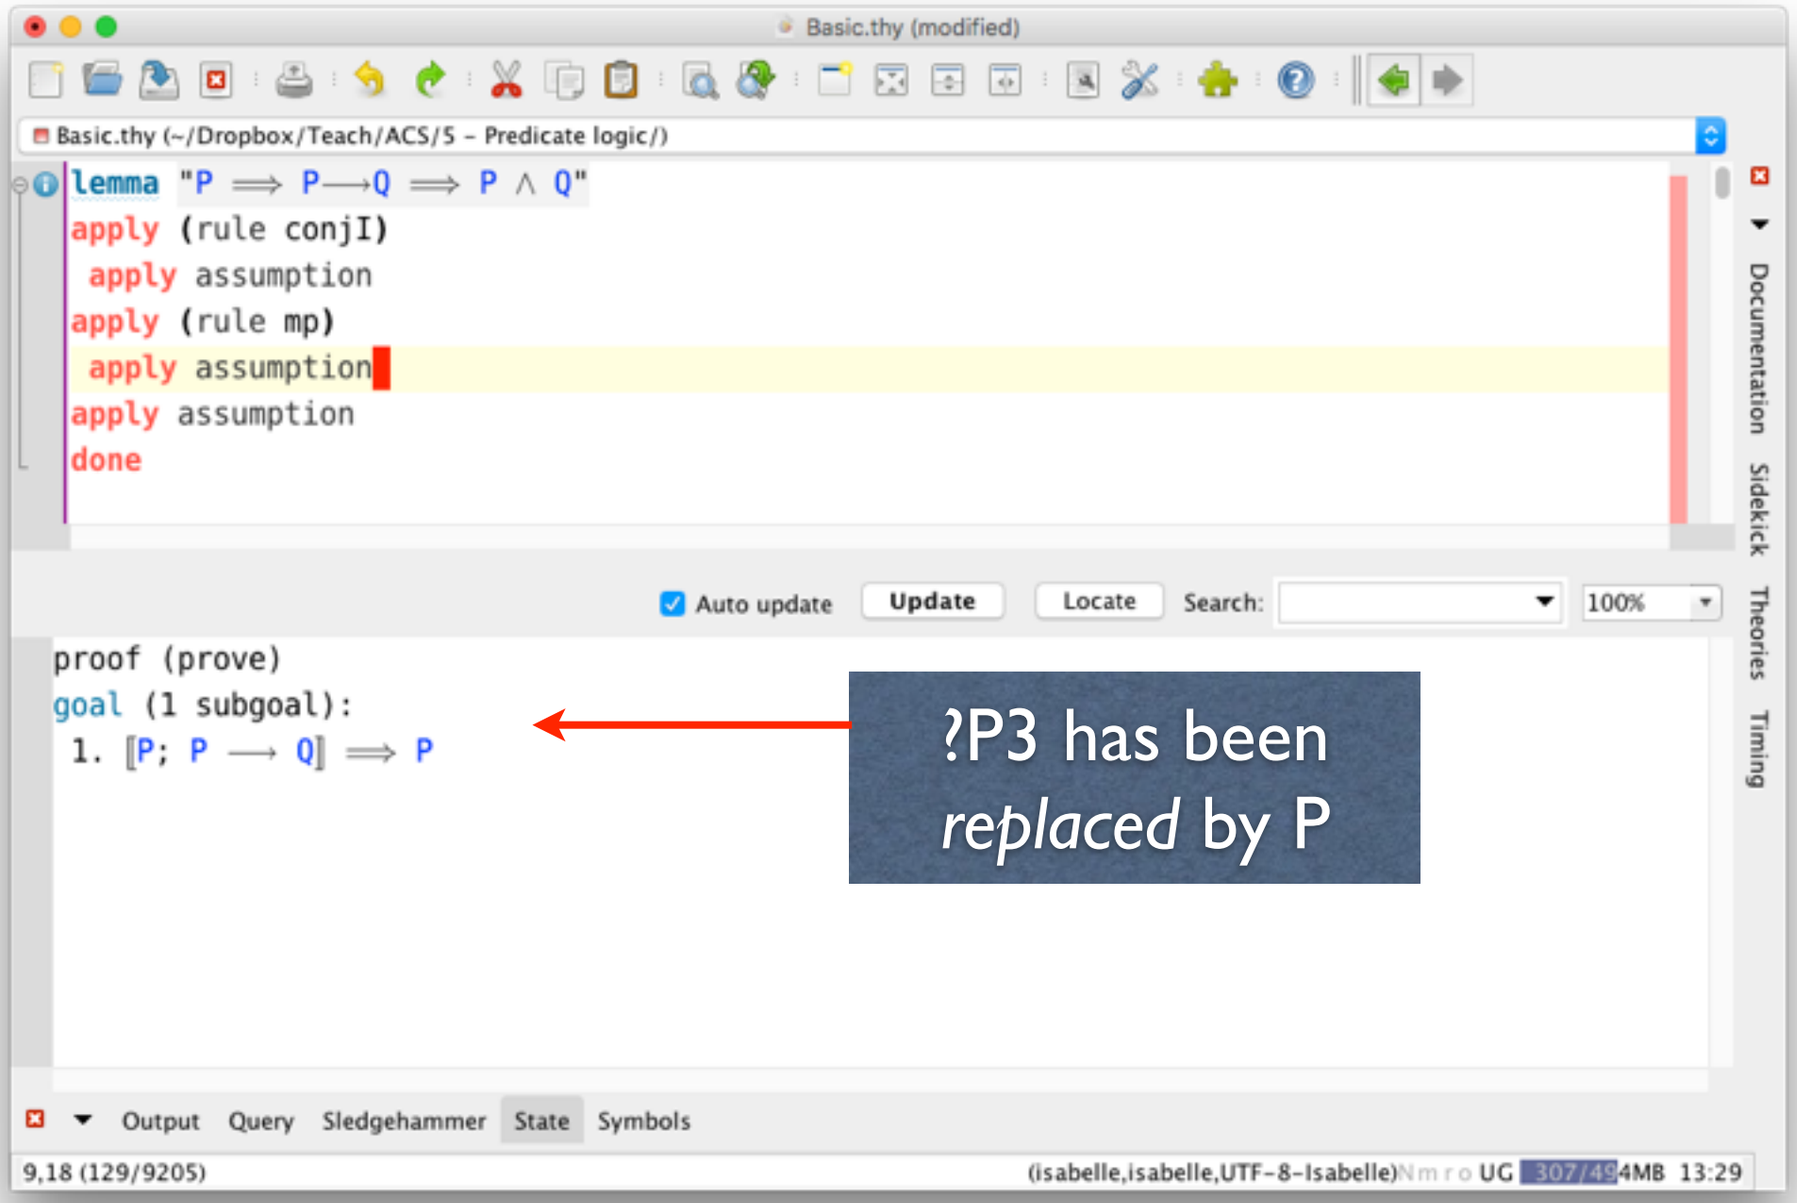

Proving ?P3→Q from the assumption P→Q performs unification, and the variable ?P3 is updated. All occurrences of the variable are updated. In this way, proving one subgoal can make another subgoal impossible to prove. Sometimes there are multiple choices and only one will allow the proof to go through.

Discharging Assumptions

$$\frac{\begin{array}{c}[P]\\ \vdots\\ Q\end{array}}{P \to Q}$$

```
(P ⇒ Q) ⇒ P⟶Q
```

$$\frac{P \vee Q \qquad \begin{array}{c}[P]\\ \vdots\\ R\end{array} \qquad \begin{array}{c}[Q]\\ \vdots\\ R\end{array}}{R}$$

```
⟦P ∨ Q; P⇒R; Q⇒R⟧ ⇒ R
```

Such rules take derivations that depend upon particular assumptions (written as [P] and [Q] above) and “discharge” those assumptions, which means that the conclusion is not regarded as depending on them. The backwards interpretation is more natural: to prove P→Q, it suffices to assume P and prove Q.

Meta-level implication (⇒) expresses the discharging of assumptions as well as the relationship between premises and conclusion.

A Proof using Assumptions

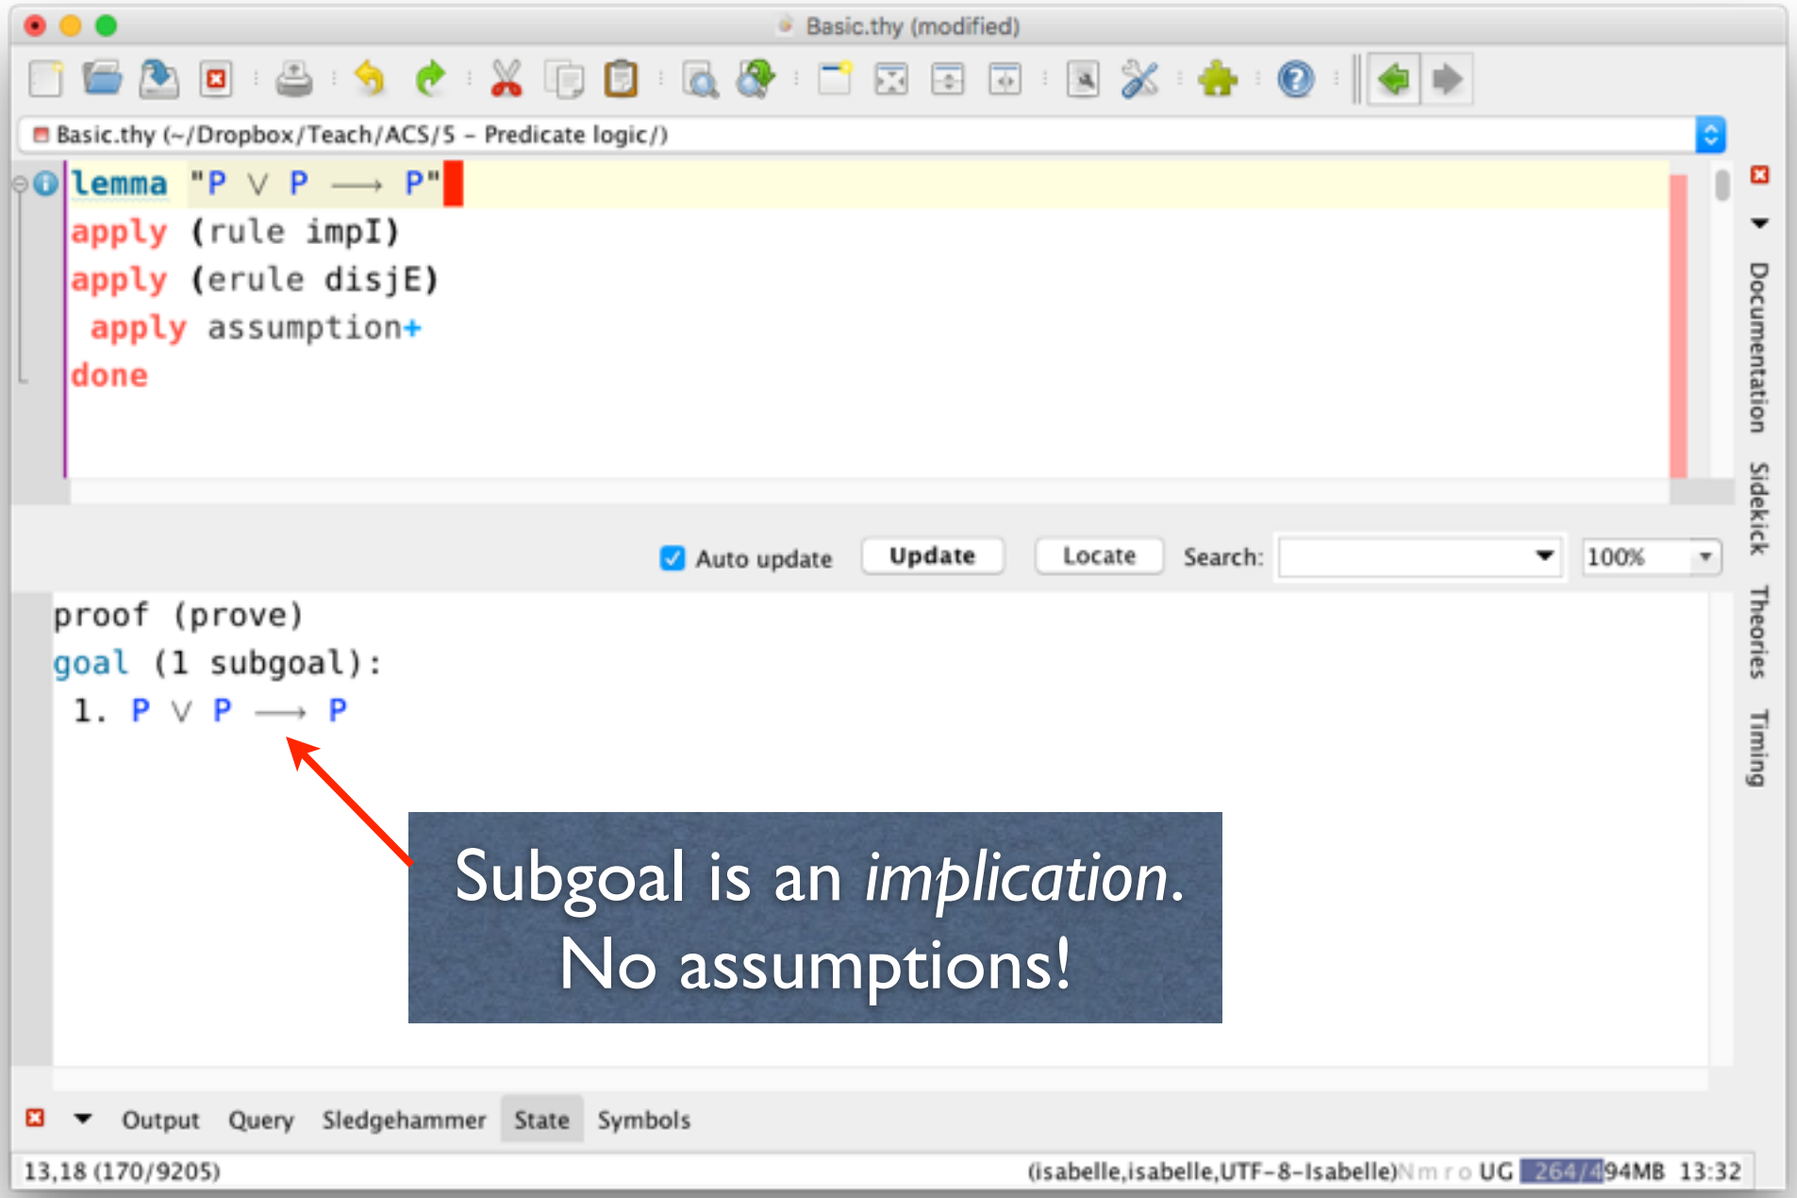

A full list of the predicate calculus rules for higher-order logic is available in **Isabelle's Logics: HOL**, an old but still useful reference manual.

After Implies-Introduction

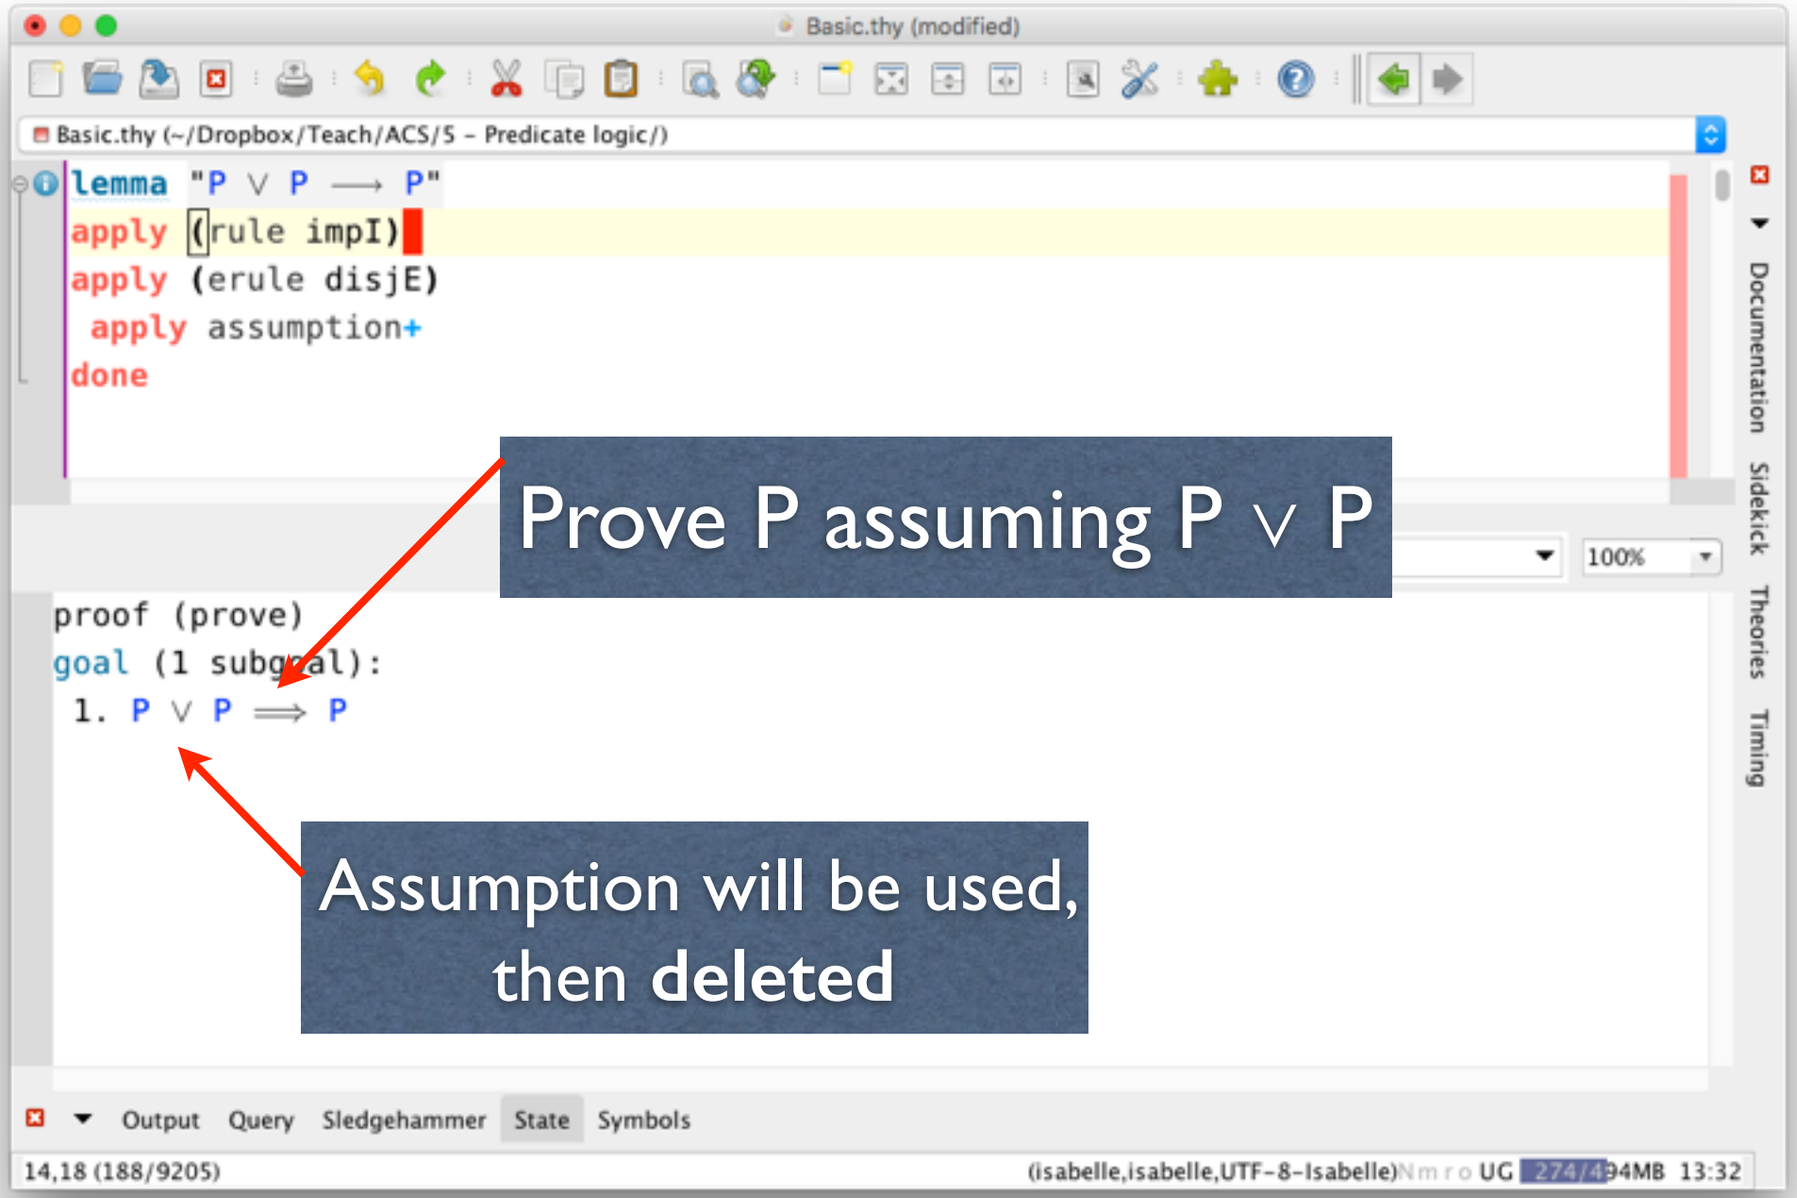

Disjunction Elimination

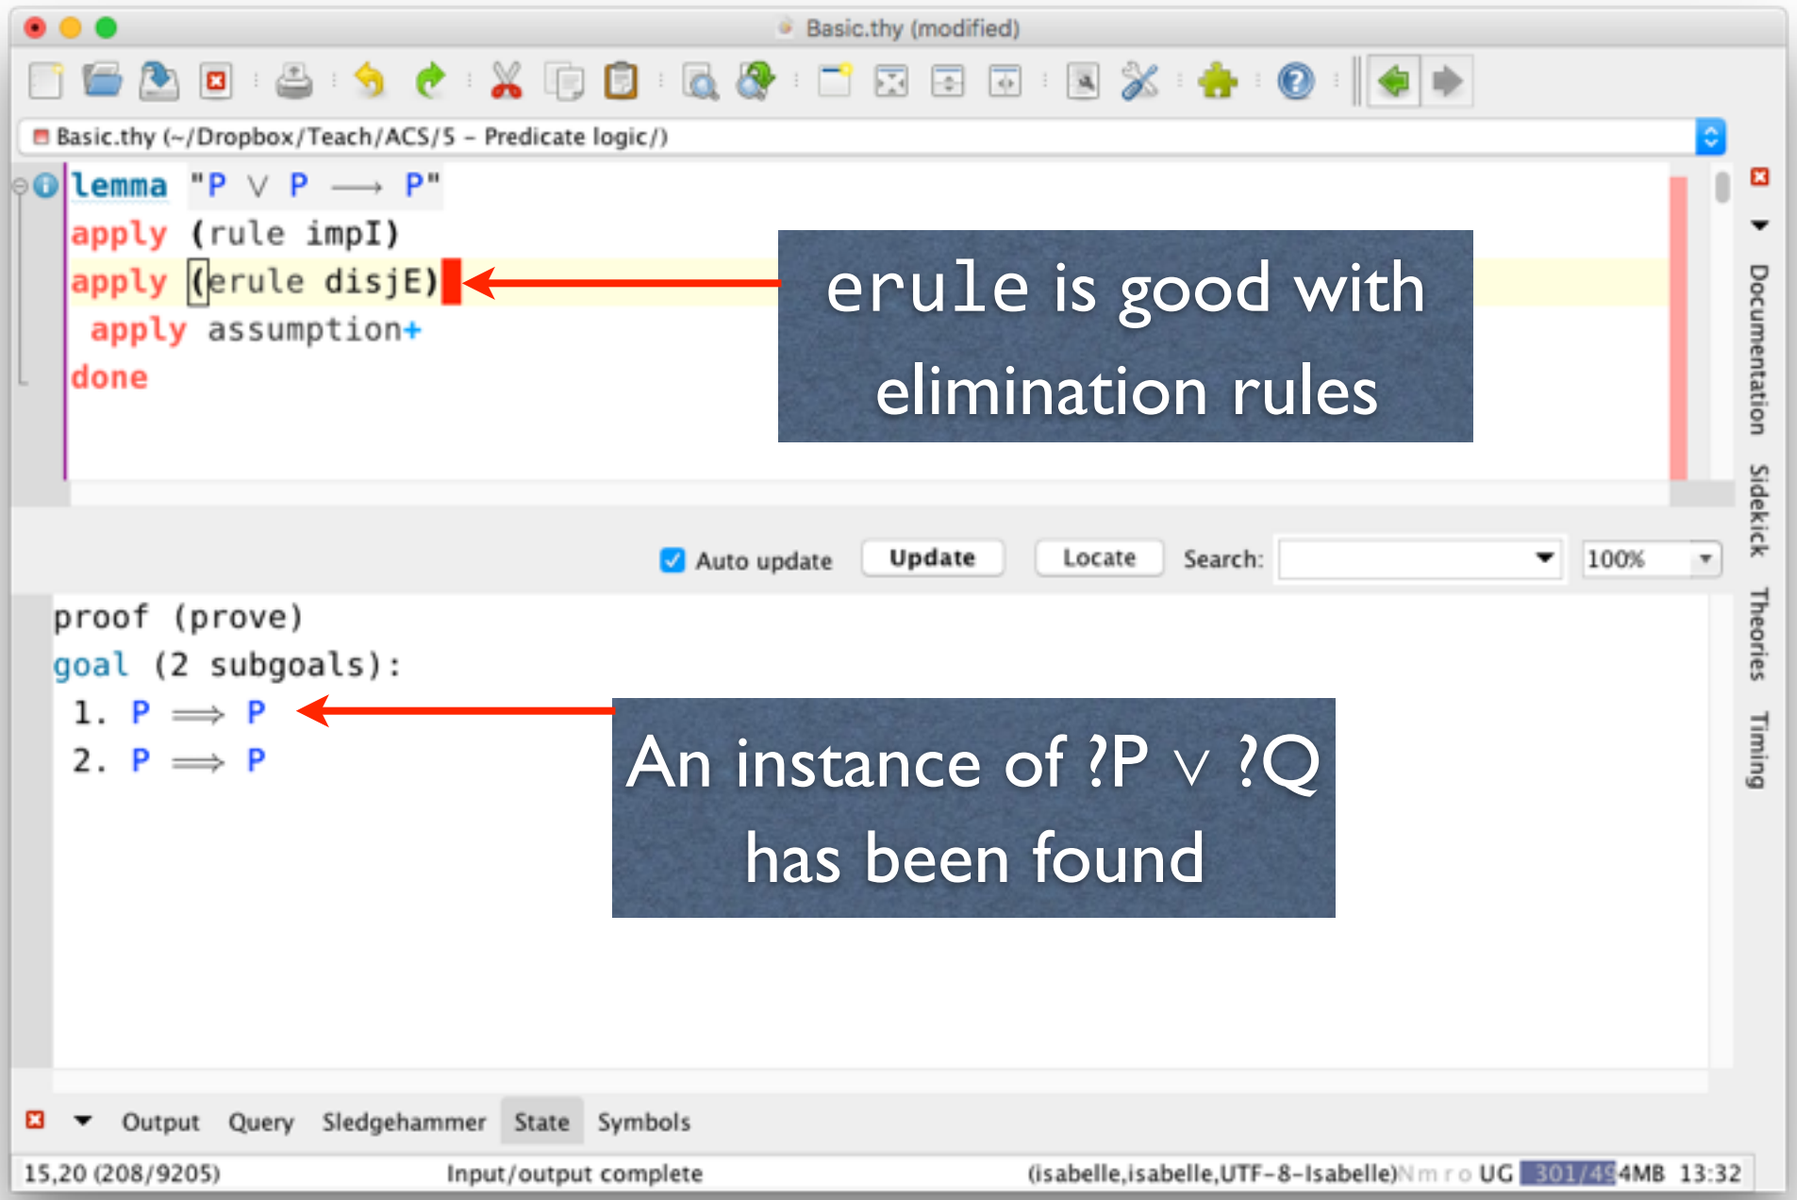

The Final Step

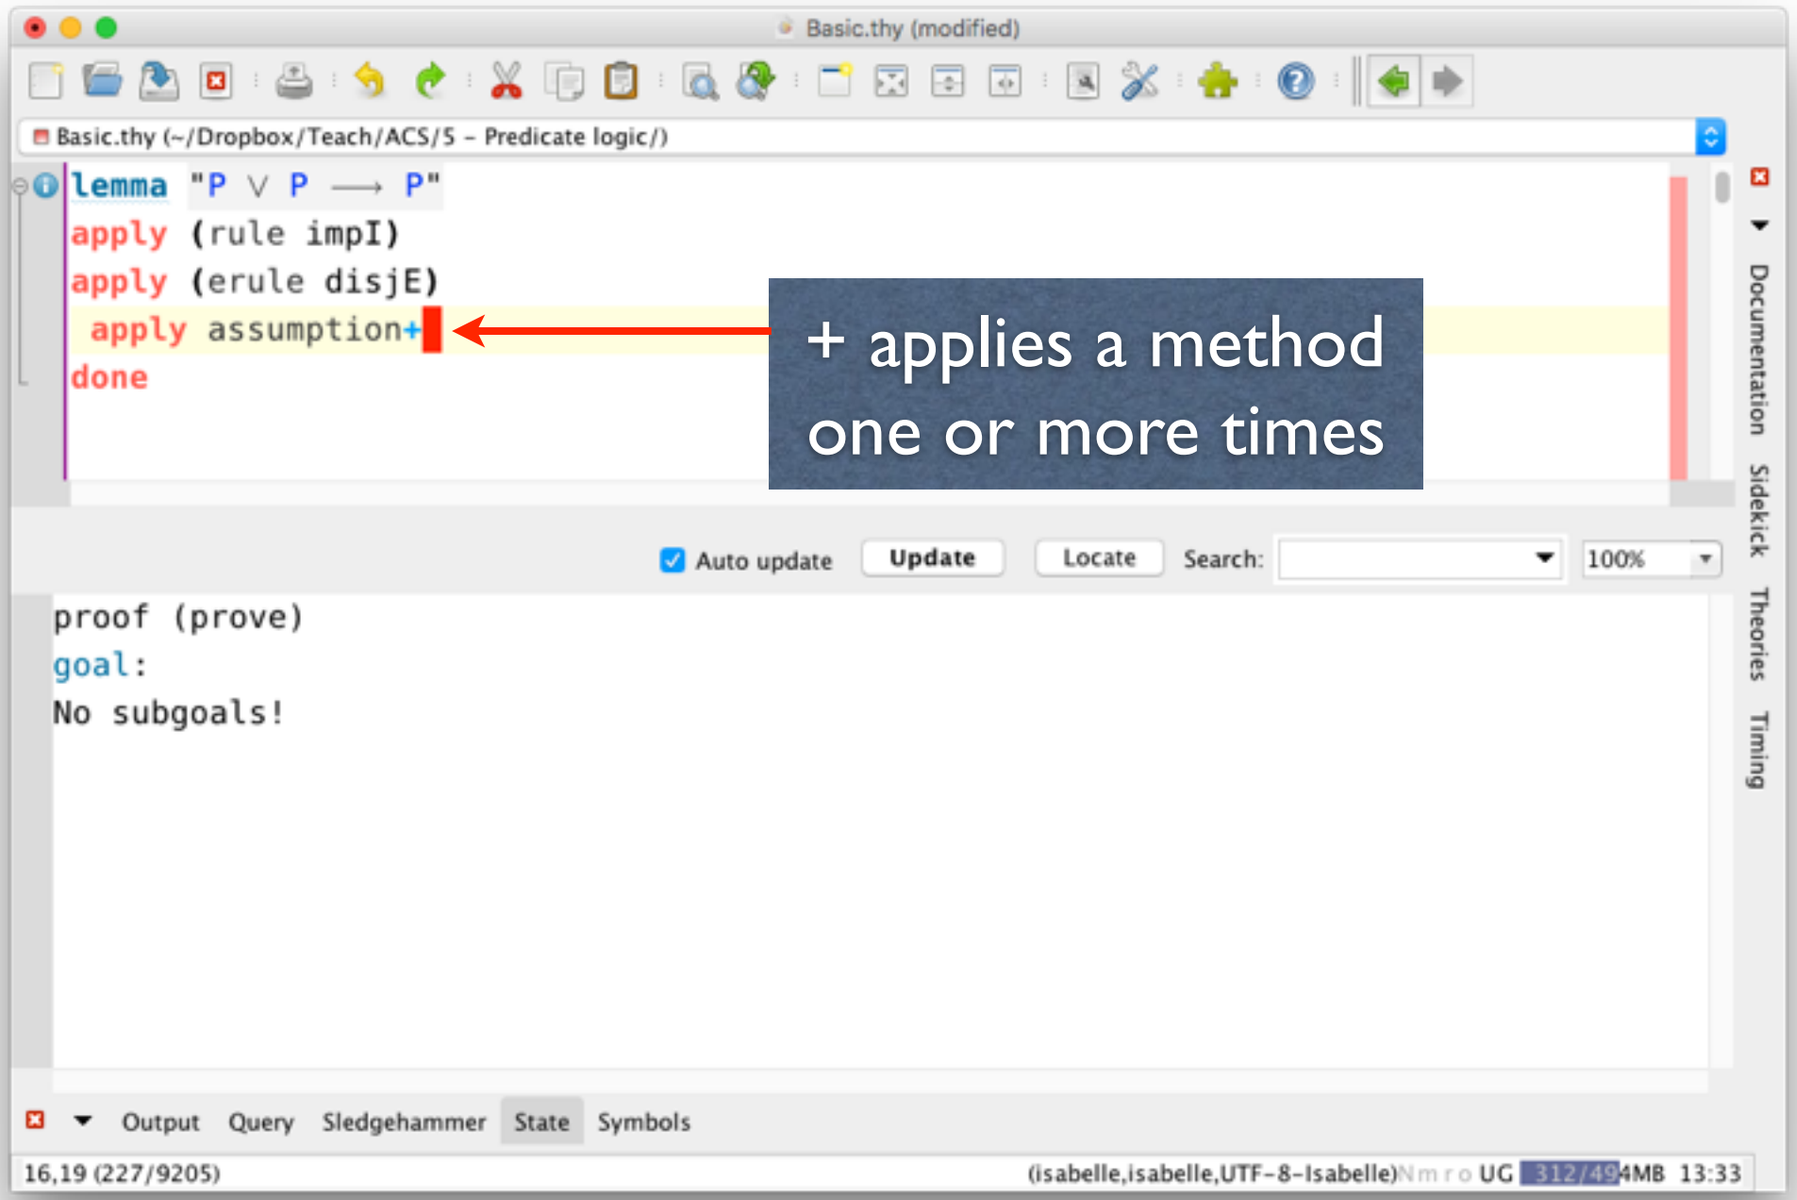

Quantifiers

$$\frac{P(t)}{\exists x.\, P(x)} \qquad \texttt{P(x) ⇒ ∃x.P(x)}$$

$$\frac{\exists x.\, P(x) \qquad \begin{array}{c}[P(x)] \\ \vdots \\ Q\end{array}}{Q} \qquad \texttt{⟦∃x.P(x); ⋀x.P(x)⇒Q⟧ ⇒ Q}$$

meta-universal quantifier
states the variable condition

Isabelle's logical framework includes the typed lambda calculus, so quantifiers can be declared as constants of appropriate type. Variable-binding syntax can also be specified.

A Tiny Quantifier Proof

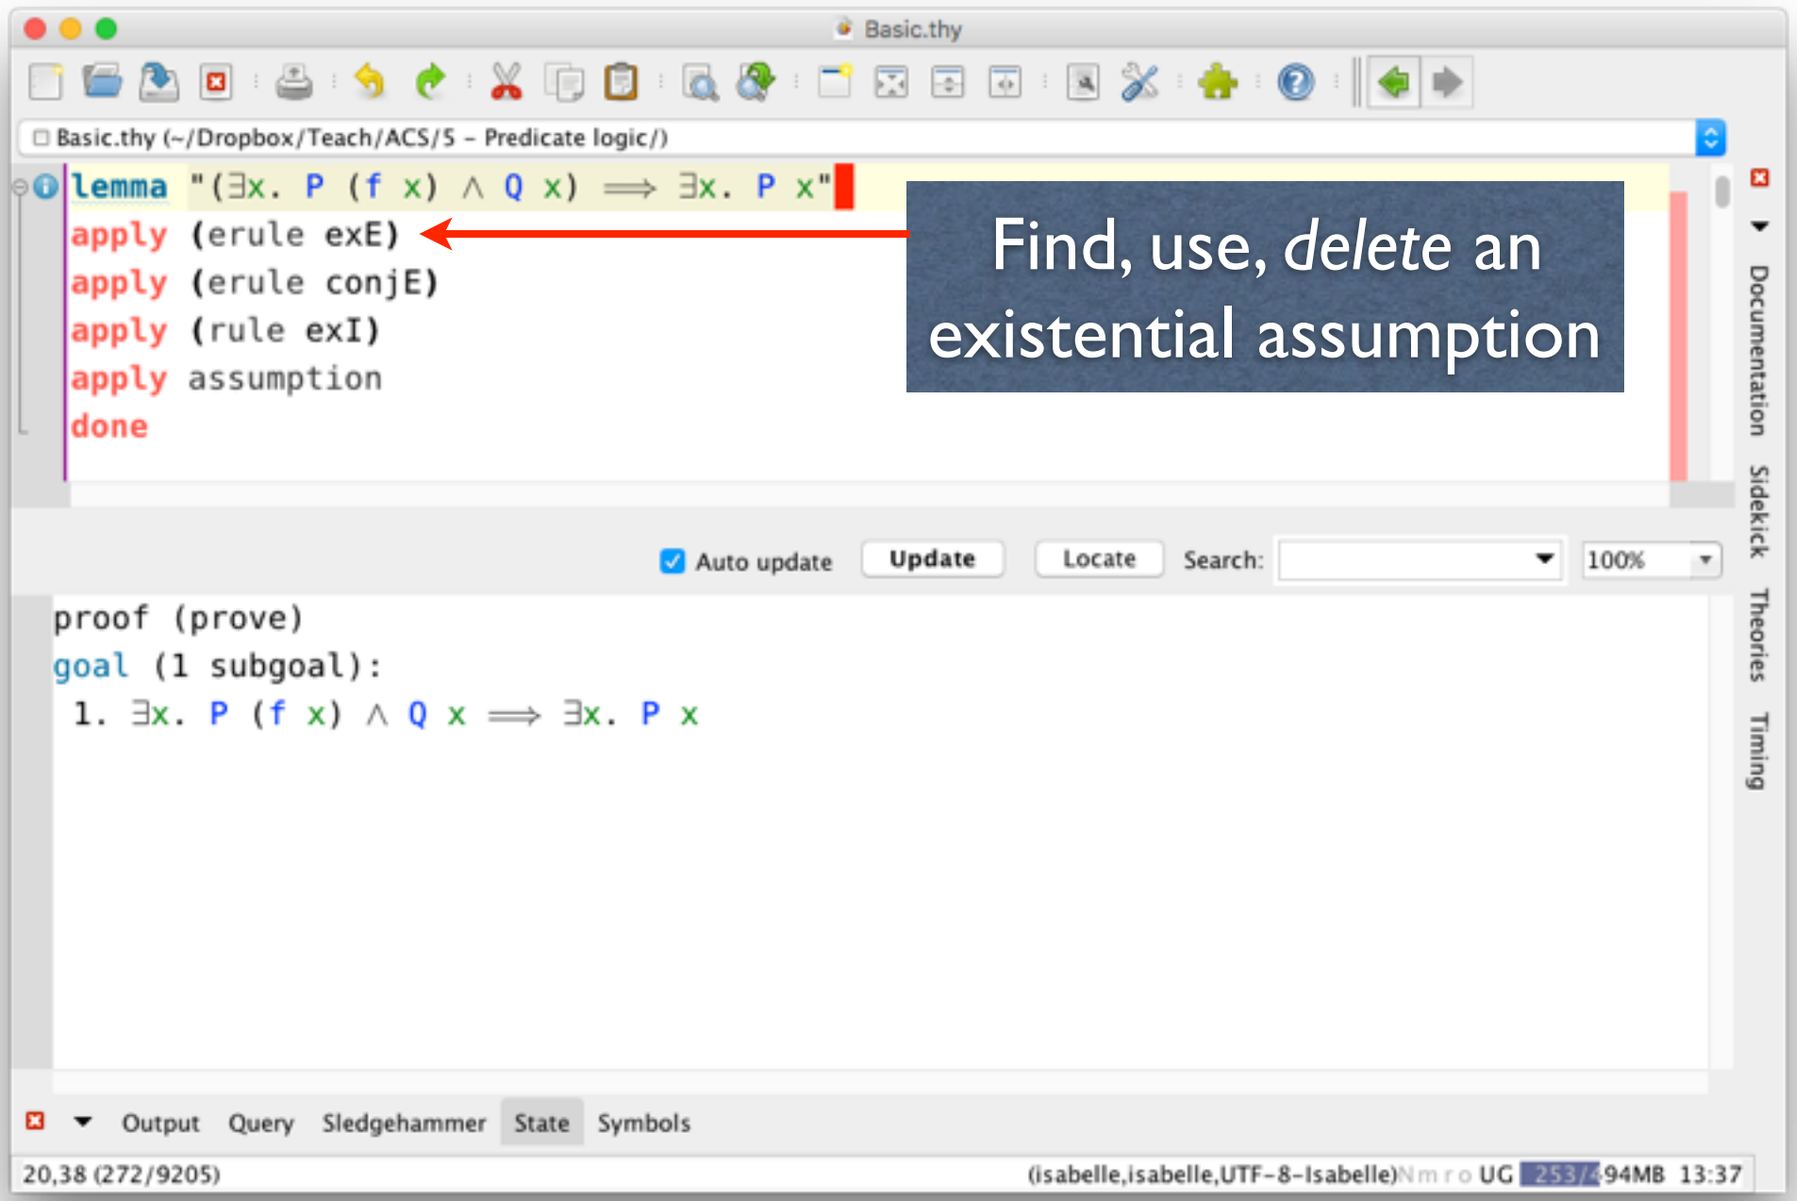

Conjunction Elimination

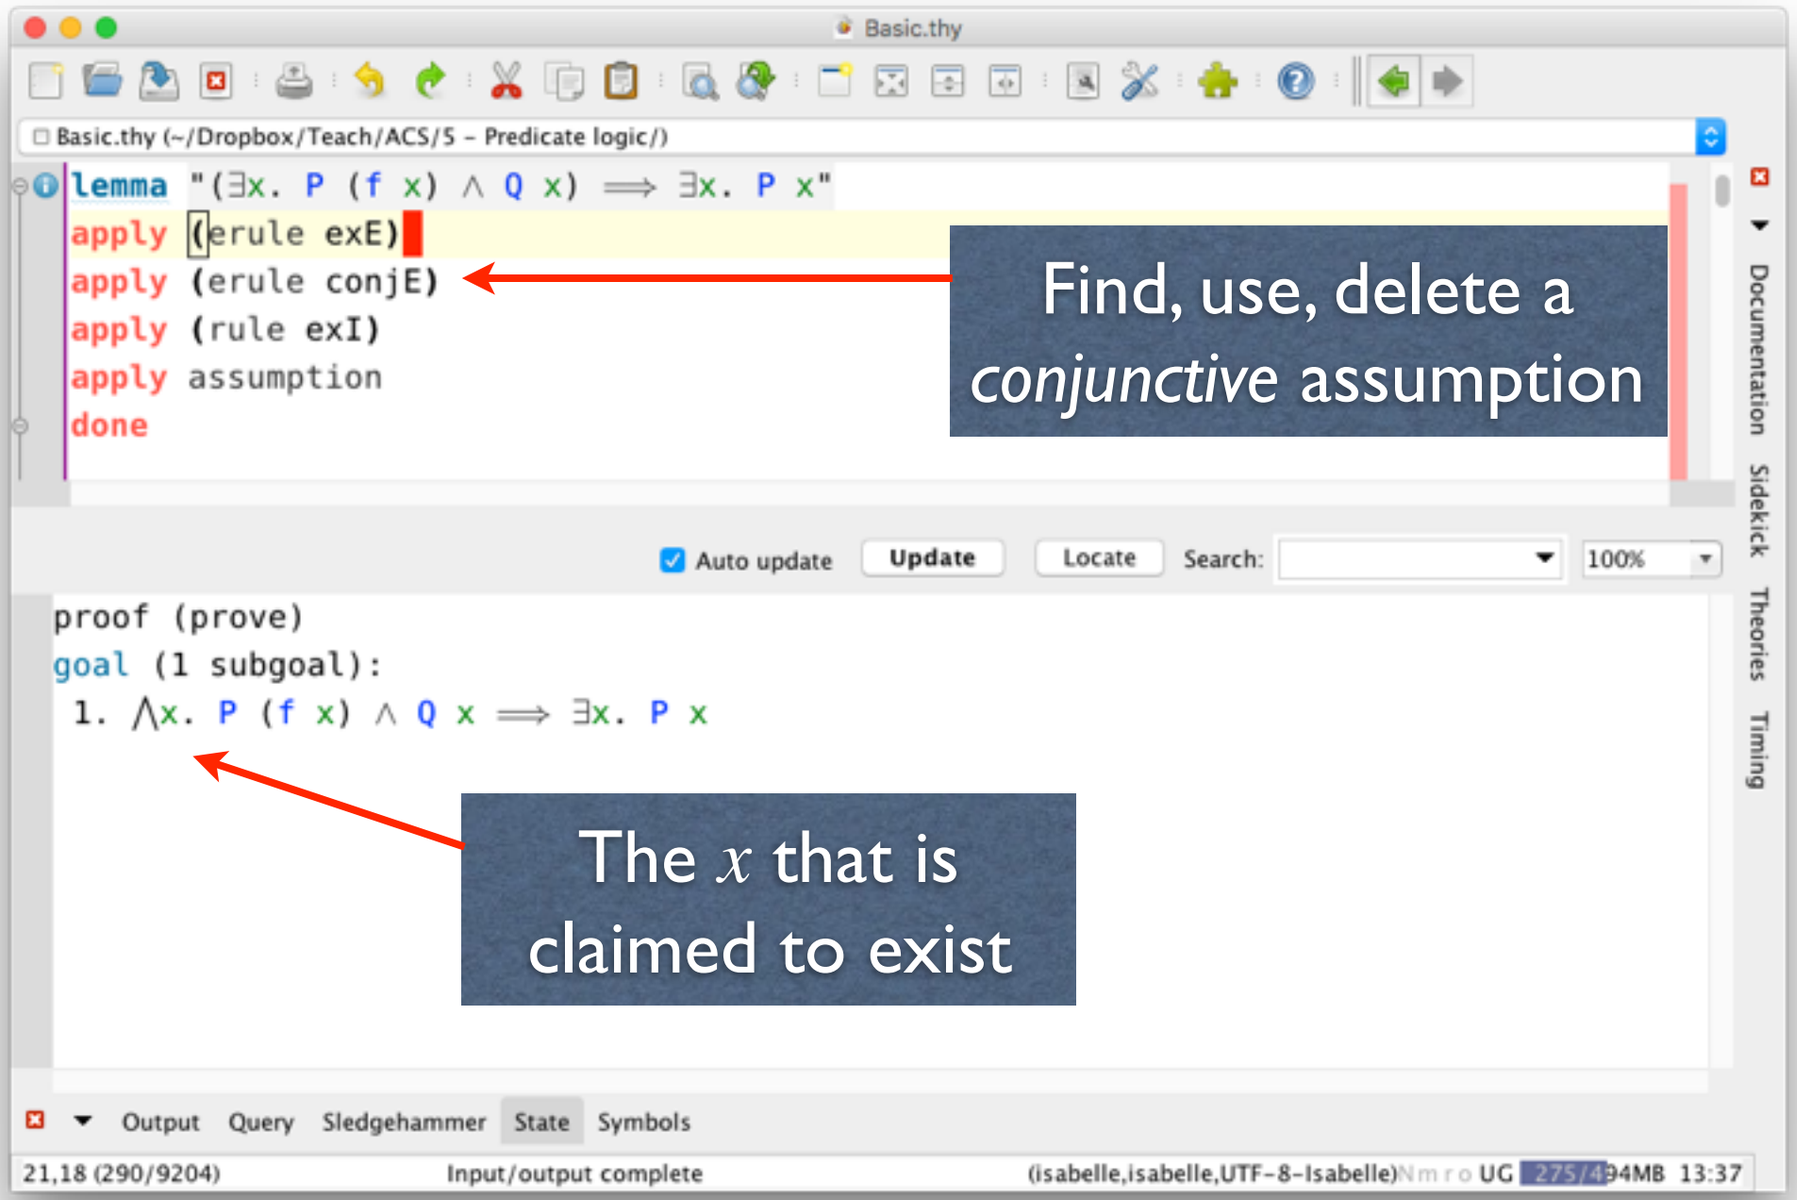

`conjE` is an alternative to the `conjunct1` and `conjunct2`. It has the standard elimination format (like `disjE` for disjunction elimination) so it can be used with the method `erule`.

Now for ∃-Introduction

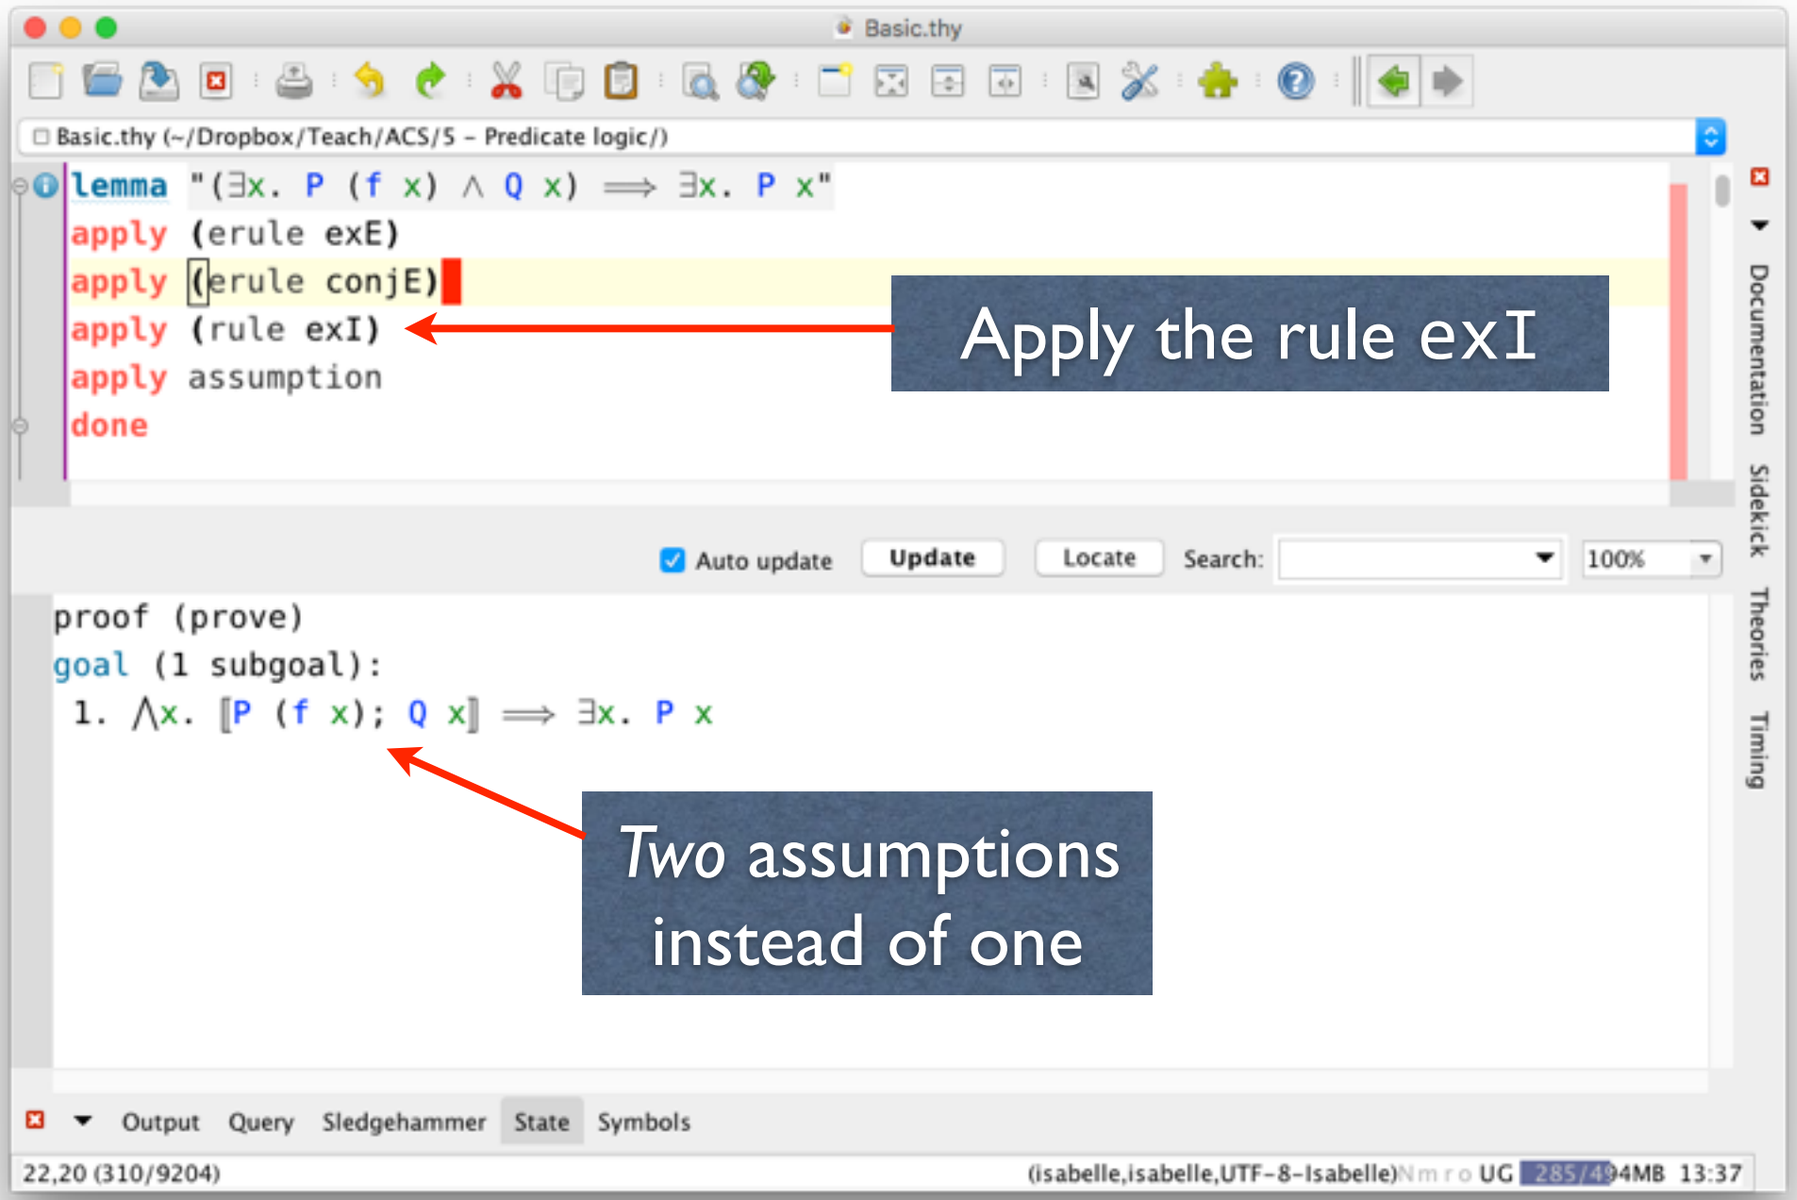

An Unknown for the Witness

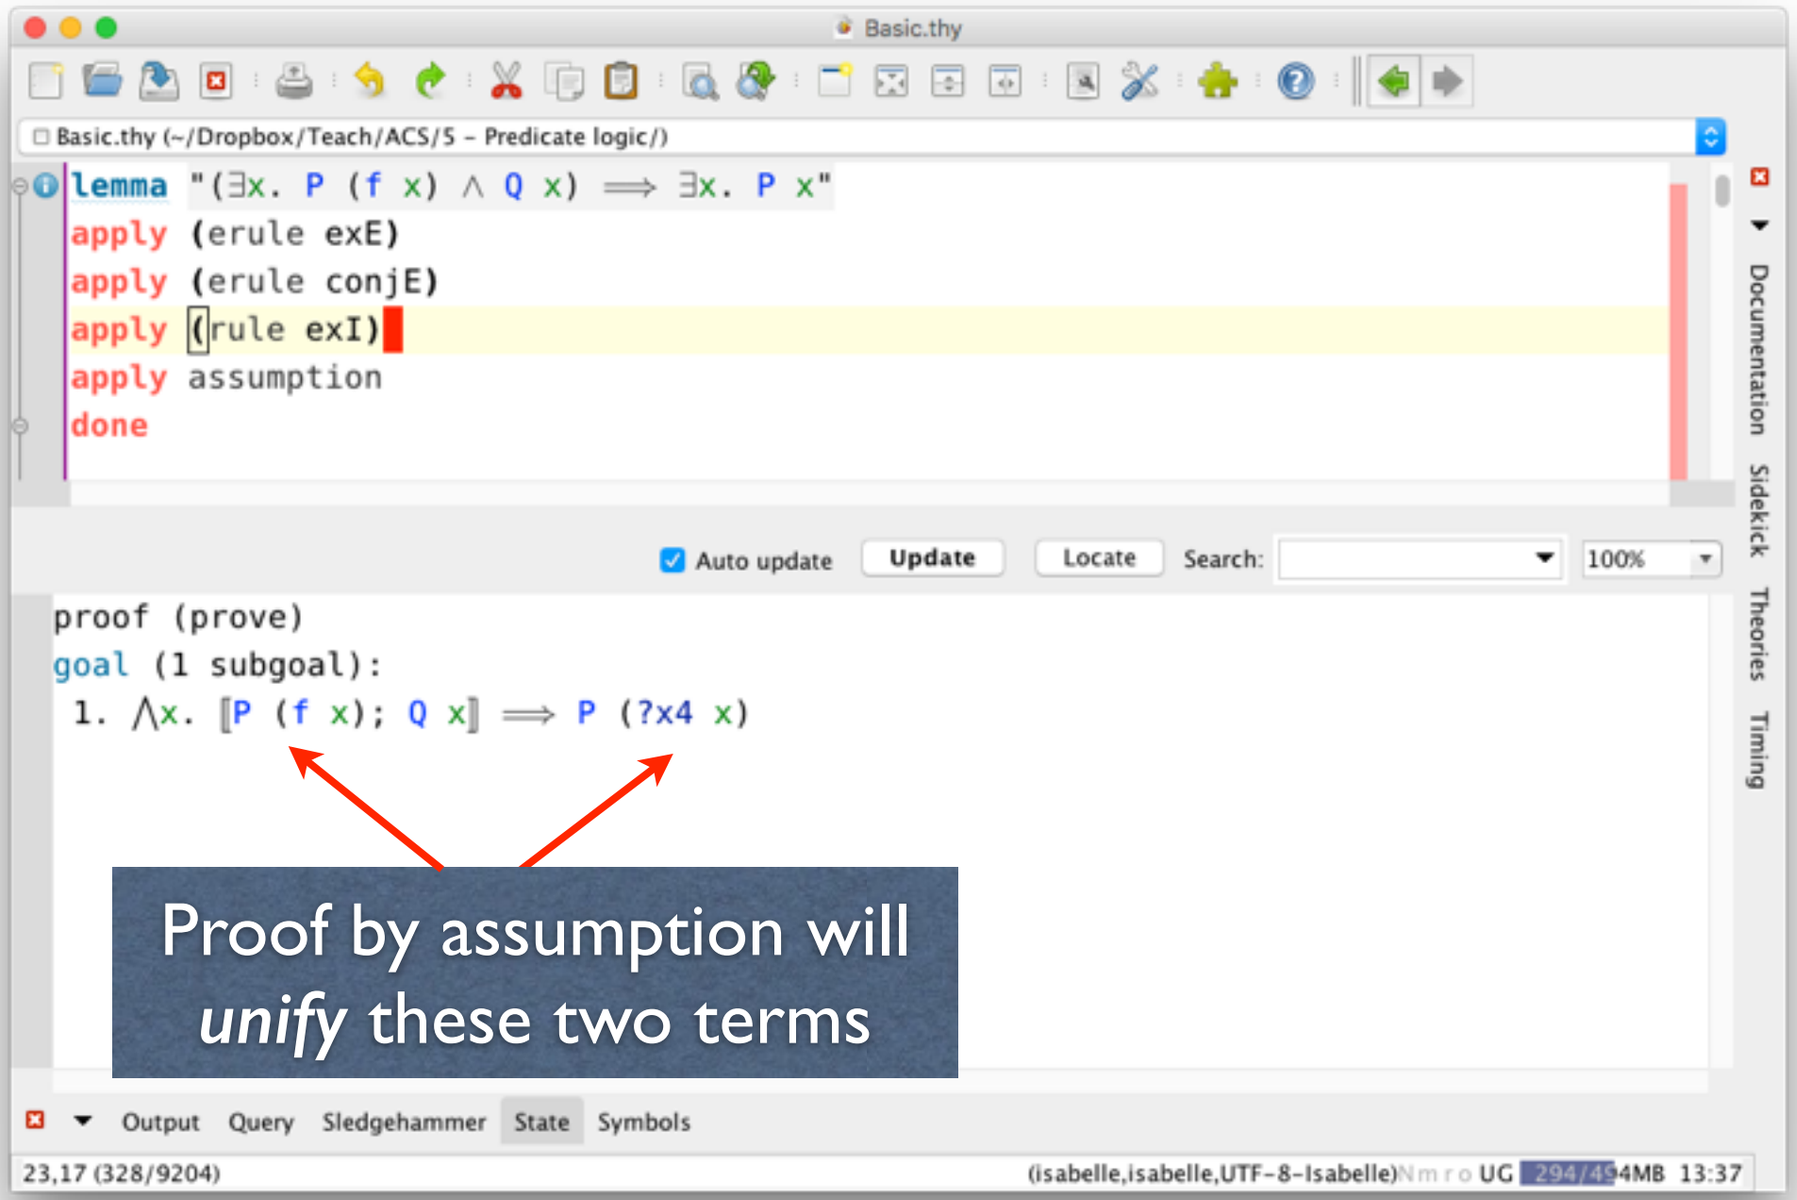

A proof of existence normally requires a **witness**, namely a specific term satisfying the required property. Isabelle allows this choice to be deferred. The structure of the term, in this case ?x4 x, holds information about which bound variables may appear in the witness. Here, x may appear in the witness.

Done!

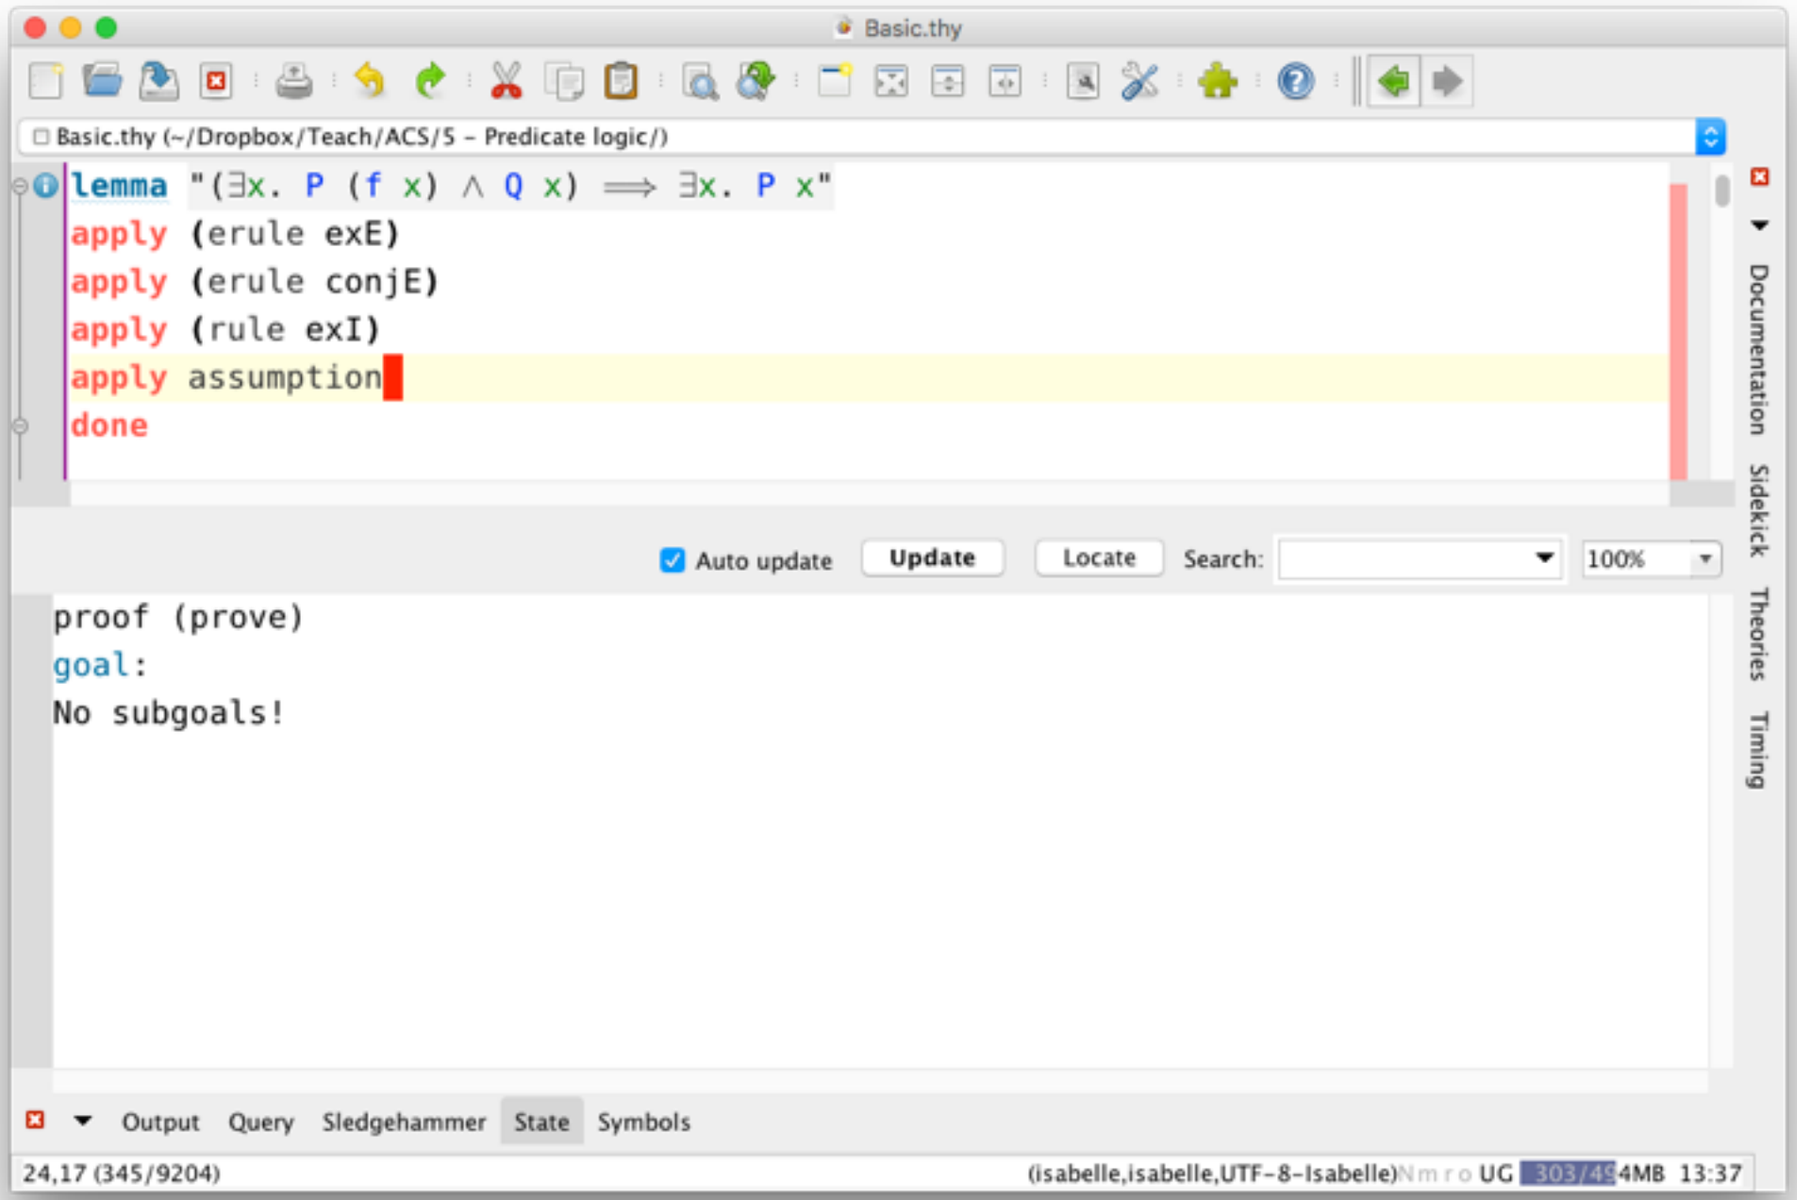
Basic.thy
Basic.thy (~/Dropbox/Teach/ACS/5 - Predicate logic/)
lemma "(∃x. P (f x) ∧ Q x) ⟹ ∃x. P x"
apply (erule exE)
apply (erule conjE)
apply (rule exI)
apply assumption
done
Auto update
Update
Locate
Search:
100%
proof (prove)
goal:
No subgoals!
Output
Query
Sledgehammer
State
Symbols
24,17 (345/9204)
(isabelle,isabelle,UTF-8-Isabelle)
UG
303/494MB
13:37
Documentation
Sidekick
Theories
Timing

Final Remarks

- There are analogous rules for universal quantifiers.
- Higher-order unification respects bound variables, ensuring that quantifier reasoning is sound.
- The examples above illustrate how things work, but typically most logical reasoning is automatic.

Interactive Formal Verification
6: Structured Proofs

Lawrence C Paulson
Computer Laboratory
University of Cambridge

A Proof about "Divides"

$$b \text{ dvd } a \;\leftrightarrow\; (\exists k.\, a = b \times k)$$

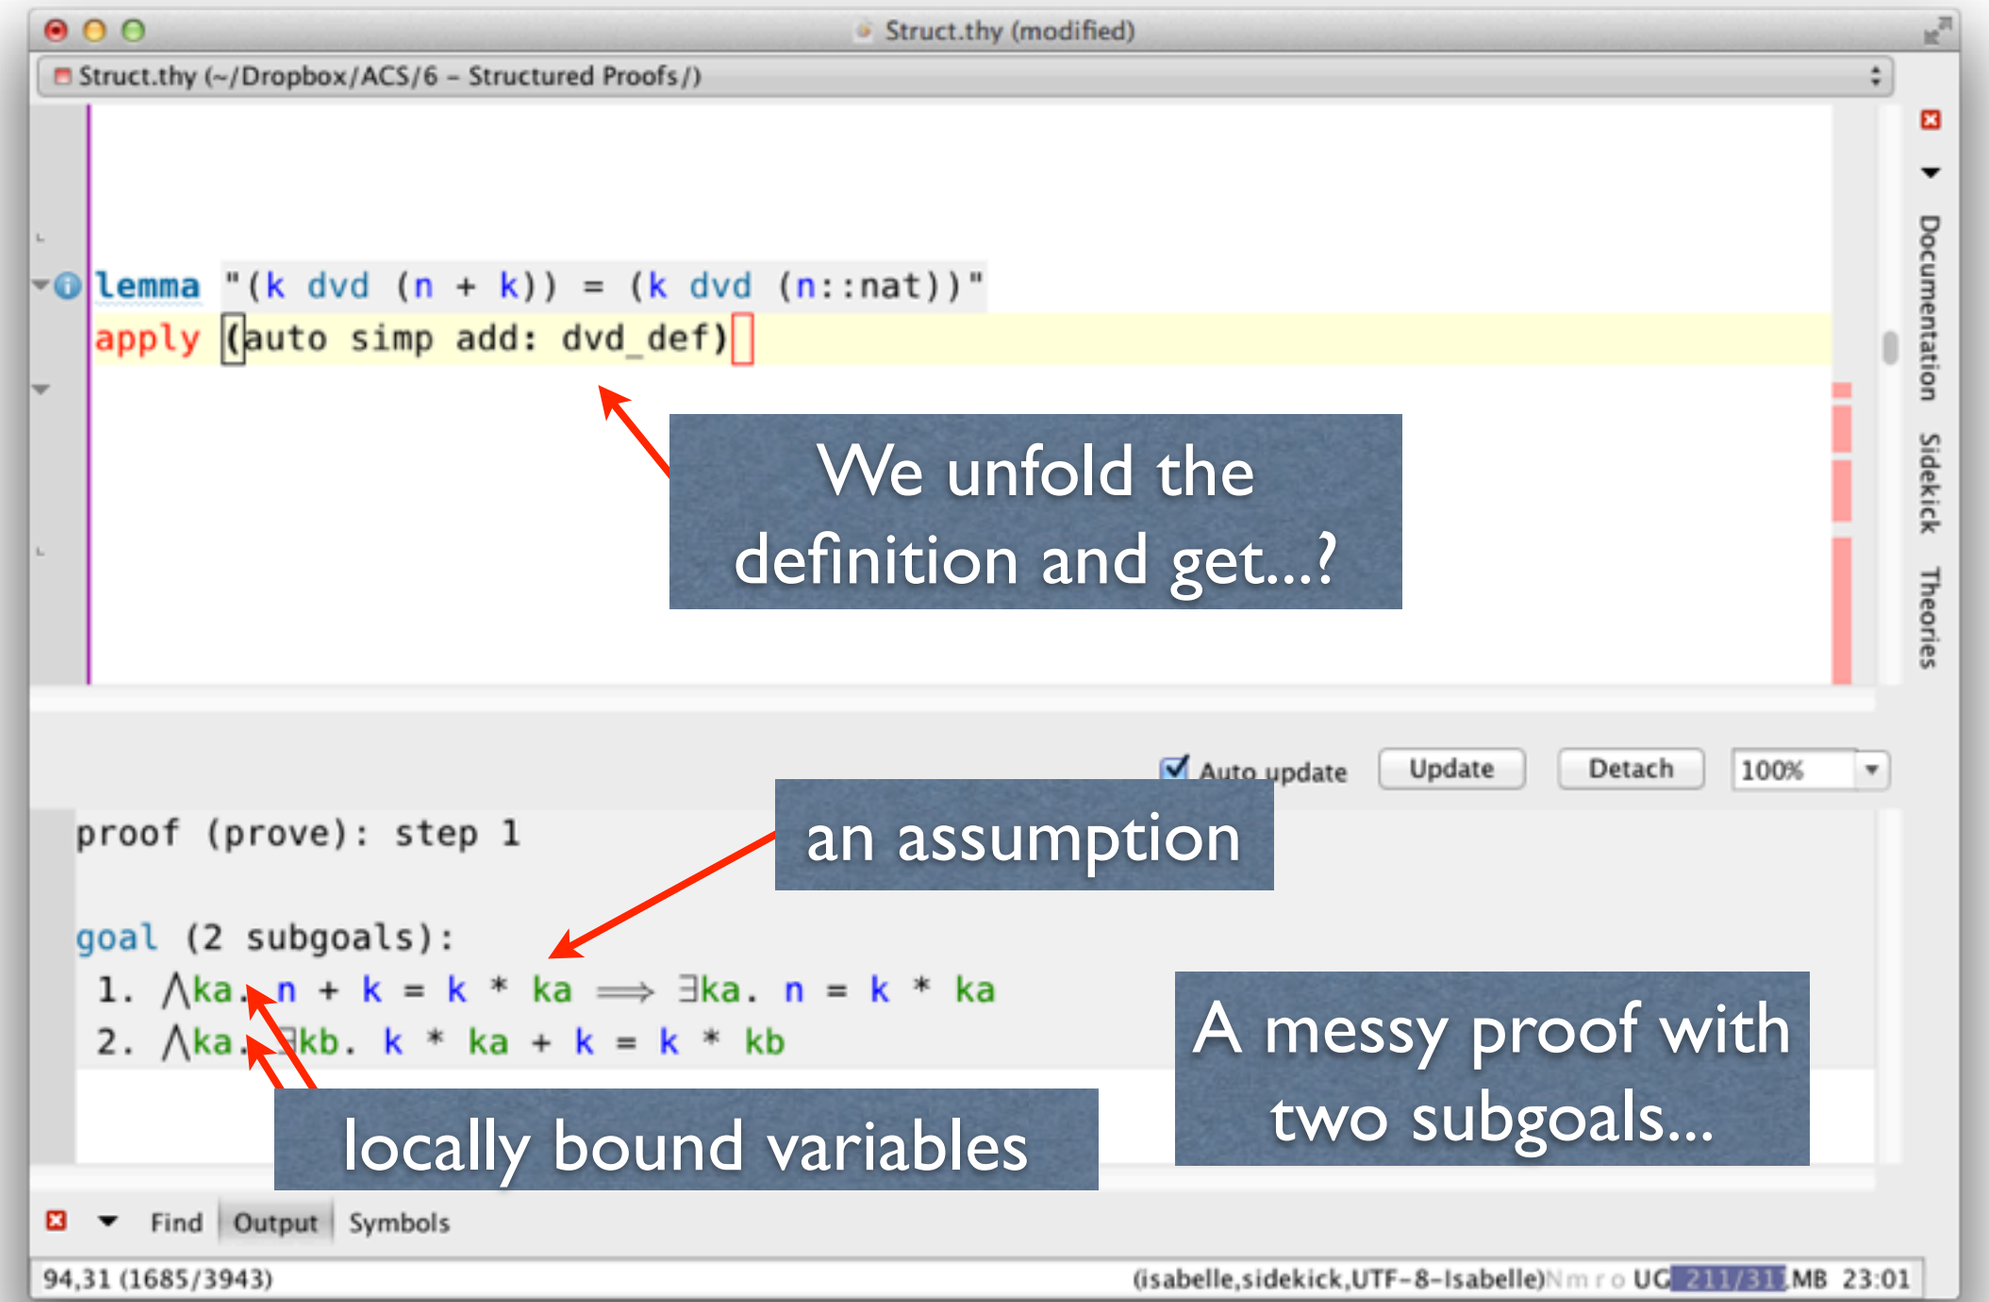

Complex Subgoals

- Isabelle provides many tactics that refer to bound variables and assumptions.
  - Assumptions are often found by matching.
  - Bound variables can be referred to by name, but these names are fragile.
  - *Structured proofs* provide a robust means of referring to these elements by name.
  - Structured proofs are typically verbose but much more readable than linear apply-proofs.

The old-fashioned tactics mentioned above, such as rule_tac, are described in the *Tutorial*, particularly from section 5.7 onwards.

A Structured Proof

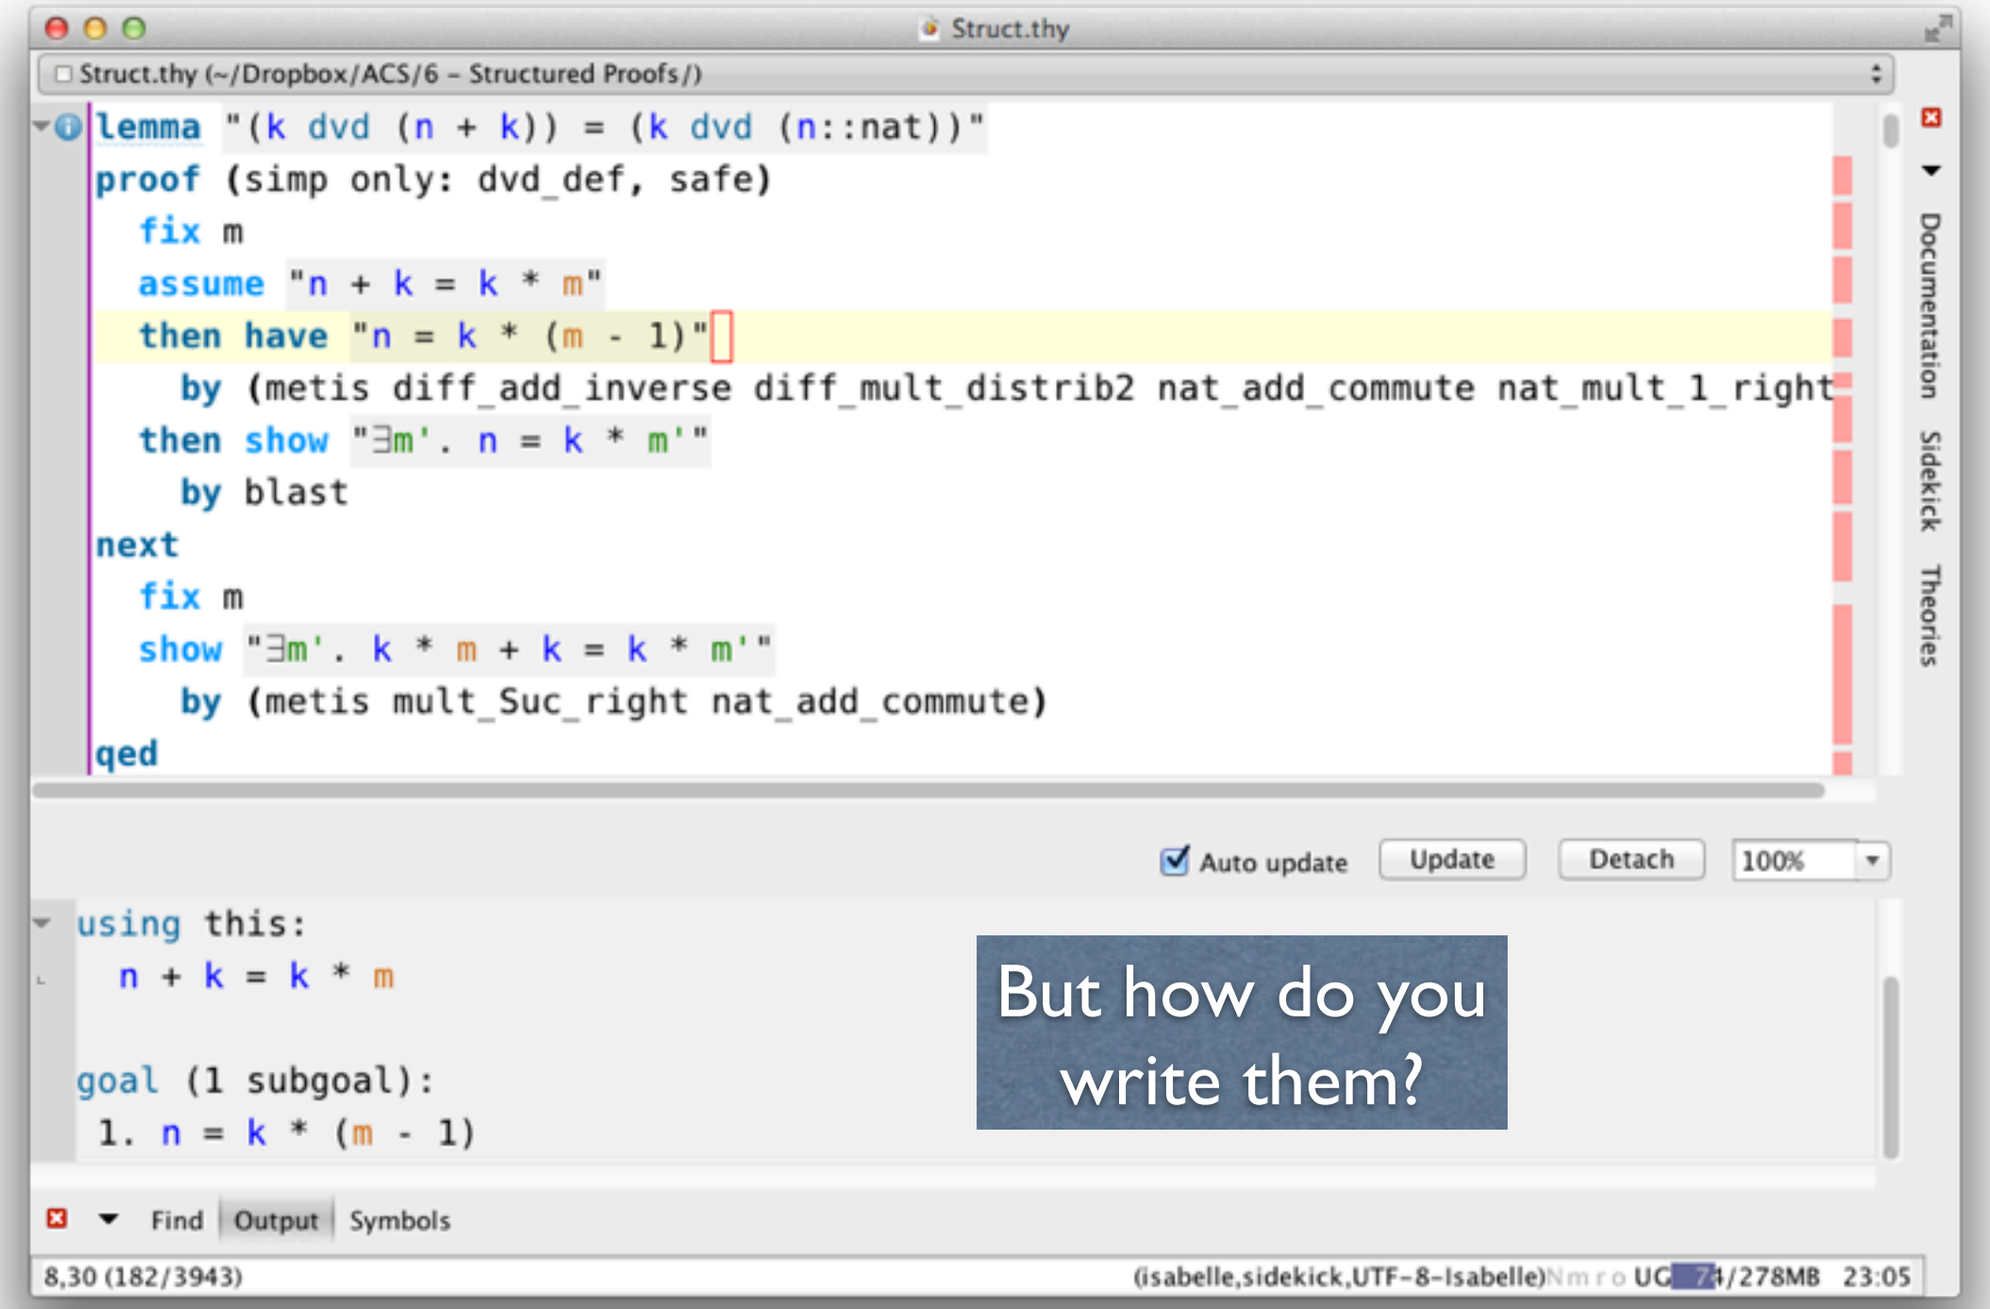

The Elements of Isar

- A *proof context* holds a local variables and assumptions of a subgoal.
  - In a context, the variables are free and the assumptions are simply theorems.
  - Closing a context yields a theorem having the structure of a subgoal.
- The Isar language lets us state and prove intermediate results, express inductions, etc.

The *Tutorial* has little to say about structured proofs. Separate introductions exist, for example, “A Tutorial Introduction to Structured Isar Proofs” by Tobias Nipkow.

Structured proofs can be tricky to write at first. Interaction with proof General is essential: it is virtually impossible to write a structured proof otherwise.

Getting Started

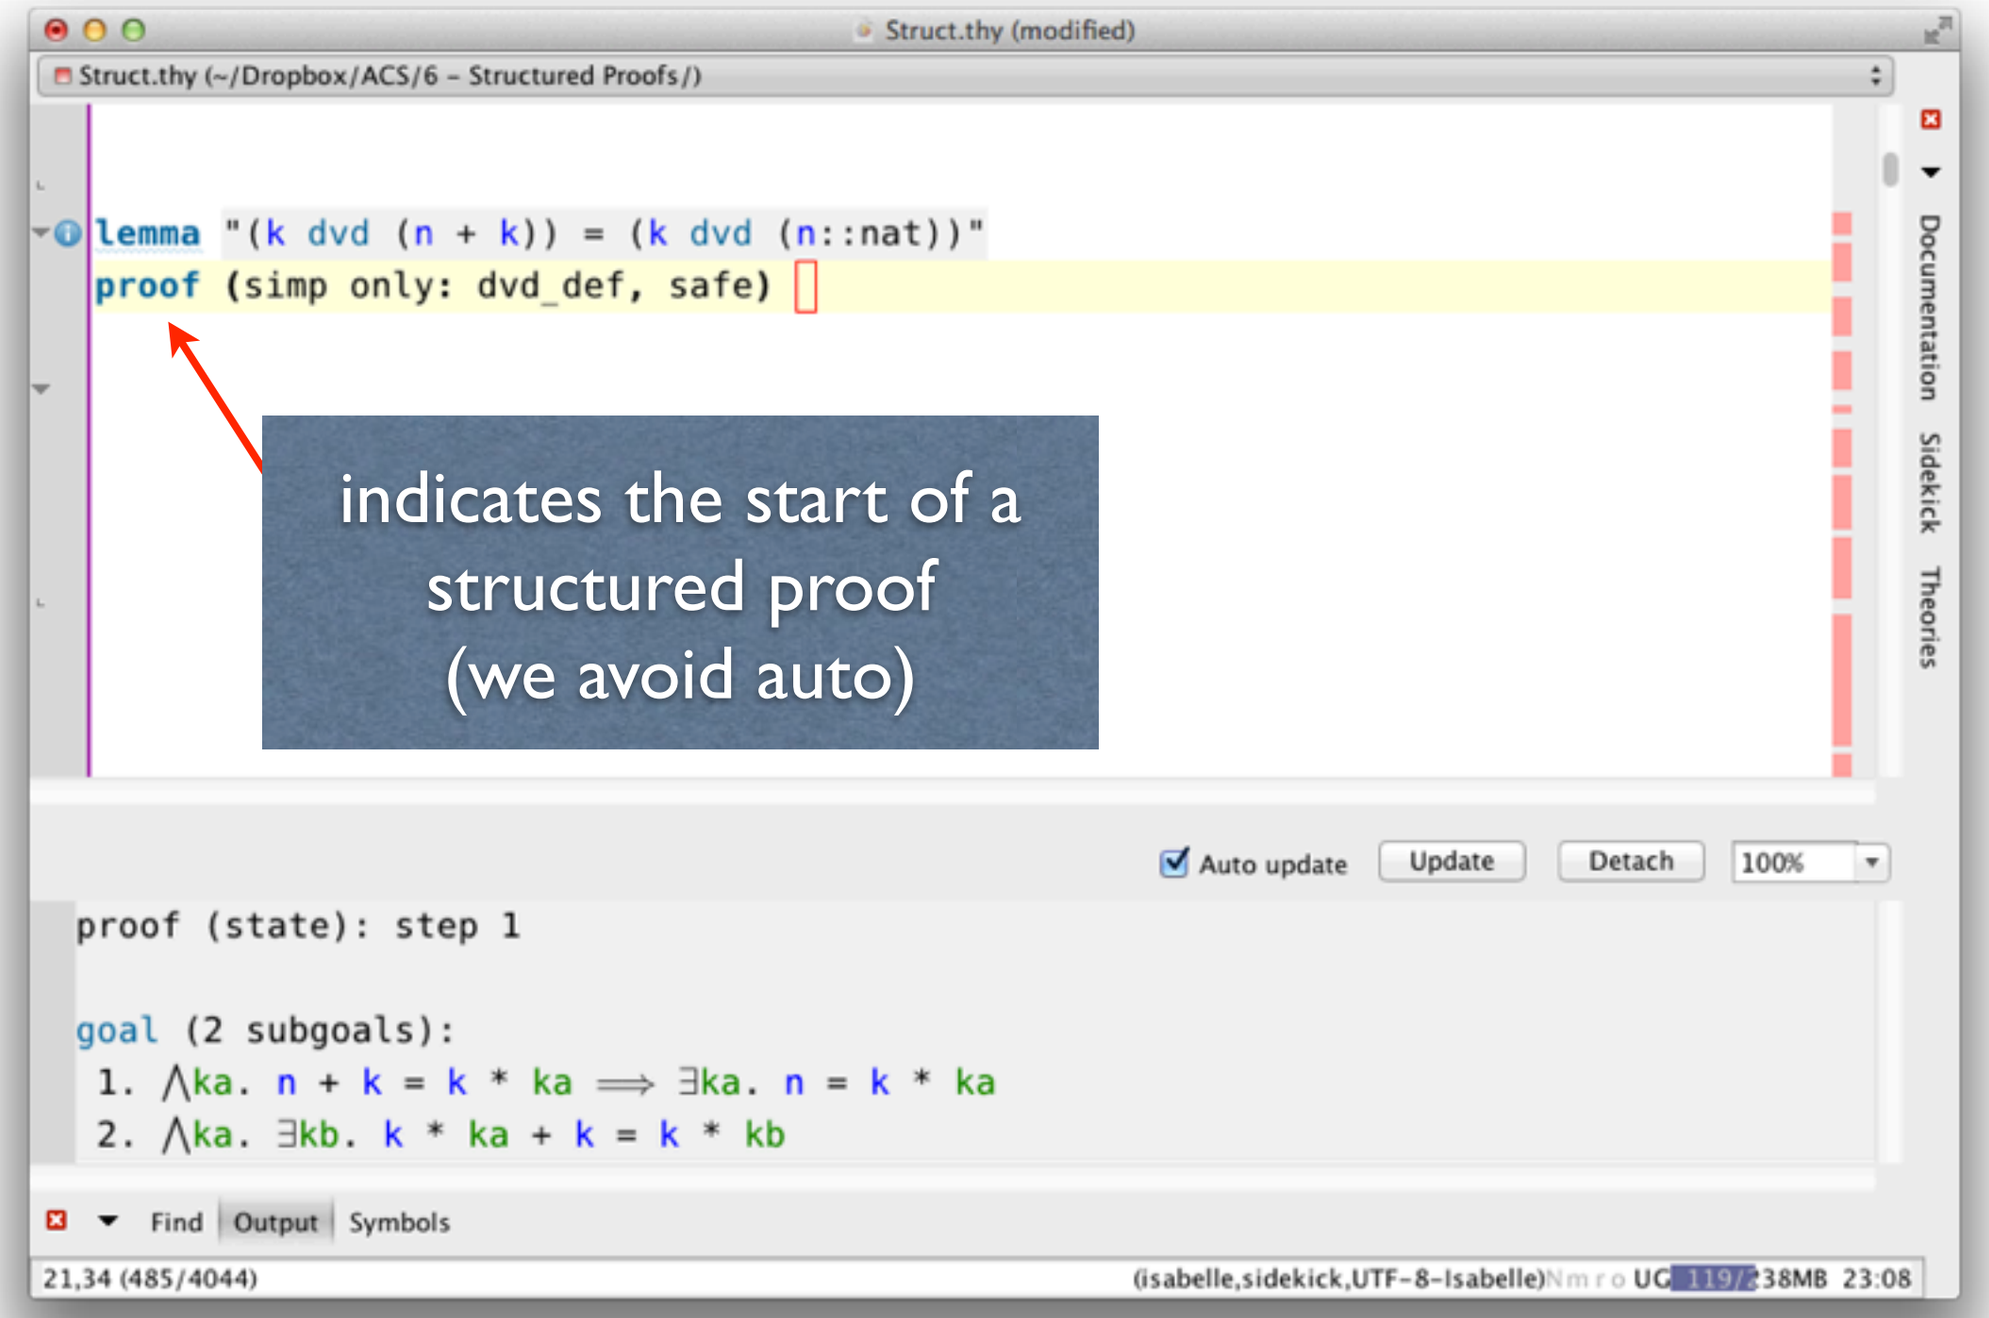

The simplest way to get started is as shown: applying auto with any necessary definitions. The resulting output will then dictate the structure of the final proof.

This style is actually rather fragile. Potentially, a change to `auto` could alter its output, causing a proof based around this precise output to fail. There are two ways of reducing this risk. One is to use a proof method less general than `auto` to unfold the definition of the divides relation and to perform basic logical reasoning. The other is to encapsulate the proofs of the two subgoals in local blocks that can be passed to `auto`; this approach requires a rather sophisticated use of Isar.

Replacing “auto” by “simp only: dvd_def, safe” produces a more robust proof, since these methods are much simpler and more stable than auto.

The Proof Skeleton

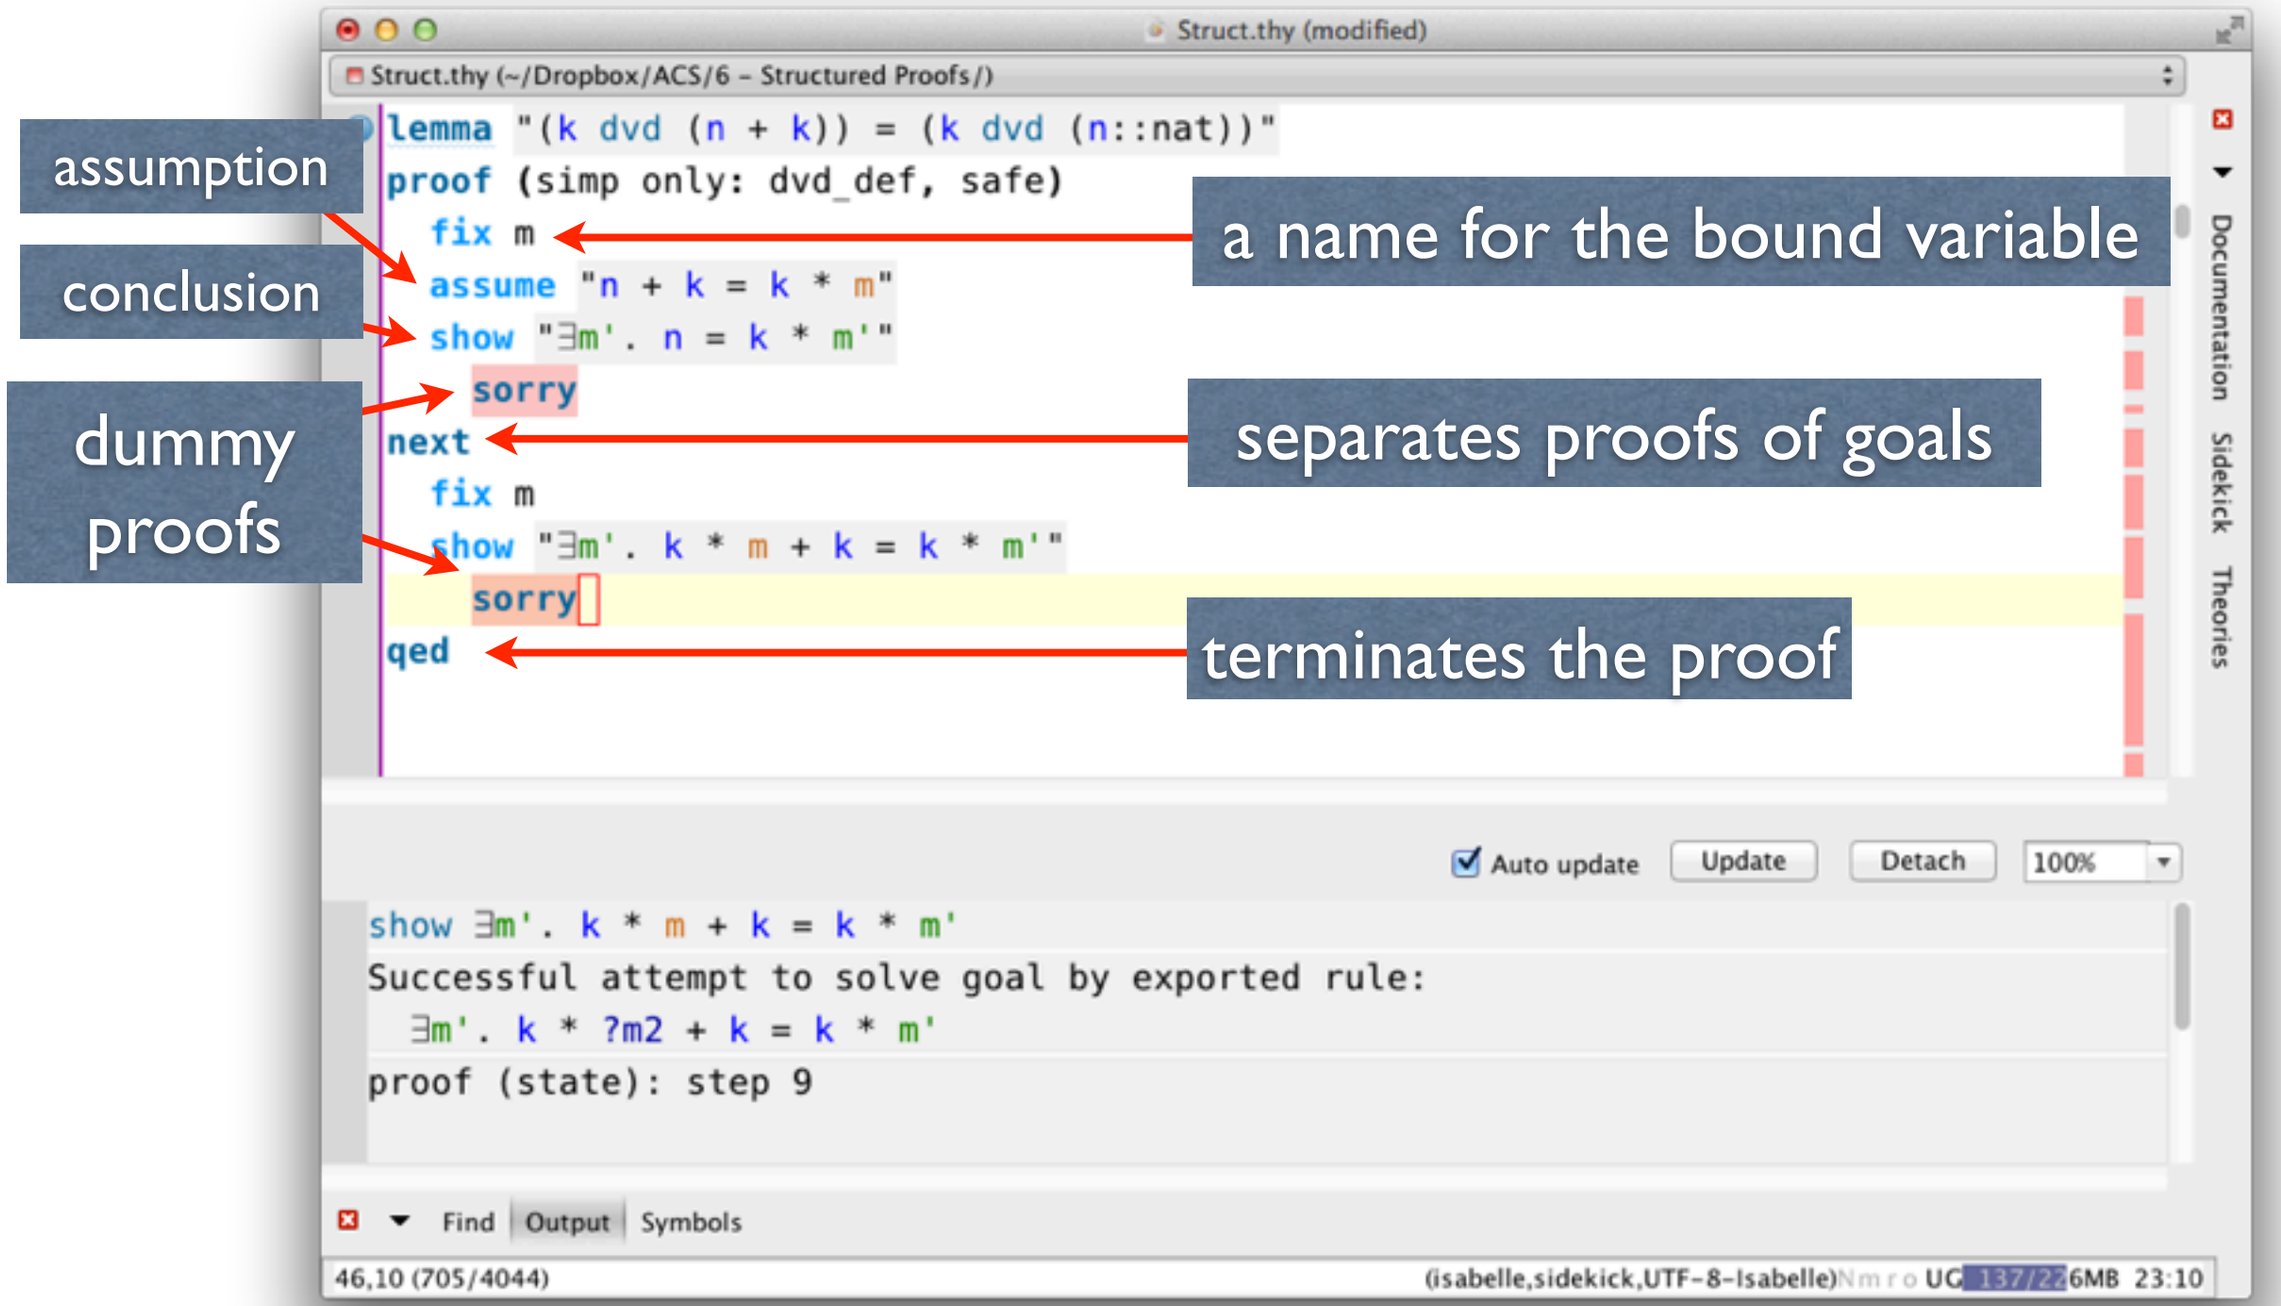

We have used `sorry` to omit the proofs. These dummy proofs allow us to construct the outer shell and confirm that it fits together. We use `show` to state (and eventually prove for real!) the subgoal's conclusion. Since we have renamed the bound variable `ka` to `m`, we must rename it in the assumption and conclusions. The context that we create with `fix`/`assume`, together with the conclusion that we state with `show`, must agree with the original subgoal. Otherwise, Isabelle will generate an error message, "*Local statement will fail to refine any pending goal*".

Fleshing Out that Skeleton

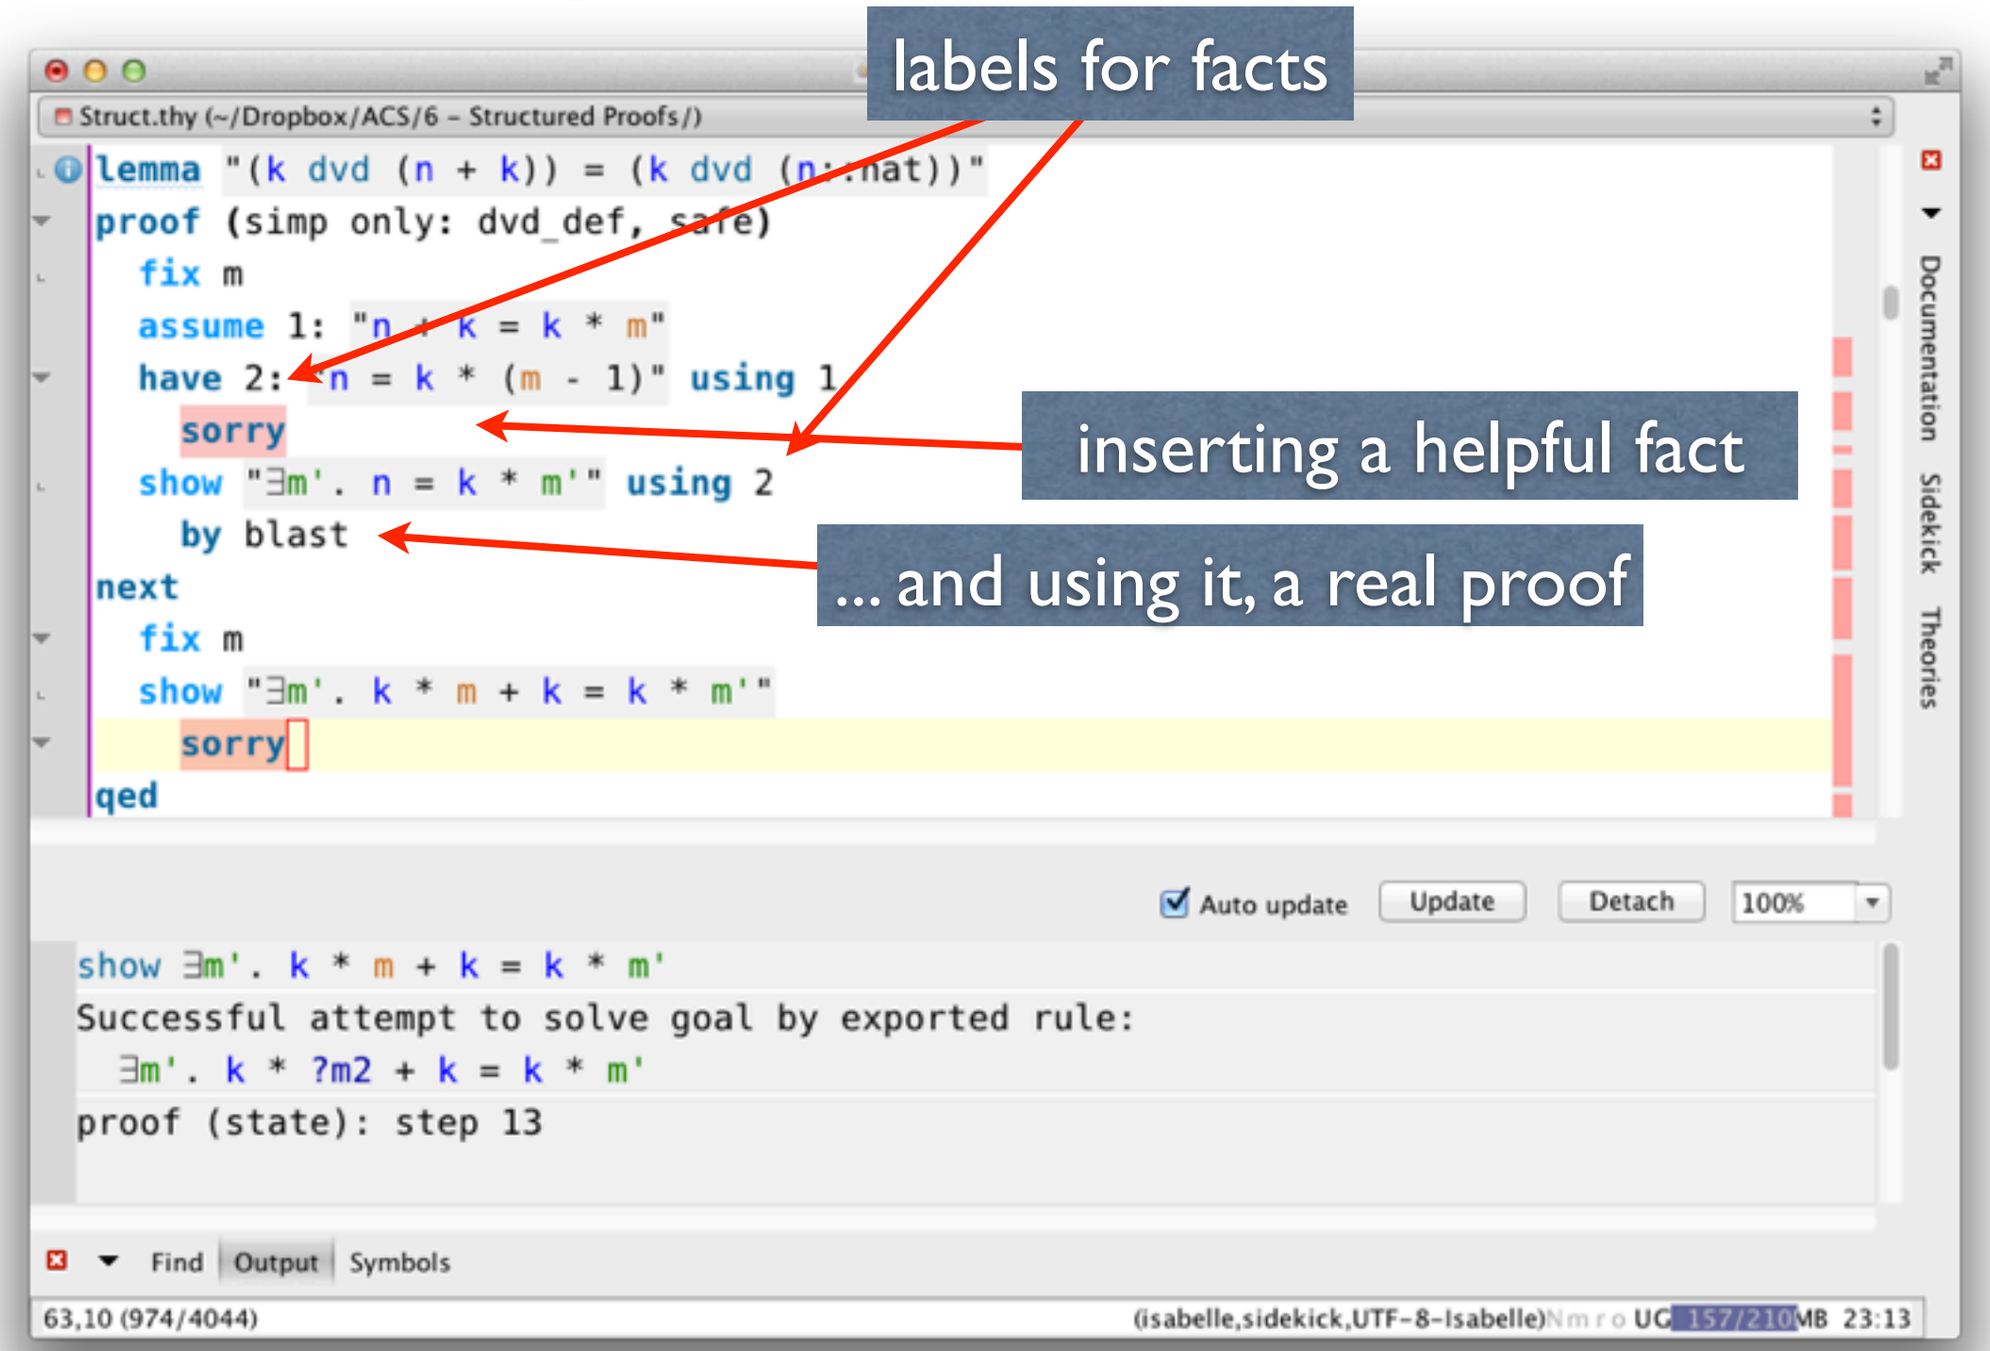

Looking at the first subgoal, we see that it would help to transform the assumption to resemble the body of the quantified formula that is the conclusion. Proving that conclusion should then be trivial, because the existential witness (`m-1`) is explicit. We use `sorry` to obtain this intermediate result, then confirm that the conclusion is provable from it using `blast`. Because it is a one line proof, we write it using “by”. It is permissible to insert a string of “`apply`” commands followed by “`done`”, but that looks ugly.

We give labels to the assumption and the intermediate result for easy reference. We can then write “`using 1`”, for example, to indicate that the proof refers to the designated fact. However, referring to the previous result is extremely common, and soon we shall streamline this proof to eliminate the labels. Also, labels do not have to be integers: they can be any Isabelle identifiers.

Completing the Proof

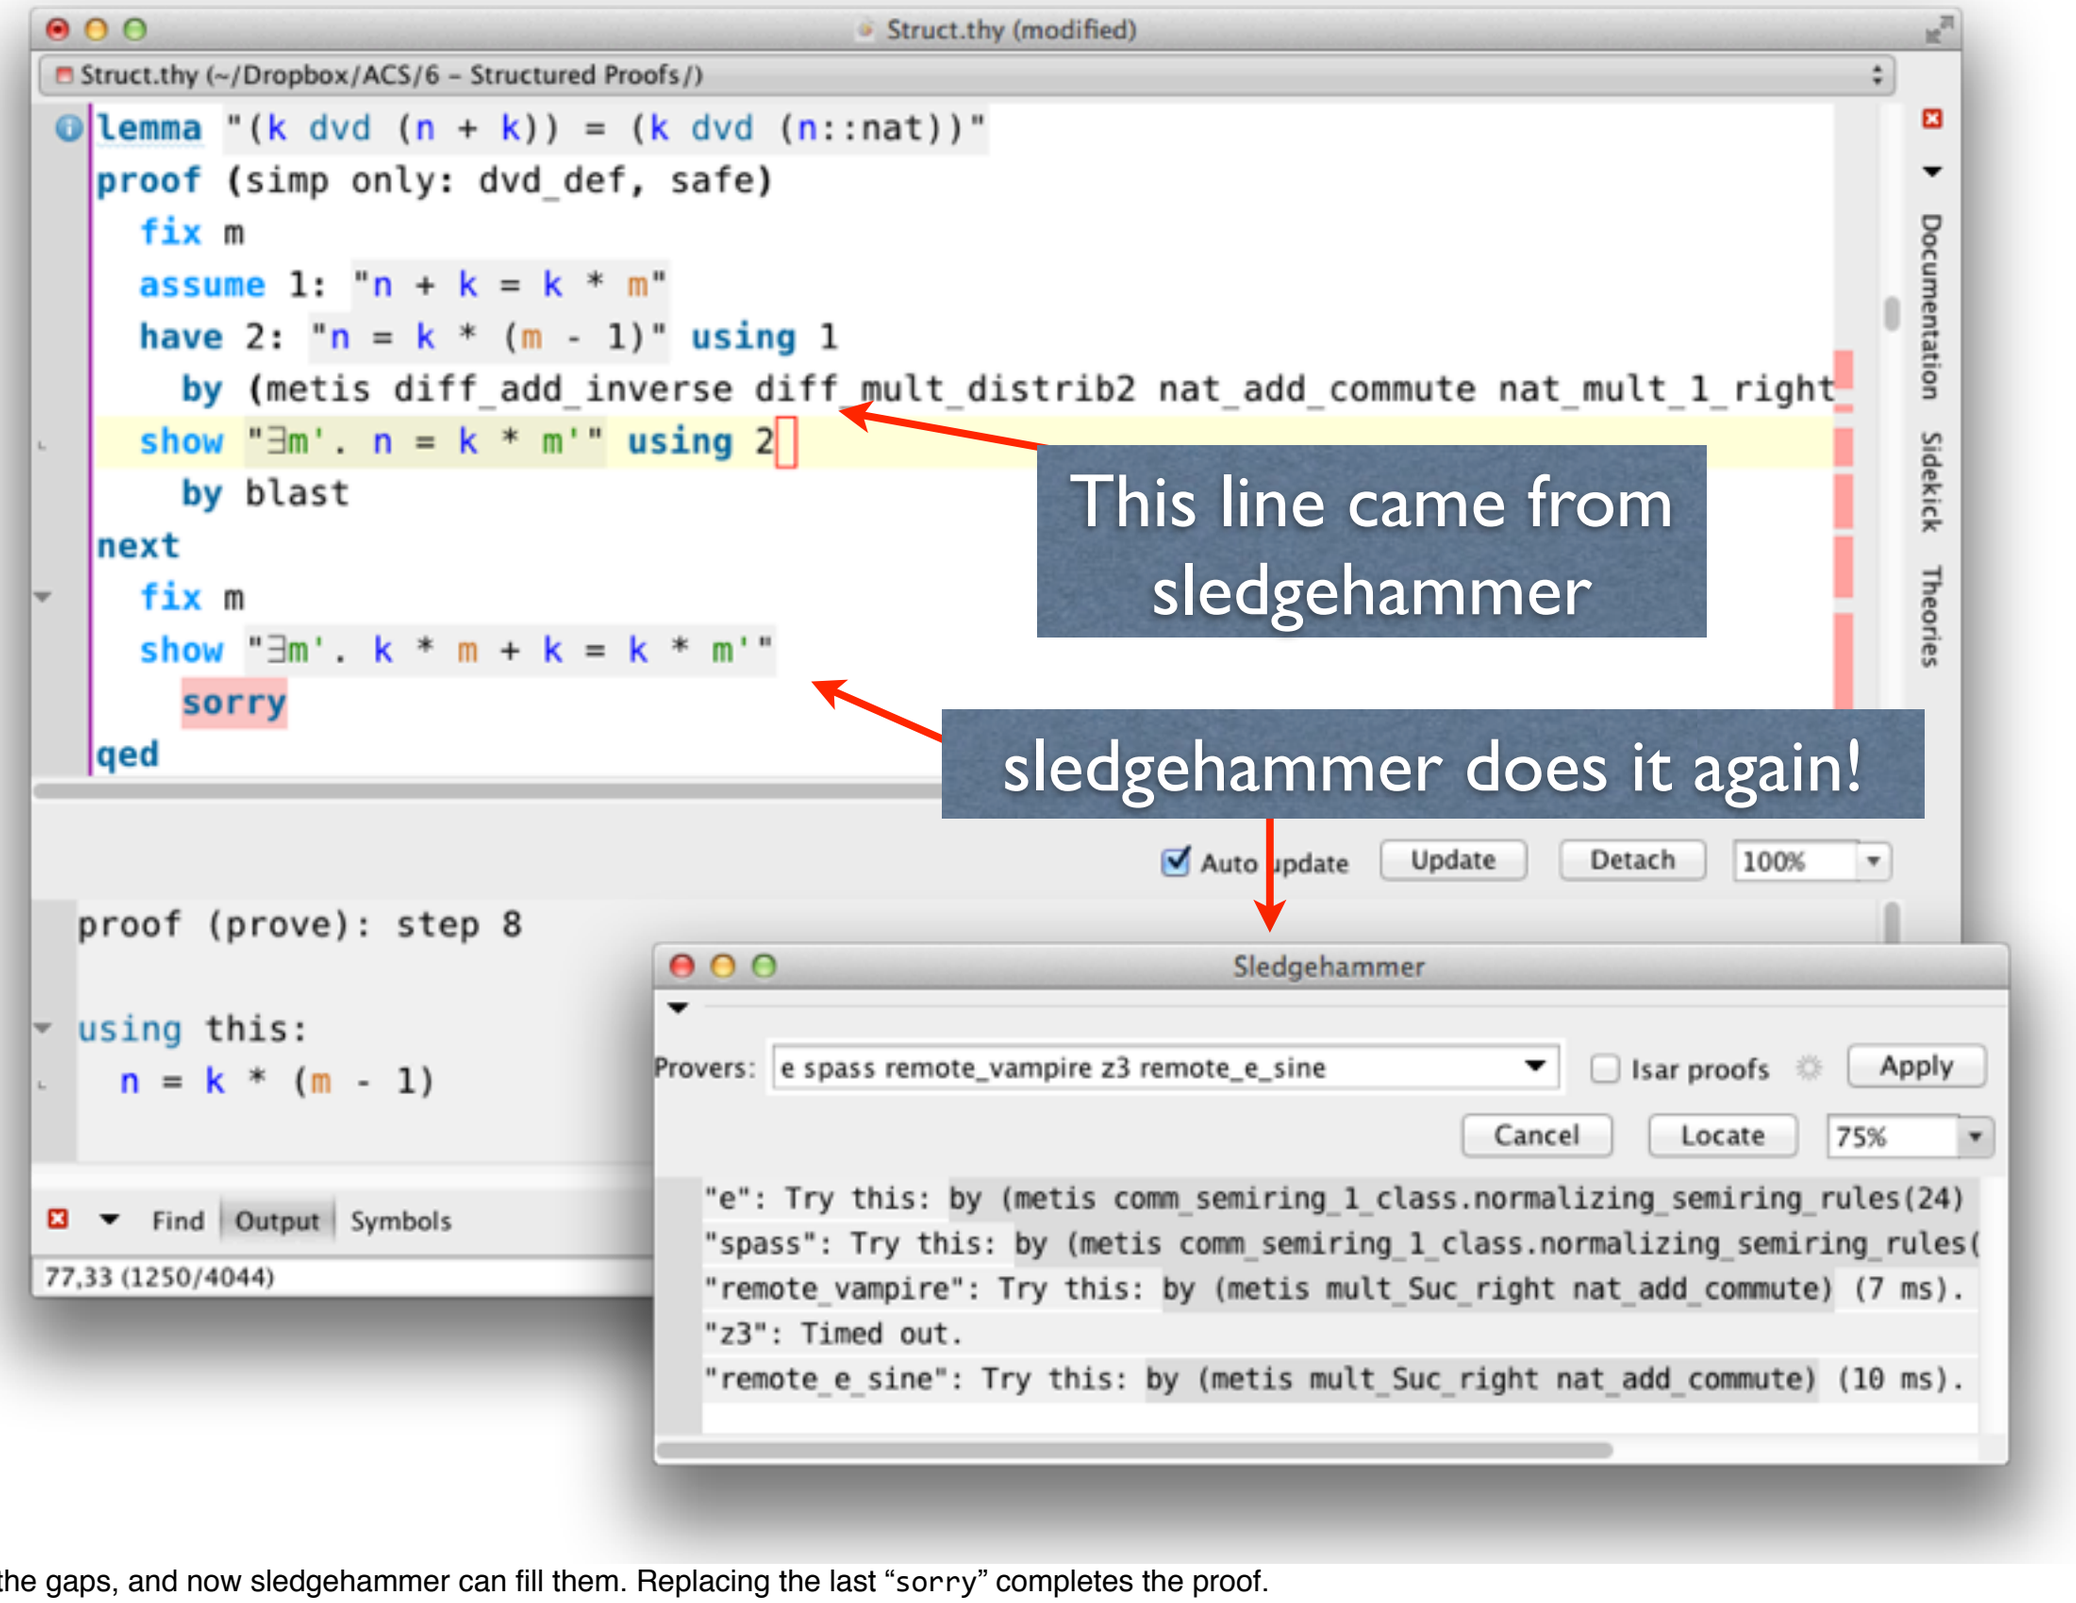

We have narrowed the gaps, and now sledgehammer can fill them. Replacing the last "`sorry`" completes the proof.

There is of course no need to follow this sort of top-down development. It is one approach that is particularly simple for beginners.

Streamlining the Proof

```
assume 1: "n + k = k * m"
have 2: "n = k * (m - 1)" using 1
  sorry
show "∃m'. n = k * m'" using 2
```

```
assume "n + k = k * m"
then have "n = k * (m - 1)"
  sorry
then show "∃m'. n = k * m'"
```

`using` the previous fact without mentioning labels

- `then have` *or* `hence`
- `then show` *or* `thus`

There are numerous other tricks of this sort!

Avoiding the contracted forms “hence” and “thus” may be better for readability, emphasising the role of “then”, which uses the previous fact.

Another Proof Skeleton

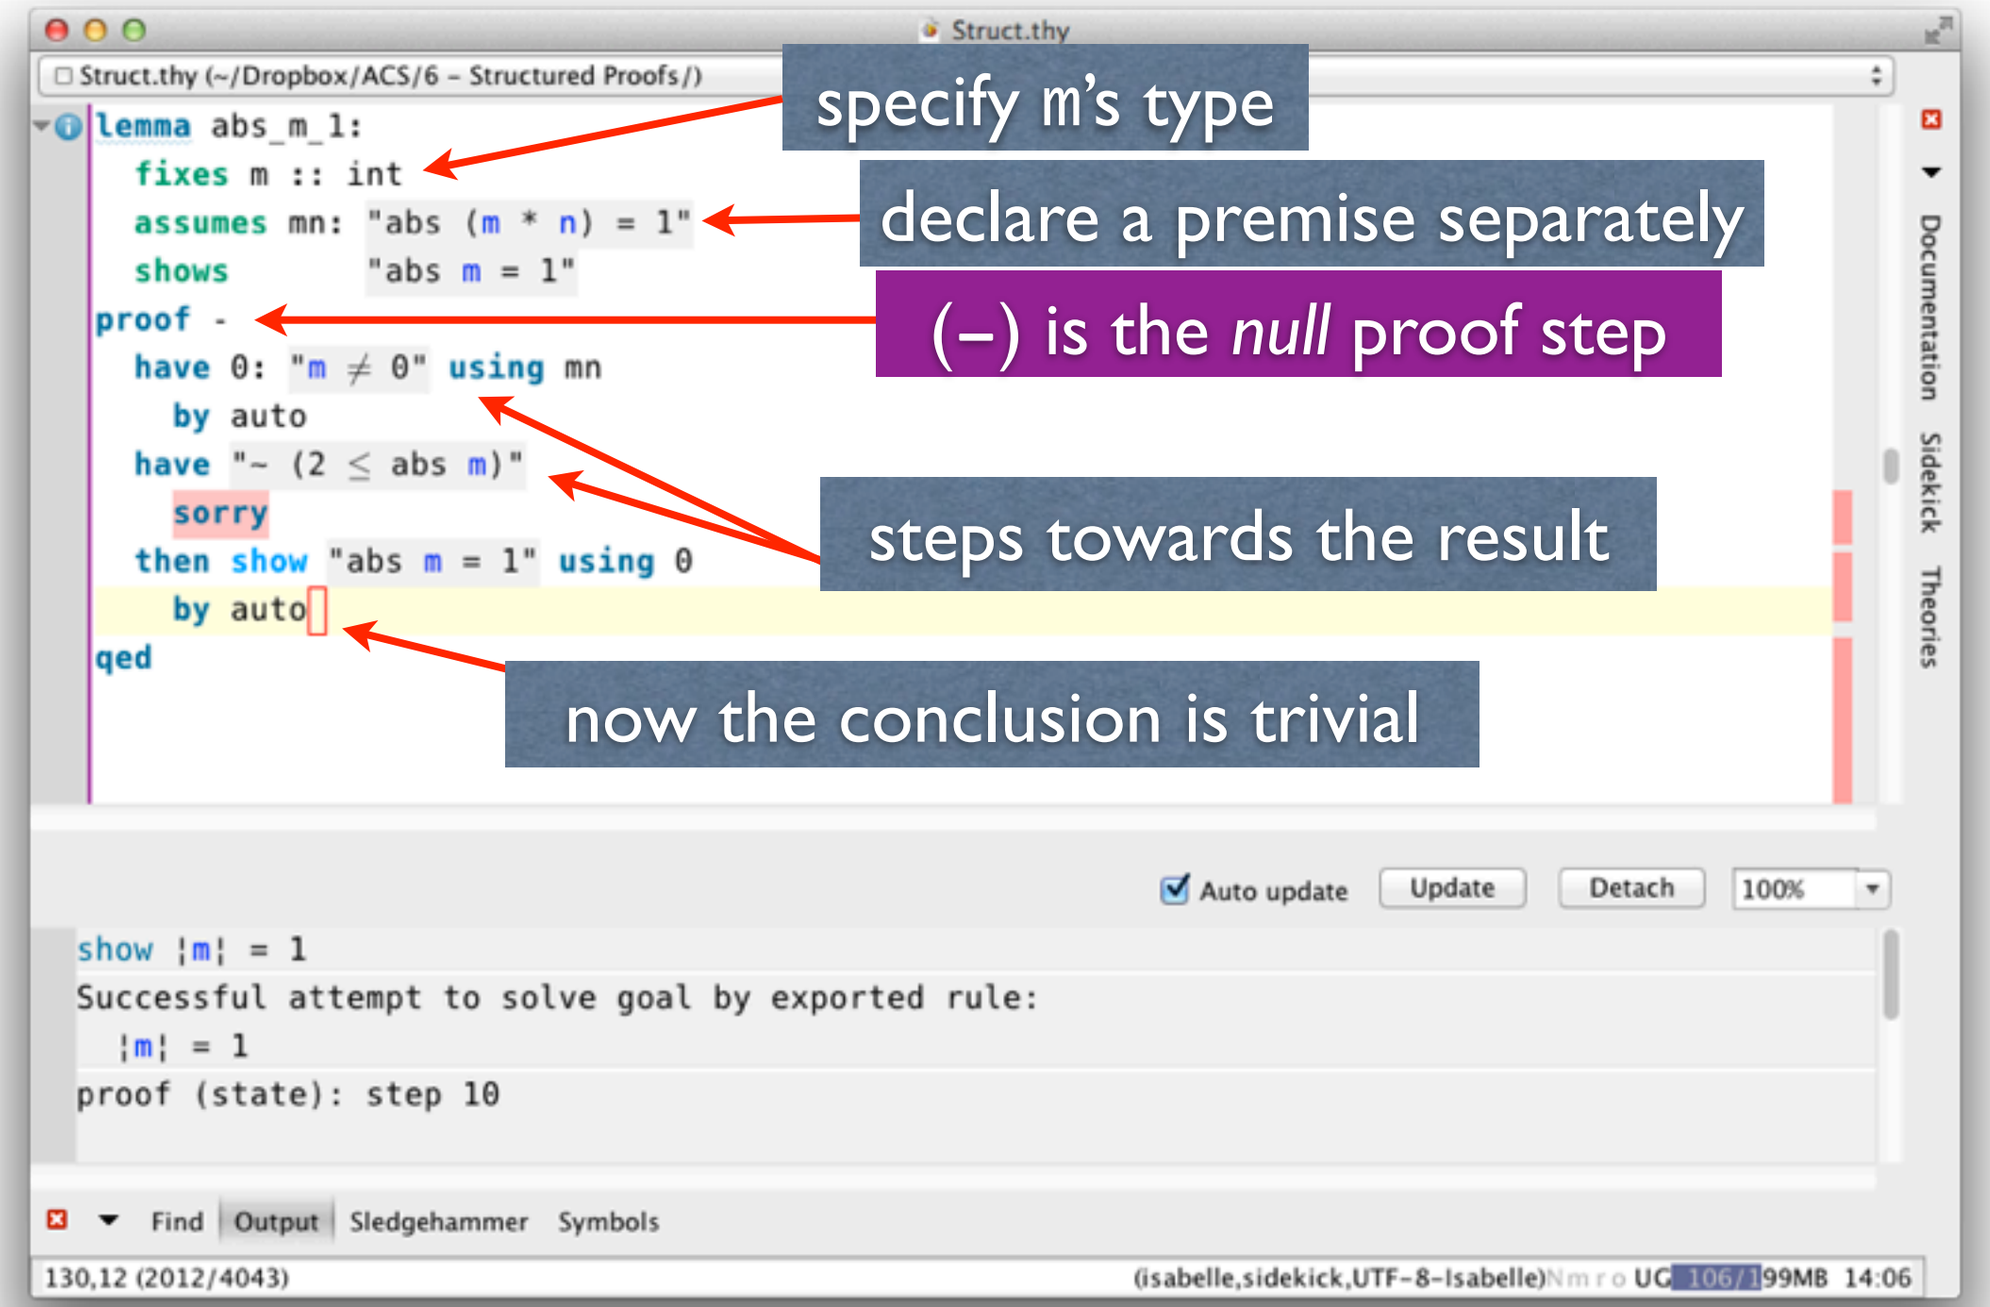

This is an example of an obvious fact is proof is not obvious. Clearly m≠0, since otherwise m*n=0. If we can also show that |m|≥2 is impossible, then the only remaining possibility is |m|=1.

In this example, `auto` can do nothing. No proof steps are obvious from the problem's syntax. So the Isar proof begins with "`-`", the null proof. This step does nothing but insert any "pending facts" from a previous step (here, there aren't any) into the proof state. It is quite common to begin with "`proof -`".

Starting a Nested Proof

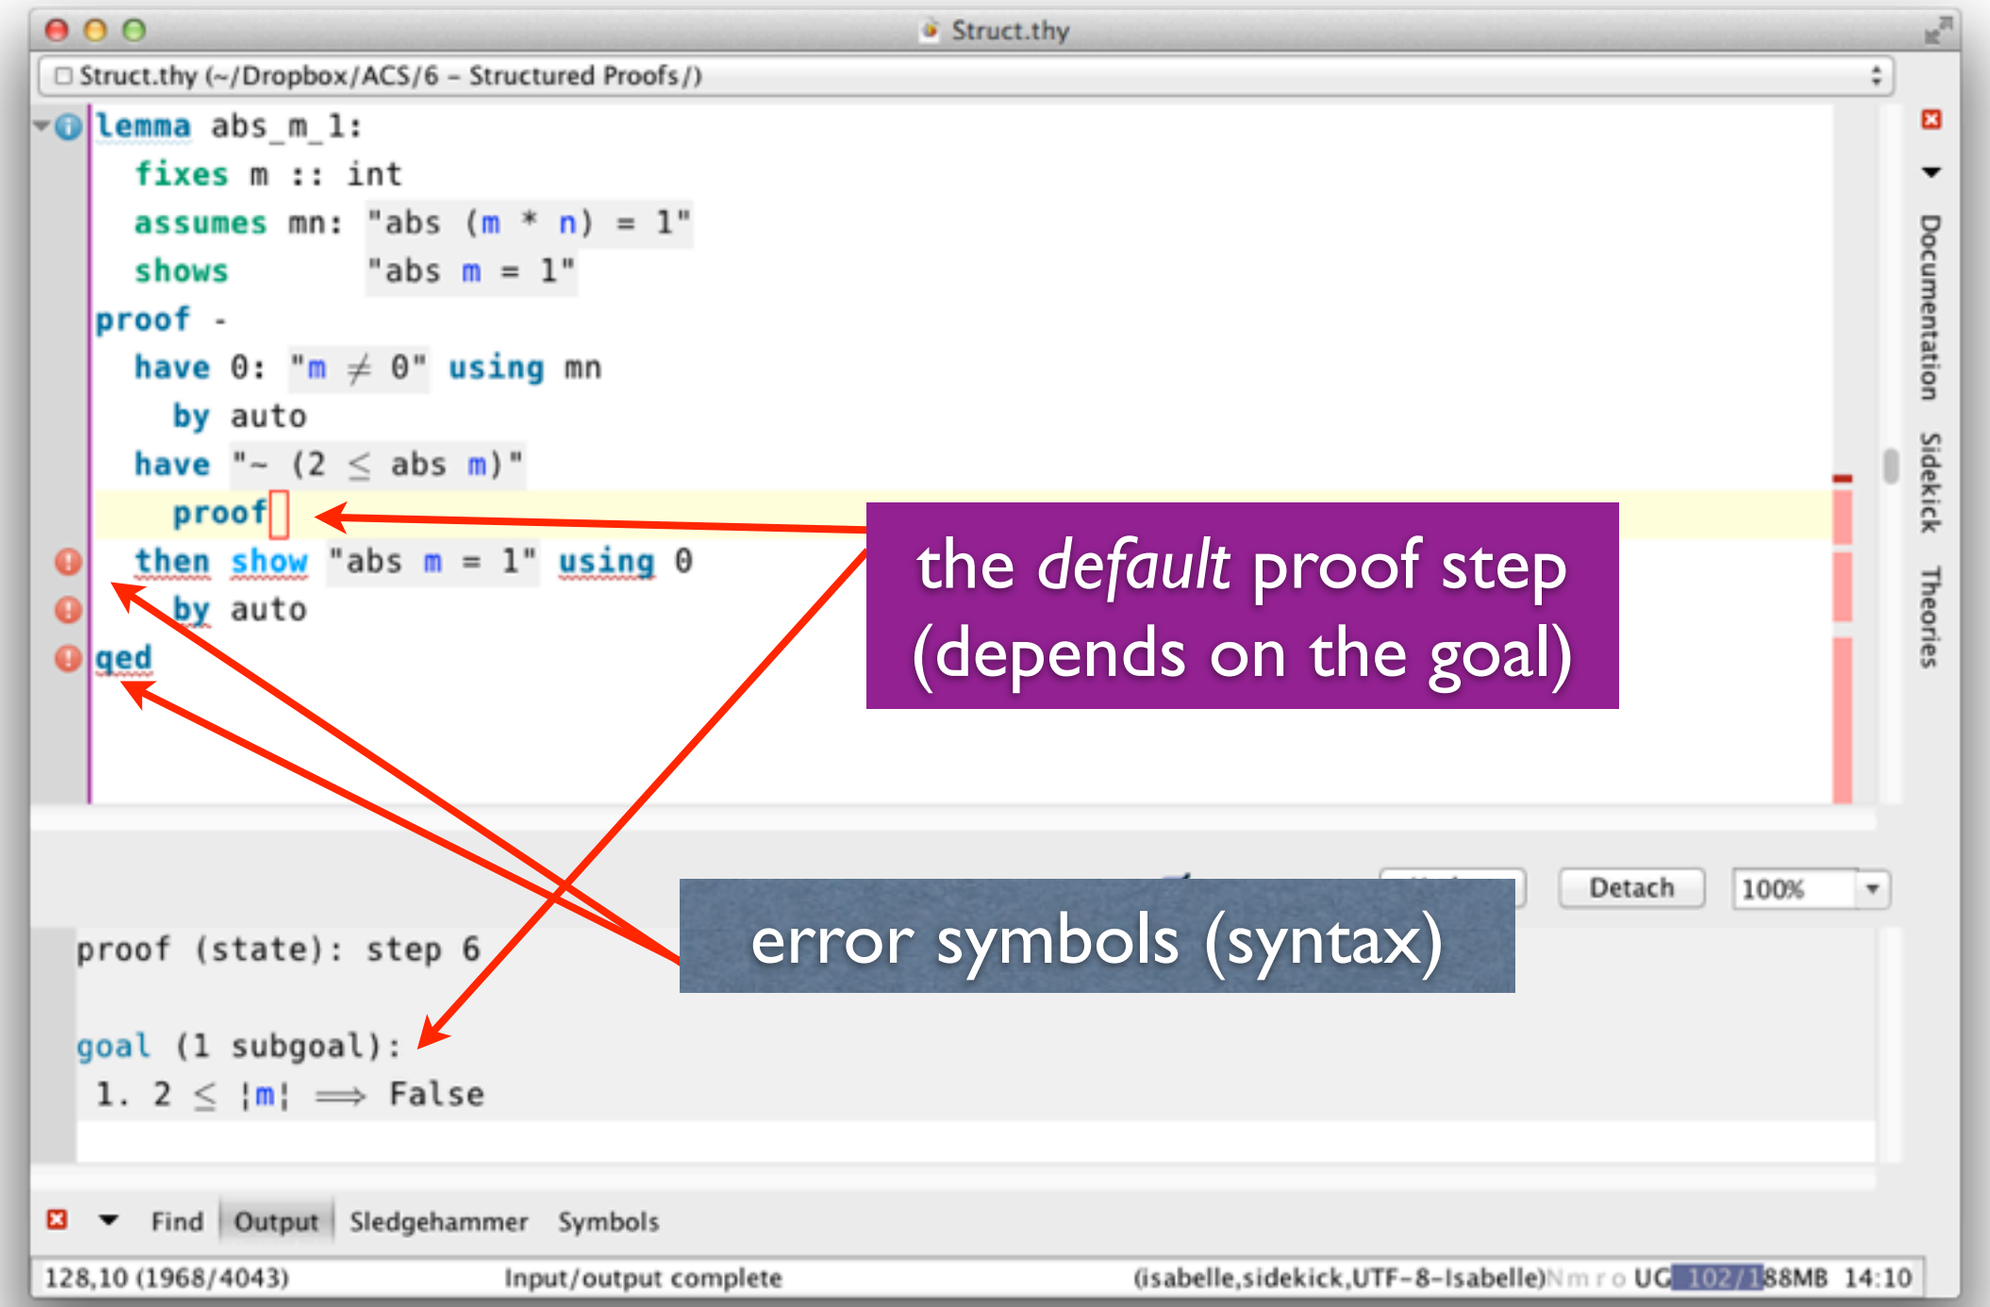

To begin with "`proof`" (not to be confused with "`proof -`") applies a default proof method. In theory, this method should be appropriate for the problem, but in practice, it is often unhelpful. The default method is determined by elementary syntactic criteria. For example, the formula "¬ (2 ≤ abs m)" begins with a negation sign, so the default method applies the corresponding logical inference: it reduces the problem to proving False under the assumption 2 ≤ abs m.

A Nested Proof Skeleton

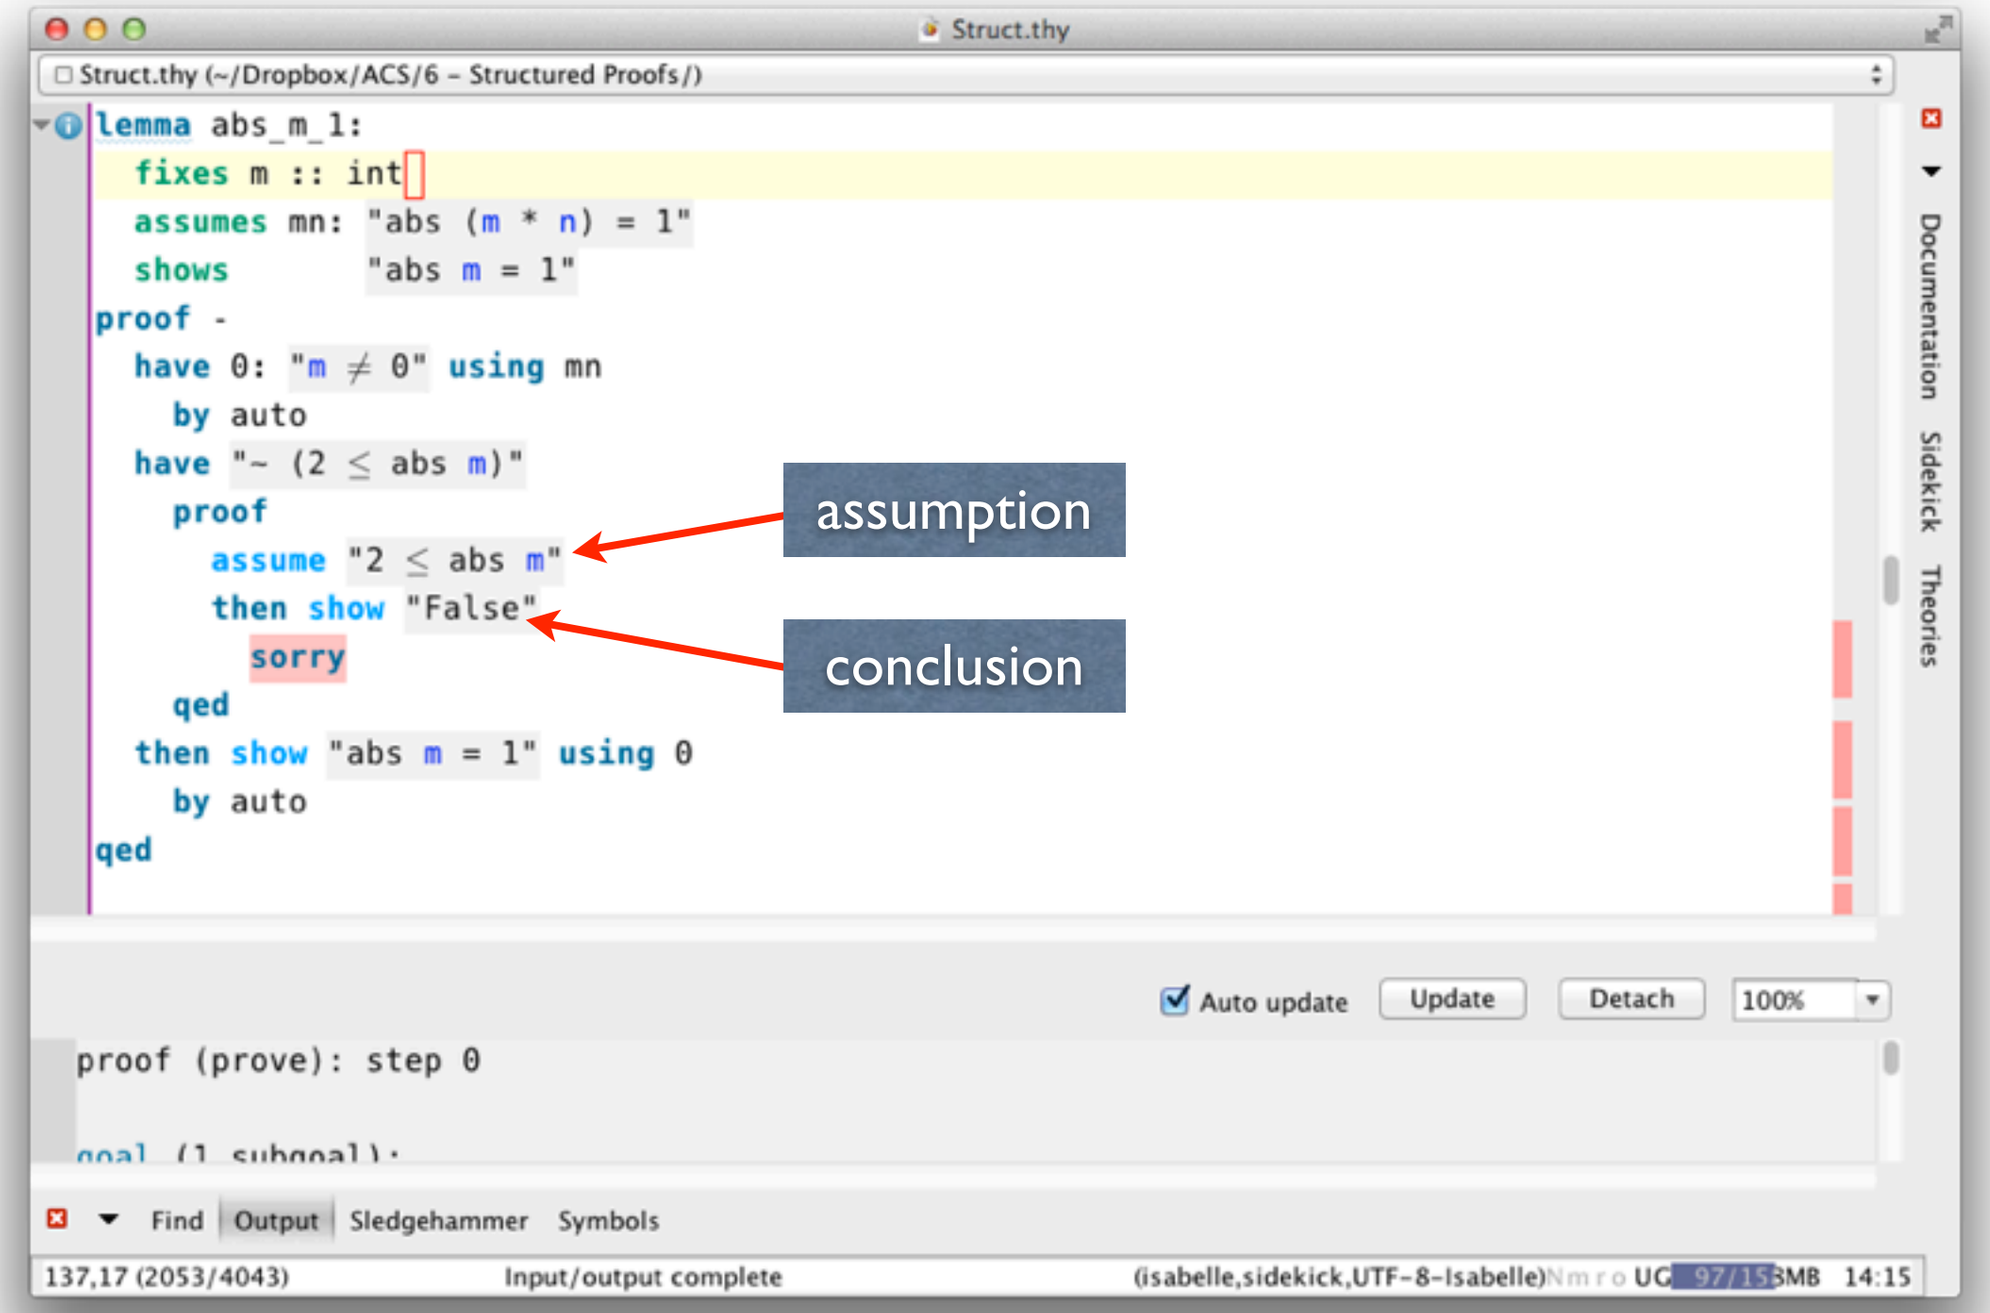

Proofs can be nested to any depth. The assumptions and conclusions of each nested proof are independent of one another. The usual scoping rules apply, and in particular the facts `mn` and `0` are visible within this inner scope.

A Complete Proof

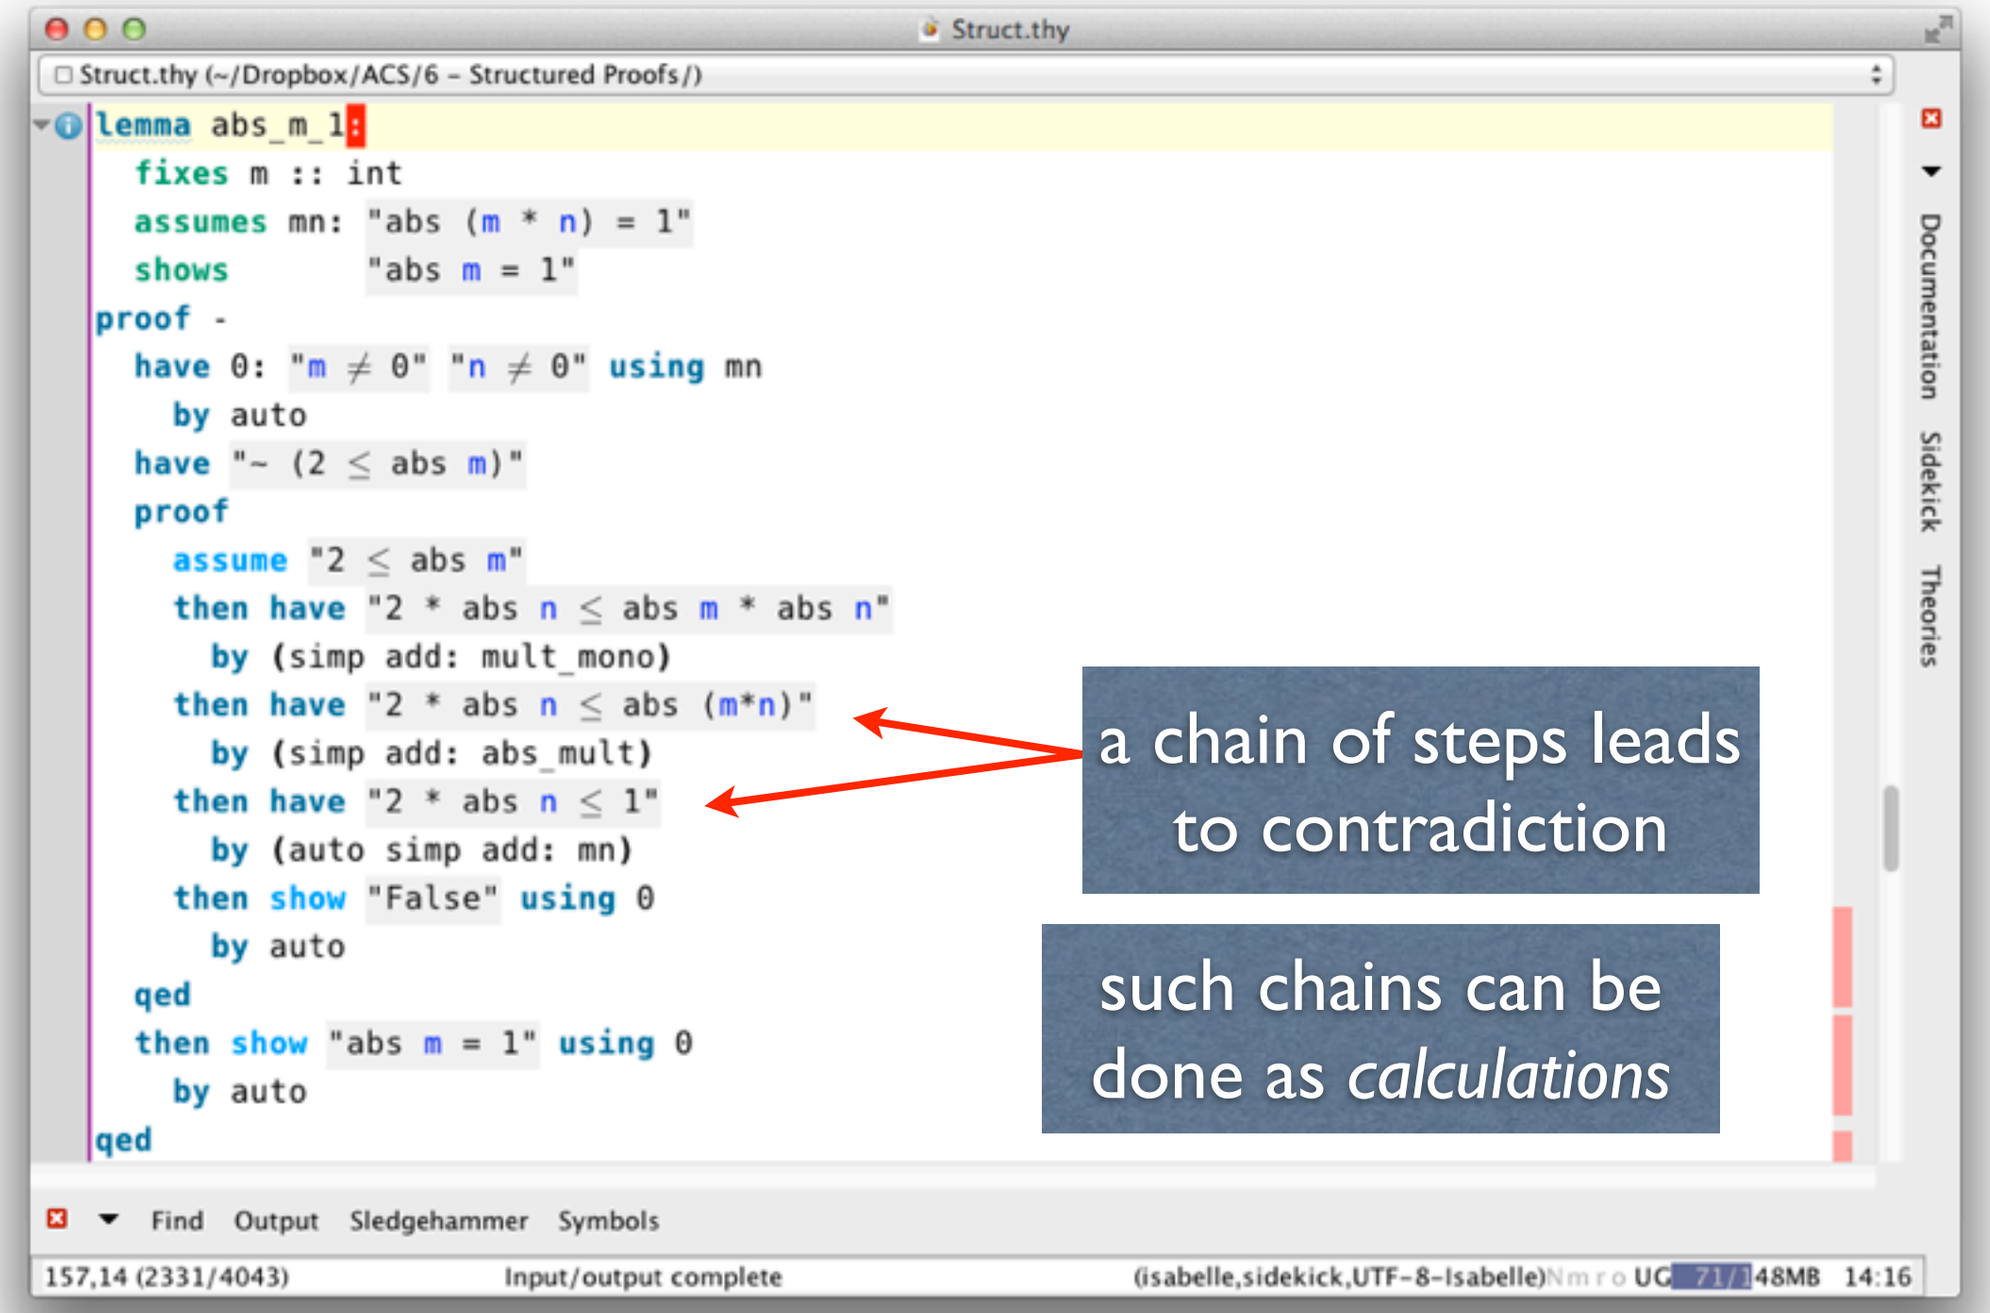

This example is typical of a structured proof. From the assumption, 2 ≤ abs m, we deduce a chain of consequences that become absurd. We connect one step to the next using "`hence`", except that we must introduce the conclusion using "`thus`".

Note that we have beefed up the fact "0" from simply m≠0 to include as well n≠0, which we need to obtain a contradiction from 2 × abs n ≤ 1. In fact, "0" here denotes a list of facts.

Calculational Proofs

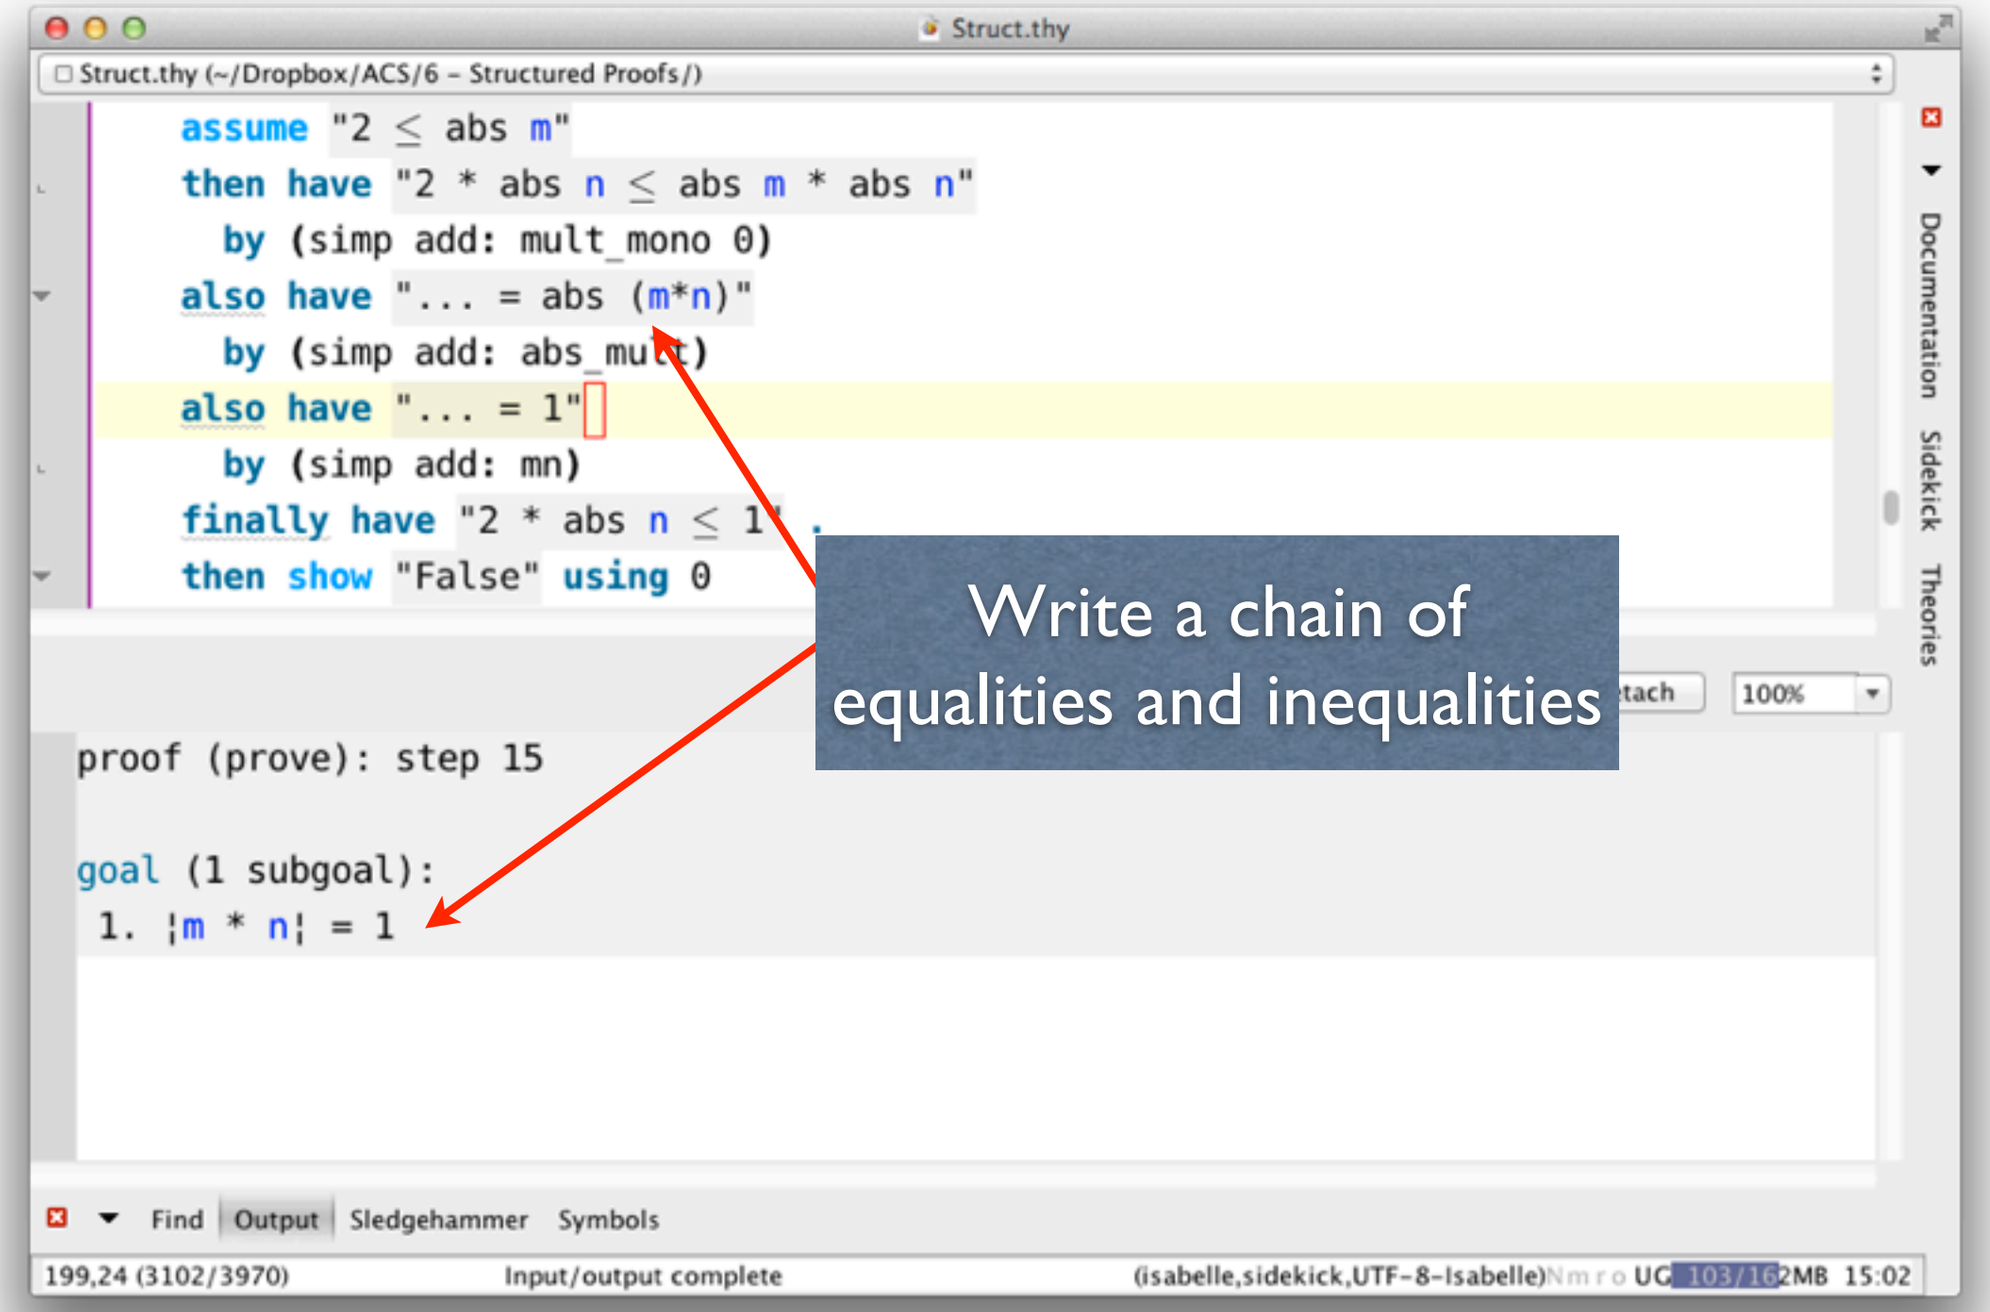

The chain of reasoning in the previous proof holds by transitivity, and in normal mathematical discourse would be written as a chain of inequalities and inequalities. Isar supports this notation.

The Next Step

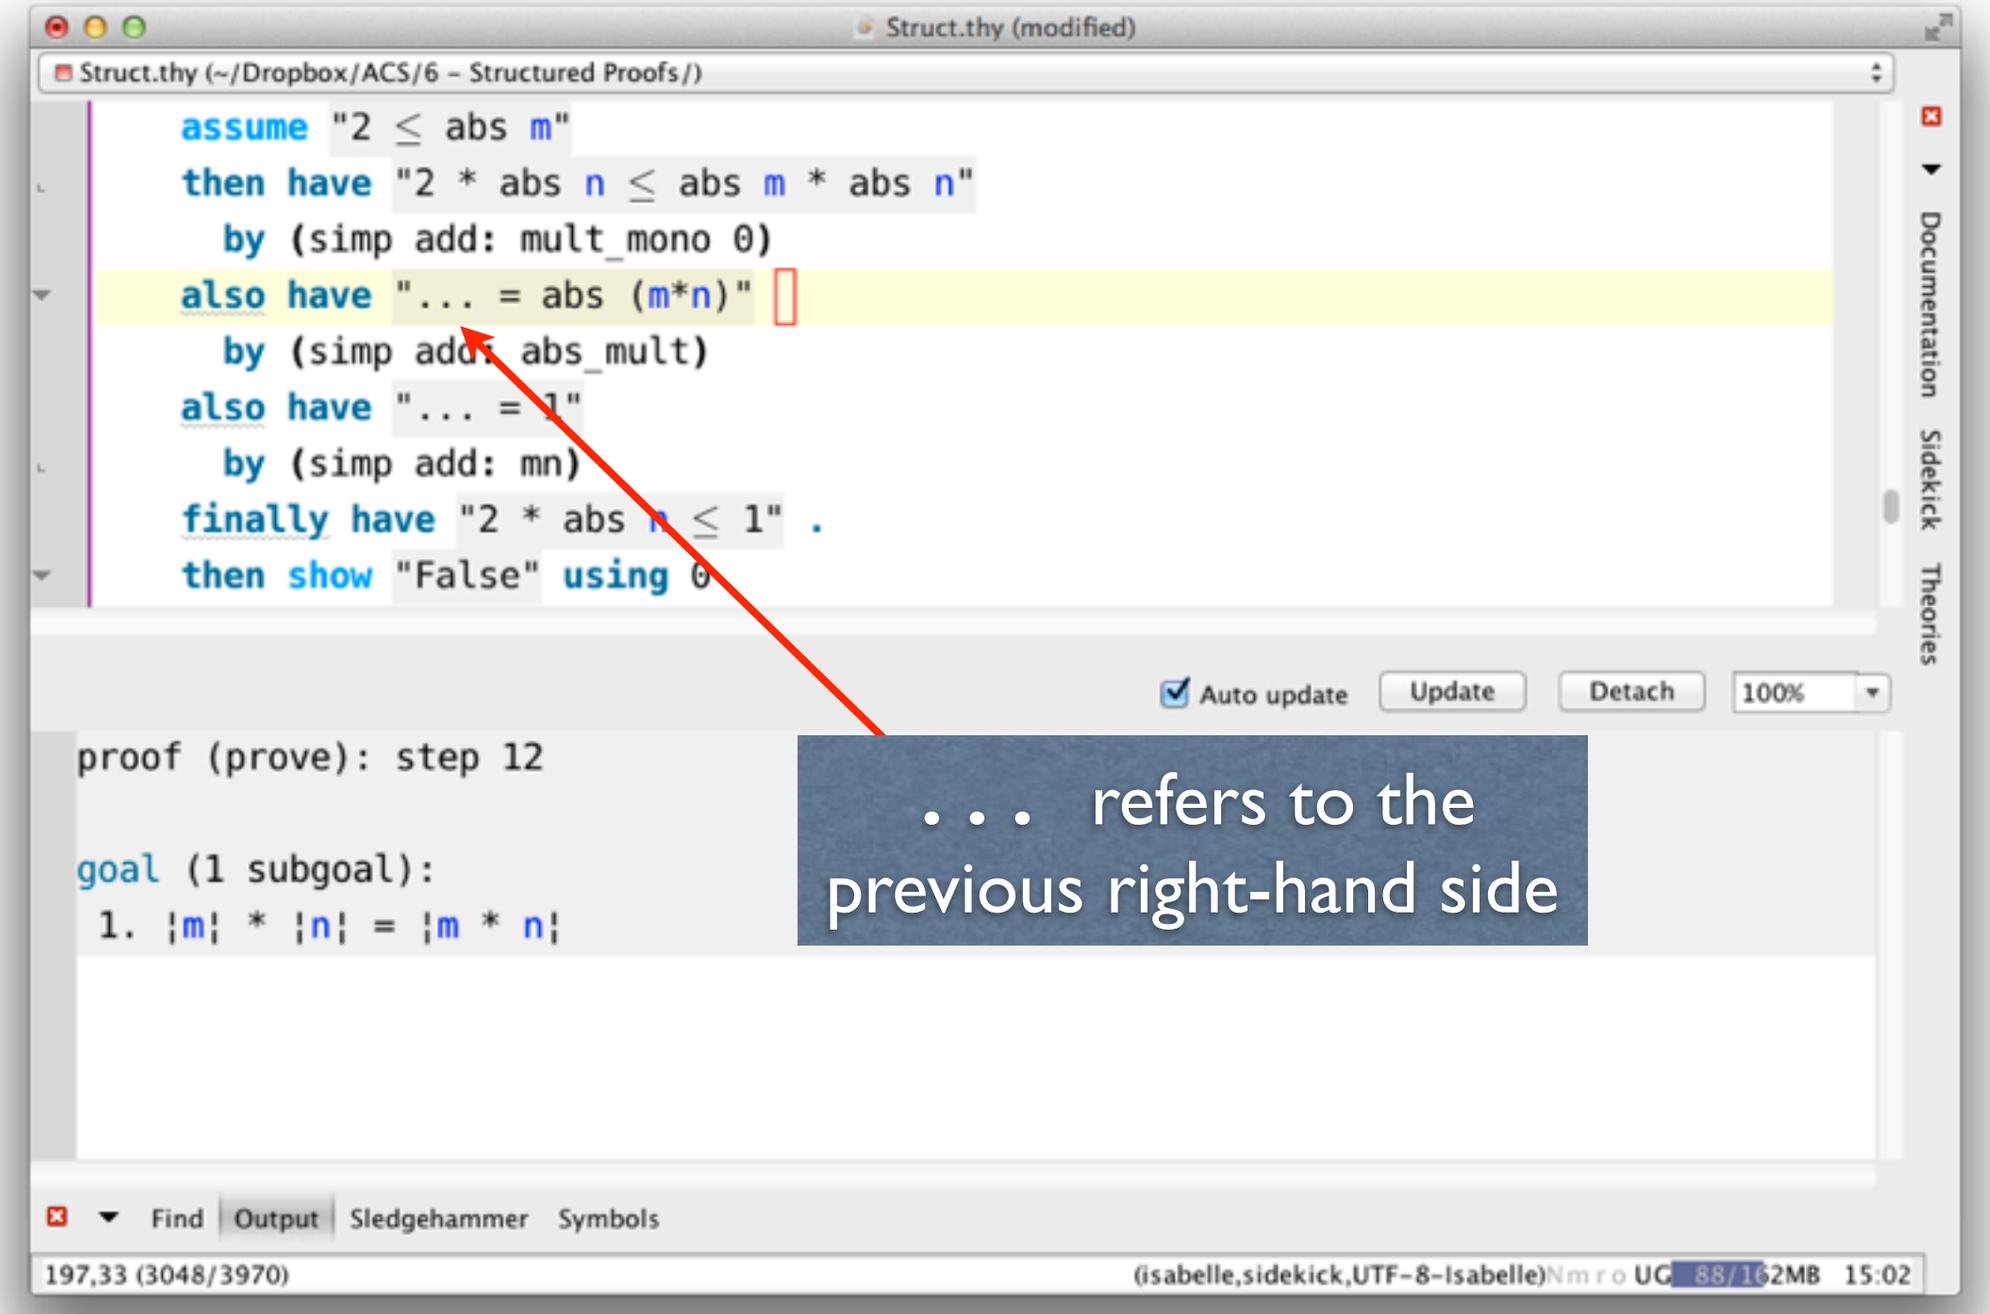

The Internal Calculation

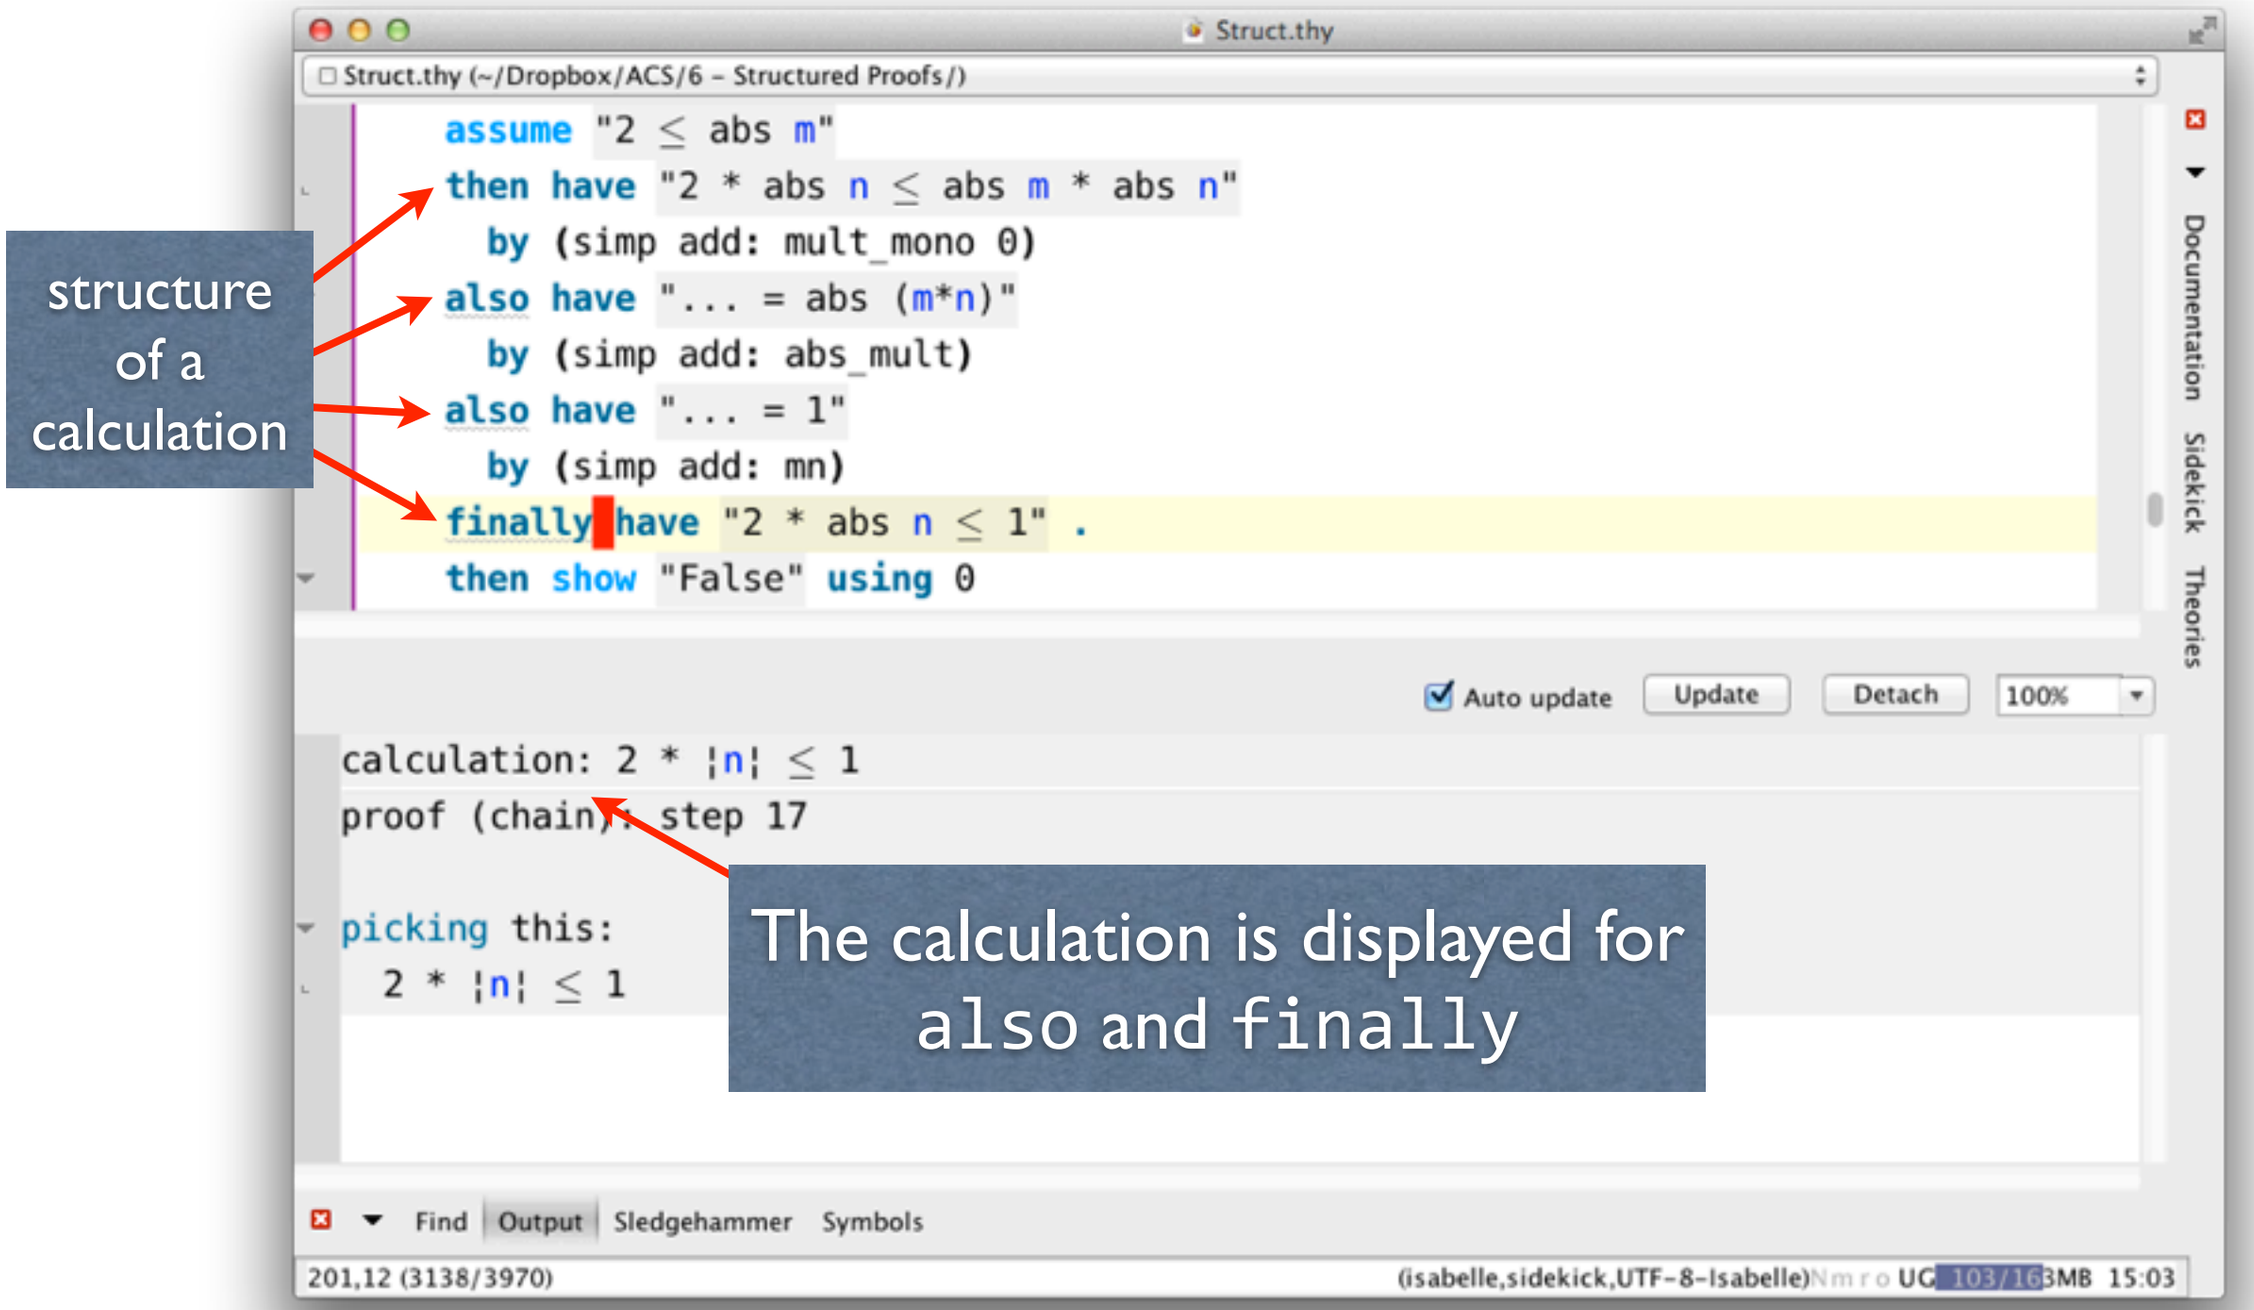

Use "`also`" to attach a new link to the chain, extending the calculation. Use "`finally`" to refer to the calculation itself. It is usual for the proof script merely to repeat explicitly what this calculation should be, as shown above. If this is done, the proof is trivial and is written in Isar as a single dot (`.`).

We could instead avoid that repetition and reach the contradiction directly as follows:

```
  also have "... = 1"
   by (simp add: mn)
finally show "False" using 0
   by auto
```

Internally, this proof is identical to the previous one. It merely differs in appearance, not bothering to note that 2 × abs n ≤ 1 has been derived.

Ending the Calculation

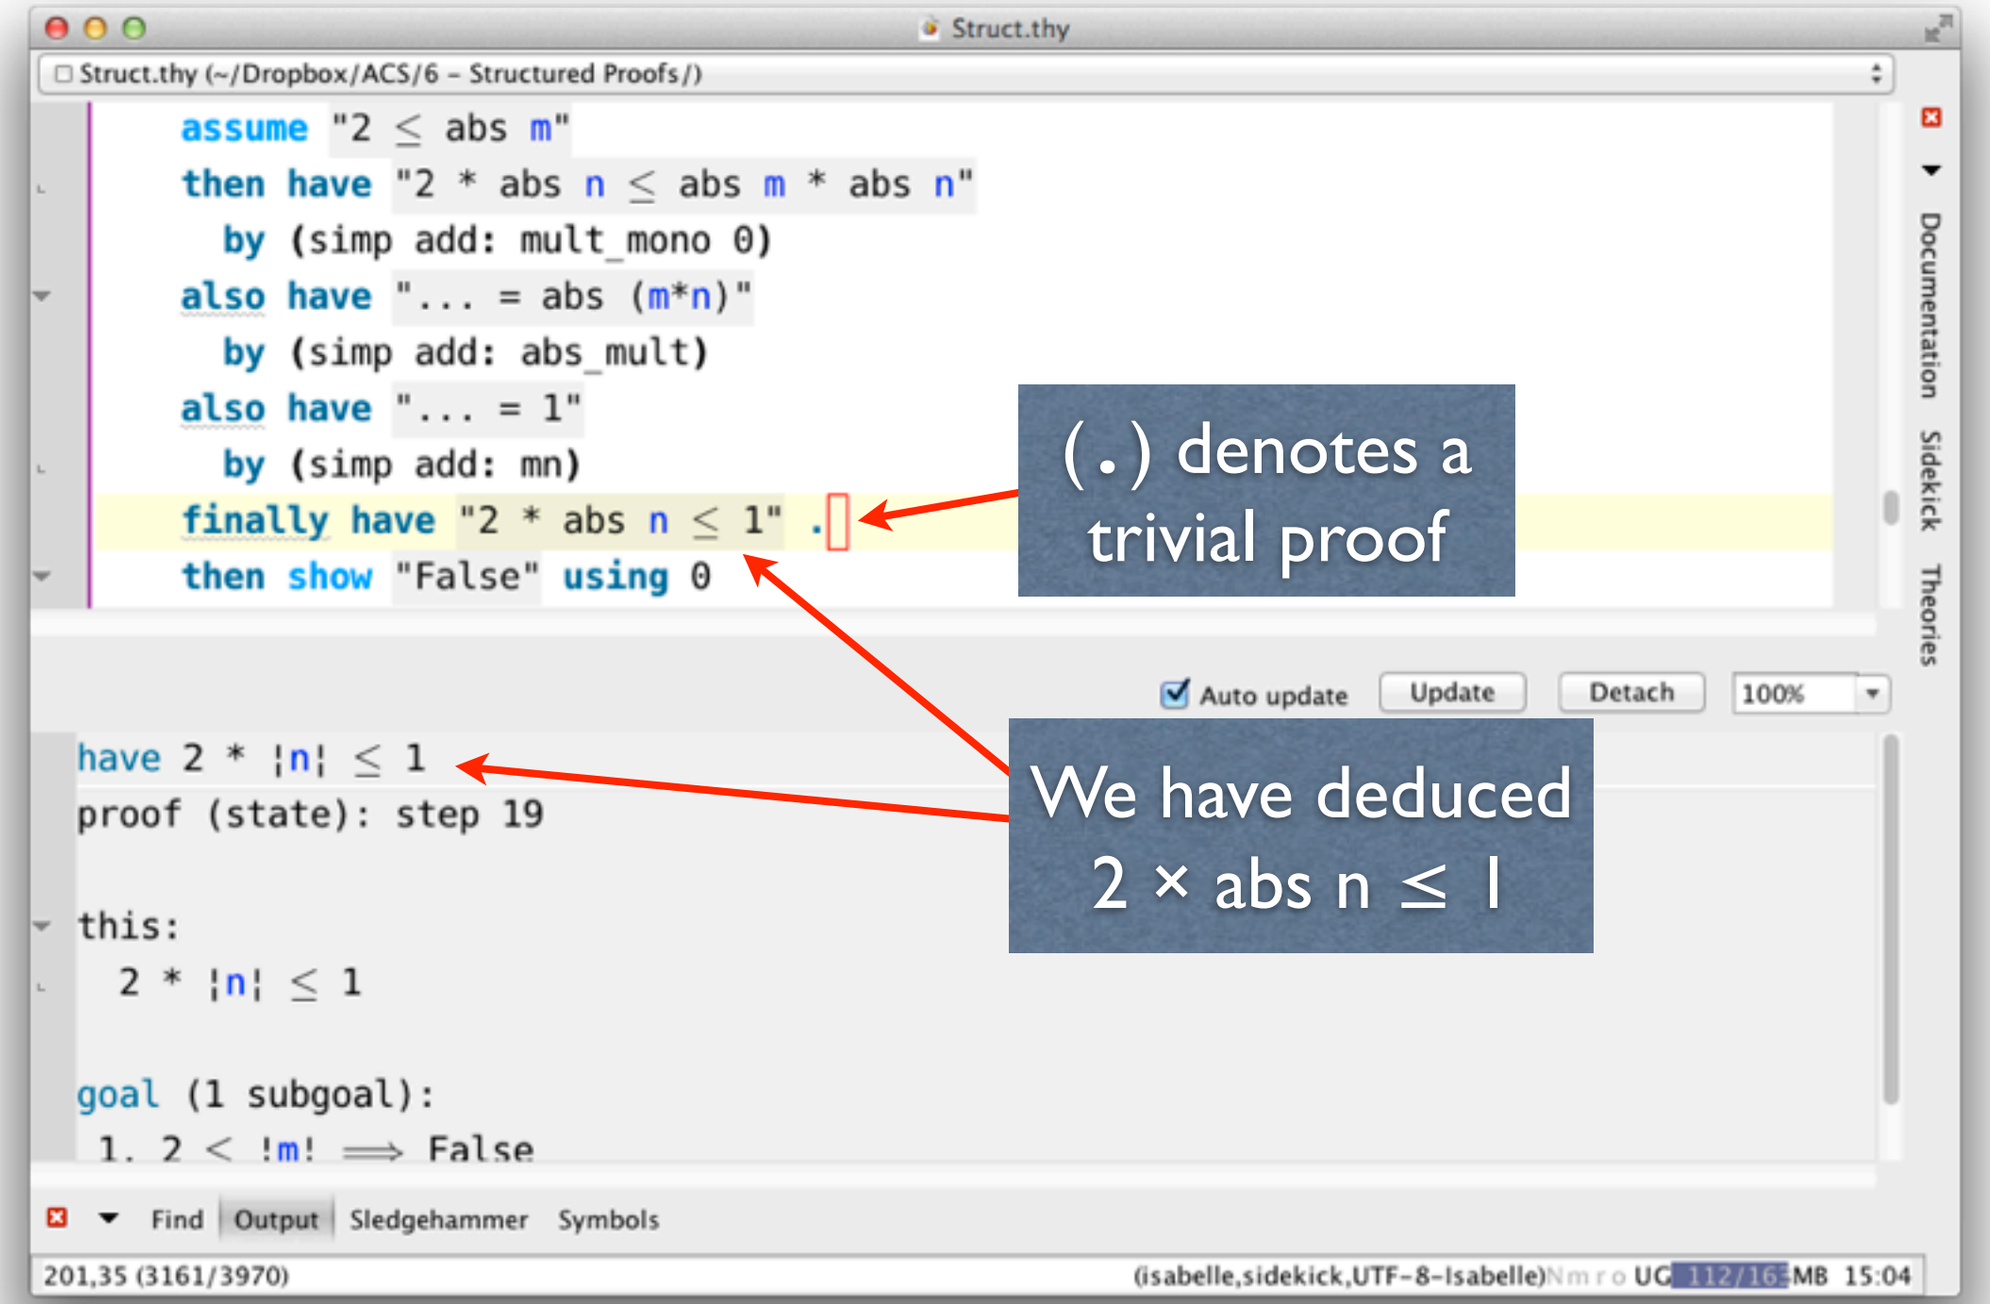

The last line of the proof is unnecessary, and merely restates what was proved by the calculation. That's why its proof is trivial. We could have concluded this proof fragment as follows, feeding the calculation straight into the desired conclusion:

```
finally show "False" using 0
      by auto
```

Structure of a Calculation

- The first line begins with `have` *or* `hence`
- Subsequent lines begin with

  `also have “... = “`

- *Any* transitive relation may be used. New ones may be declared.
- The concluding line begins with

  `finally have` *or* `show`

- It repeats the calculation and terminates with `(.)`

Interactive Formal Verification
7: Sets

Lawrence C Paulson
Computer Laboratory
University of Cambridge

Set Notation in Isabelle

- Set notation is crucial to mathematical discourse.
  - Operators such as union, intersection, powerset and image naturally express many complex constructions.
  - Functions, relations, and concepts such as transitive closure are available.
- A set in higher-order logic is similar to a *boolean-valued map*: in other words, to a logical predicate.
- The elements of a set must all have the *same type!*

See the *Tutorial*, section 6.1 Sets.

Set Theory Primitives

- The type `α set`, which is similar to `α⇒bool`
- The membership relation: ∈
- The subset relation: ⊆
  - Reflexive, anti-symmetric, transitive
- The empty set: { }
- The universal set: `UNIV`

Basic Set Theory Operations

$$
\begin{aligned}
e \in \{x.\, P(x)\} &\iff P(e) \\
e \in \{x \in A.\, P(x)\} &\iff e \in A \wedge P(e) \\
e \in -A &\iff e \notin A \\
e \in A \cup B &\iff e \in A \vee e \in B \\
e \in A \cap B &\iff e \in A \wedge e \in B \\
e \in \mathtt{Pow}(A) &\iff e \subseteq A
\end{aligned}
$$

Please note that we do not write {xIP(x)}. Isabelle would interpret the I as expressing disjunction and the expression as denoting the singleton set containing the element xIP(x)!

The logical equivalences shown above are effectively the definitions of the primitives shown, and any occurrences of the left-hand side formula will be replaced by the right-hand side by Isabelle's simplifier.

Big Union and Intersection

$$e \in \left(\bigcup x.\, B(x)\right) \iff \exists x.\, e \in B(x)$$

$$e \in \left(\bigcup x \in A.\, B(x)\right) \iff \exists x \in A.\, e \in B(x)$$

$$e \in \bigcup A \iff \exists x \in A.\, e \in x$$

And the analogous forms of intersections...

Once again, the logical equivalences are essentially definitions.

The third form of union is seldom seen.

A Simple Set Theory Proof

Special symbols can be inserted using the Symbols panel. ASCII can simply be typed; auto-completions for symbols will be offerered.

The main point of this example is that many such proofs are trivial, using `auto` or other automatic proof methods.

Also: look for icons in the left-hand "gutter", since they indicate errors, warnings or information.

Functions

$$e \in (f\text{`}A) \iff \exists x \in A.\, e = f(x)$$

$$e \in (f^{-1}\text{`}A) \iff f(e) \in A$$

$$f(x{:=}y) = (\lambda z.\ \texttt{if}\ z = x\ \texttt{then}\ y\ \texttt{else}\ f(z))$$

- Also `inj`, `surj`, `bij`, `inv`, etc. (injective,...)
- Don’t *re-invent* image and inverse image!!

Inverse image is also known as pre-image. Using the actual image primitives gives access to the many theorems proved about them.

Finite Set Notation

$$\{a_1, \ldots, a_n\} = \text{insert } a_1 \ (\ldots \ (\text{insert } a_n \ \{\})\ldots)$$

where

$$x \in \text{insert } a \ B \iff x{=}a \ \vee \ x \in B$$

Finite sets can be written explicitly, enumerating their elements in the obvious way.

Finite Sets

A finite set is defined *inductively*
in terms of `{}` and `insert`

$$\texttt{finite}(A \cup B) = (\texttt{finite}\, A \wedge \texttt{finite}\, B)$$

$$\texttt{finite}\, A \Longrightarrow \texttt{card}(\texttt{Pow}\, A) = 2^{\texttt{card}\, A}$$

Defining functions over finite sets is tricky, because your definition has to make sense regardless of the order of the elements and regardless of whether they are repeated or not, because the sets {x,y}, {y,x} and {x,y,x} are all equal. The notion of cardinality is built-in.

Intervals, Sums and Products

```
{..<u} == {x. x < u}
 {..u} == {x. x ≤ u}
{l<..} == {x. l<x}
 {l..} == {x. l≤x}
{l<..<u} == {l<..} ∩ {..<u}
{l..<u} == {l..} ∩ {..<u}
```

`setsum f A` and `setprod f A`

$\sum i \in I.\, f$ and $\prod i \in I.\, f$

Isabelle provides syntax for bounded and unbounded intervals. These are polymorphic: they are defined over all types that admit an ordering, and in particular they are applicable to intervals over the natural numbers, integers, rationals or reals.

Sums and products of functions over sets can also be written.

A Harder Proof Involving Sets

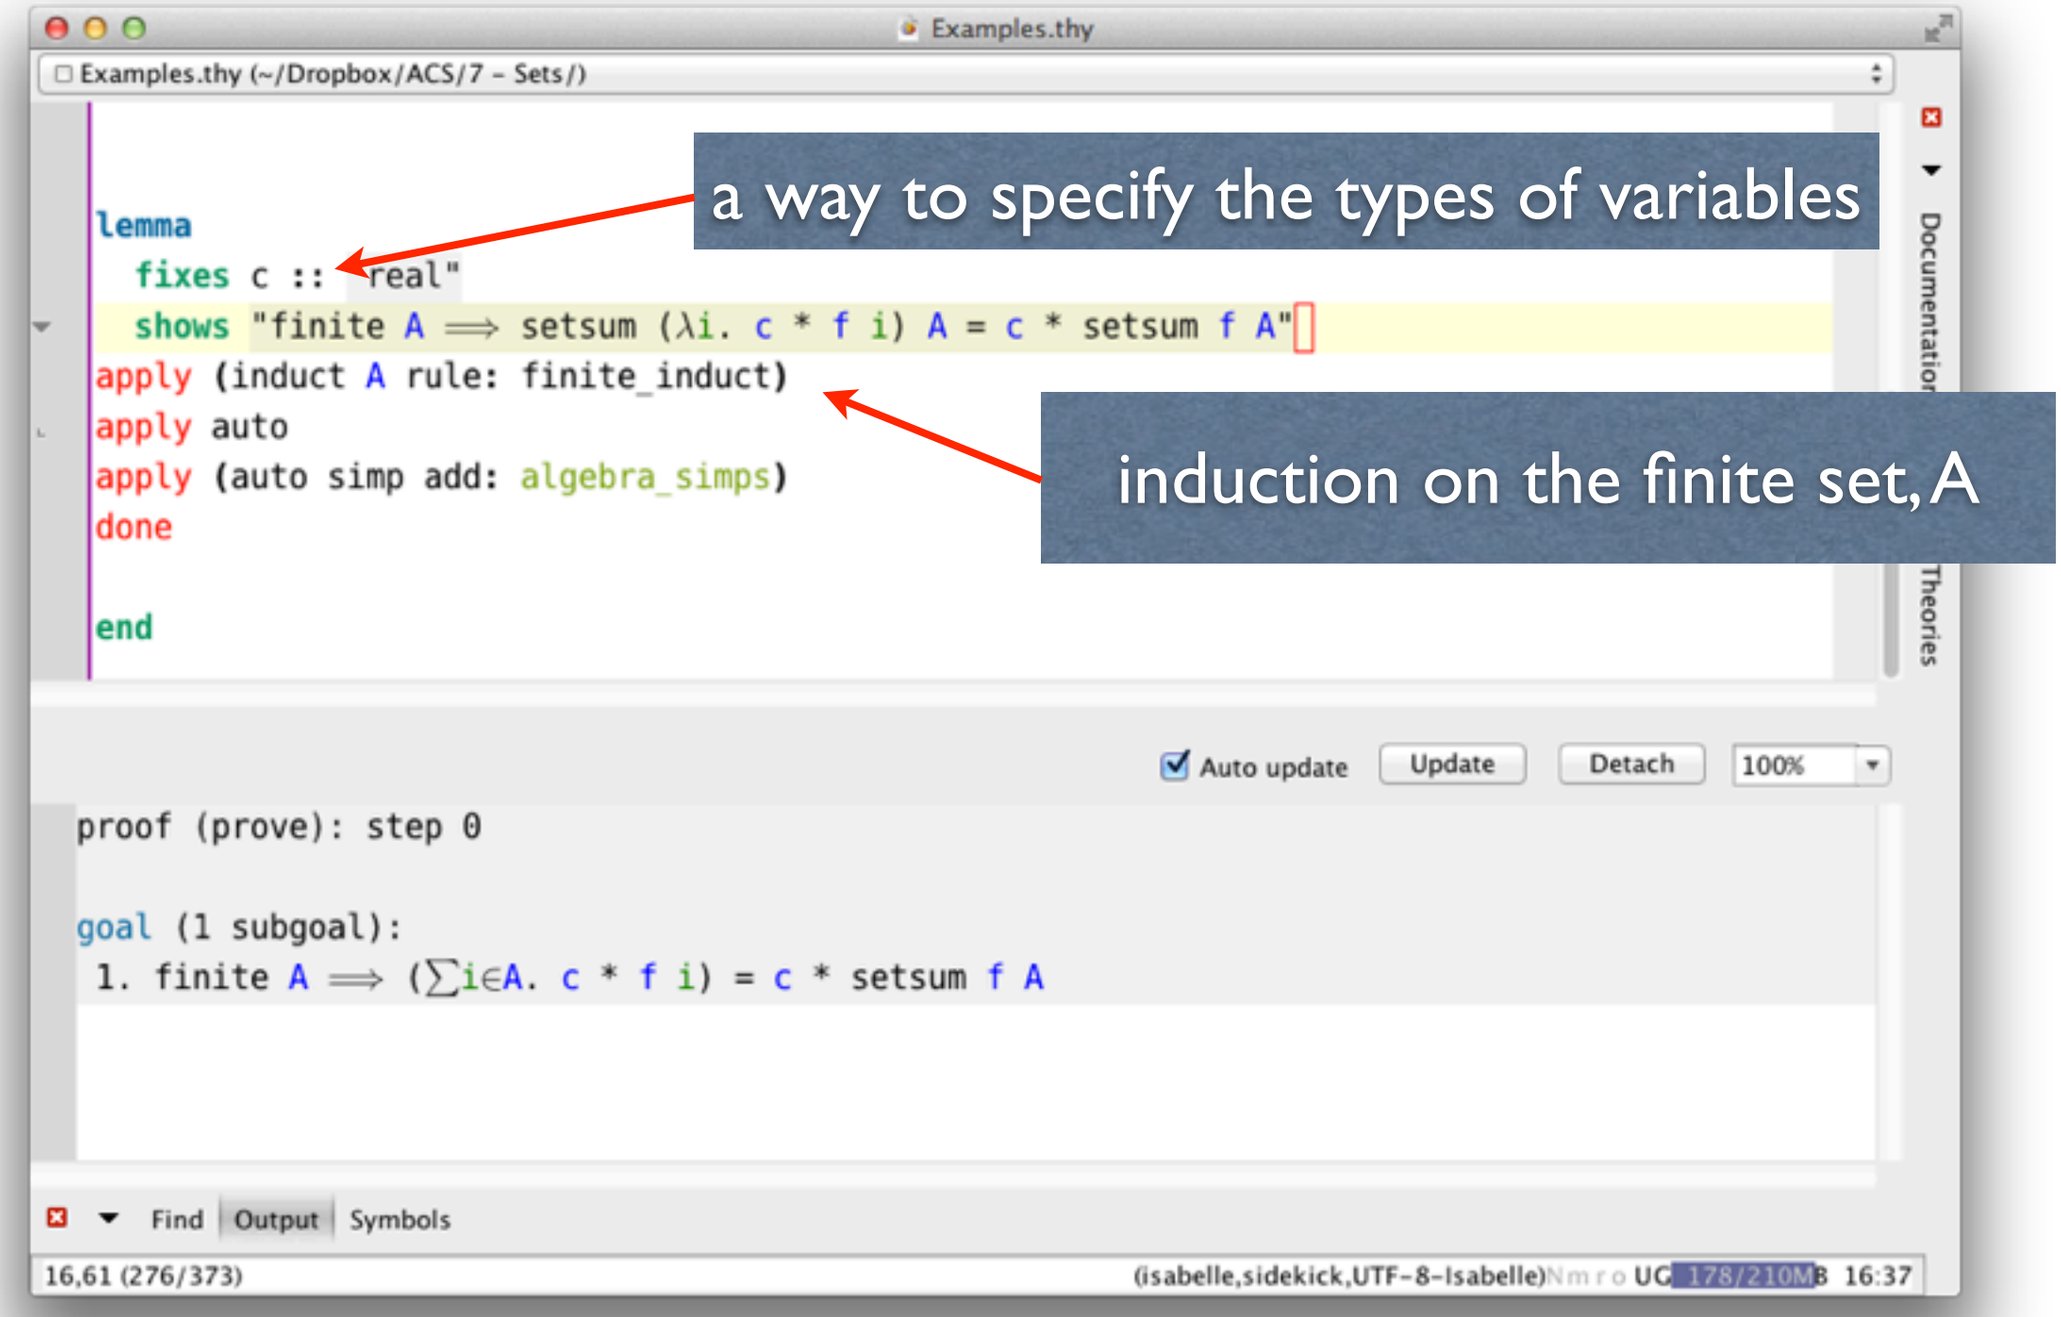

This example needs a type constraint because arithmetic concepts such as sum and product are heavily overloaded. If you use `fixes`, then you must also use `shows`!

Isabelle's type classes allow this theorem to be proved in an overloaded form, but for simplicity here we restrict ourselves to type `real`.

Outcome of the Induction

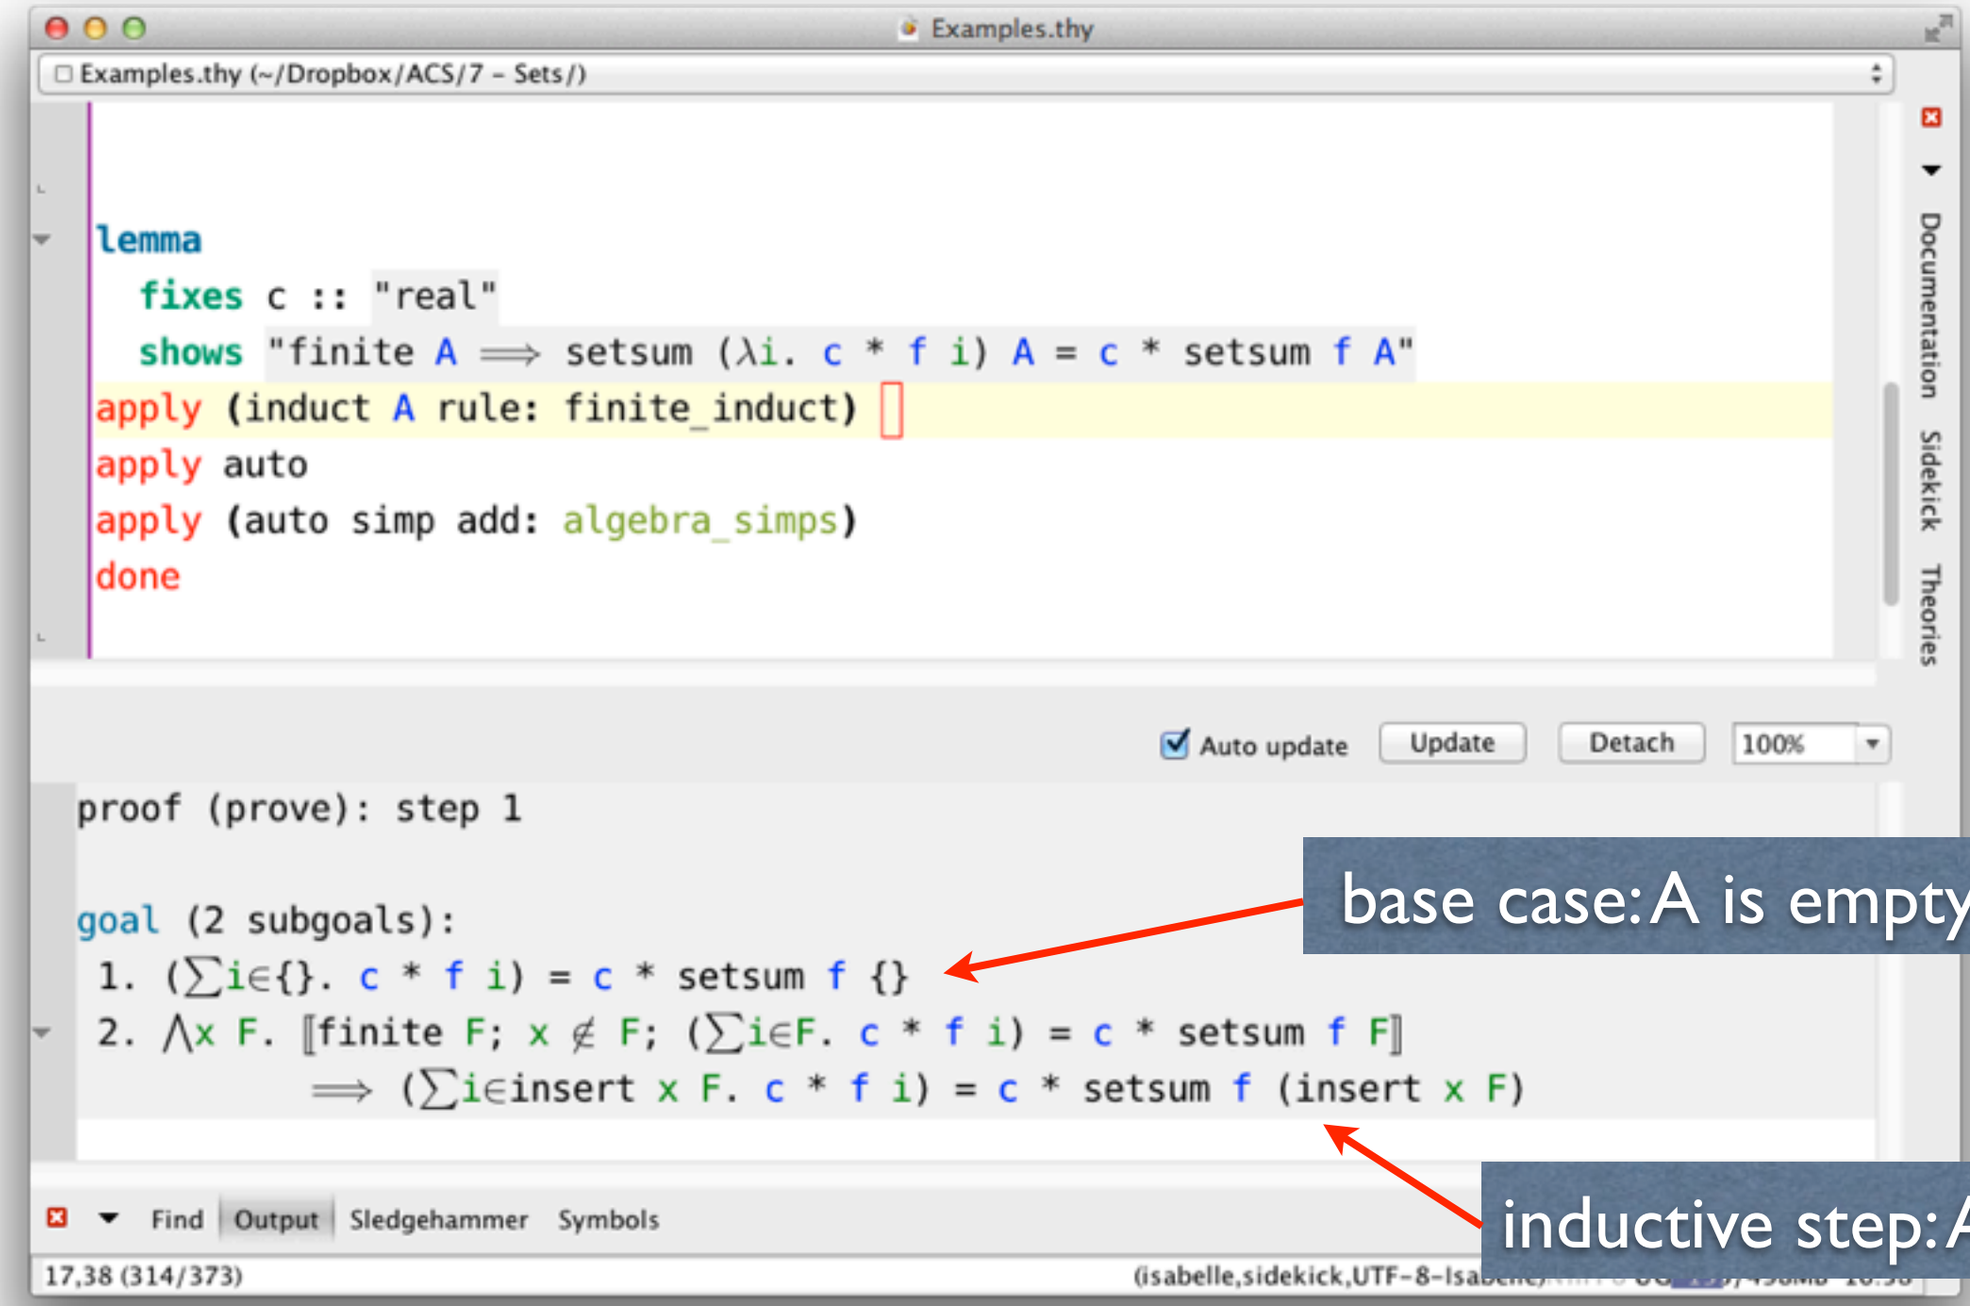

The base case is trivial, because both sides of the equality clearly equal zero. In the induction step, the induction hypothesis (which concerns the set `F`) will be applicable, because

```
setsum f (insert a F) = f a + setsum f F
```

Note that Isabelle uses a fancy notation for summations, but only if the body of the summation is nontrivial.

Almost There!

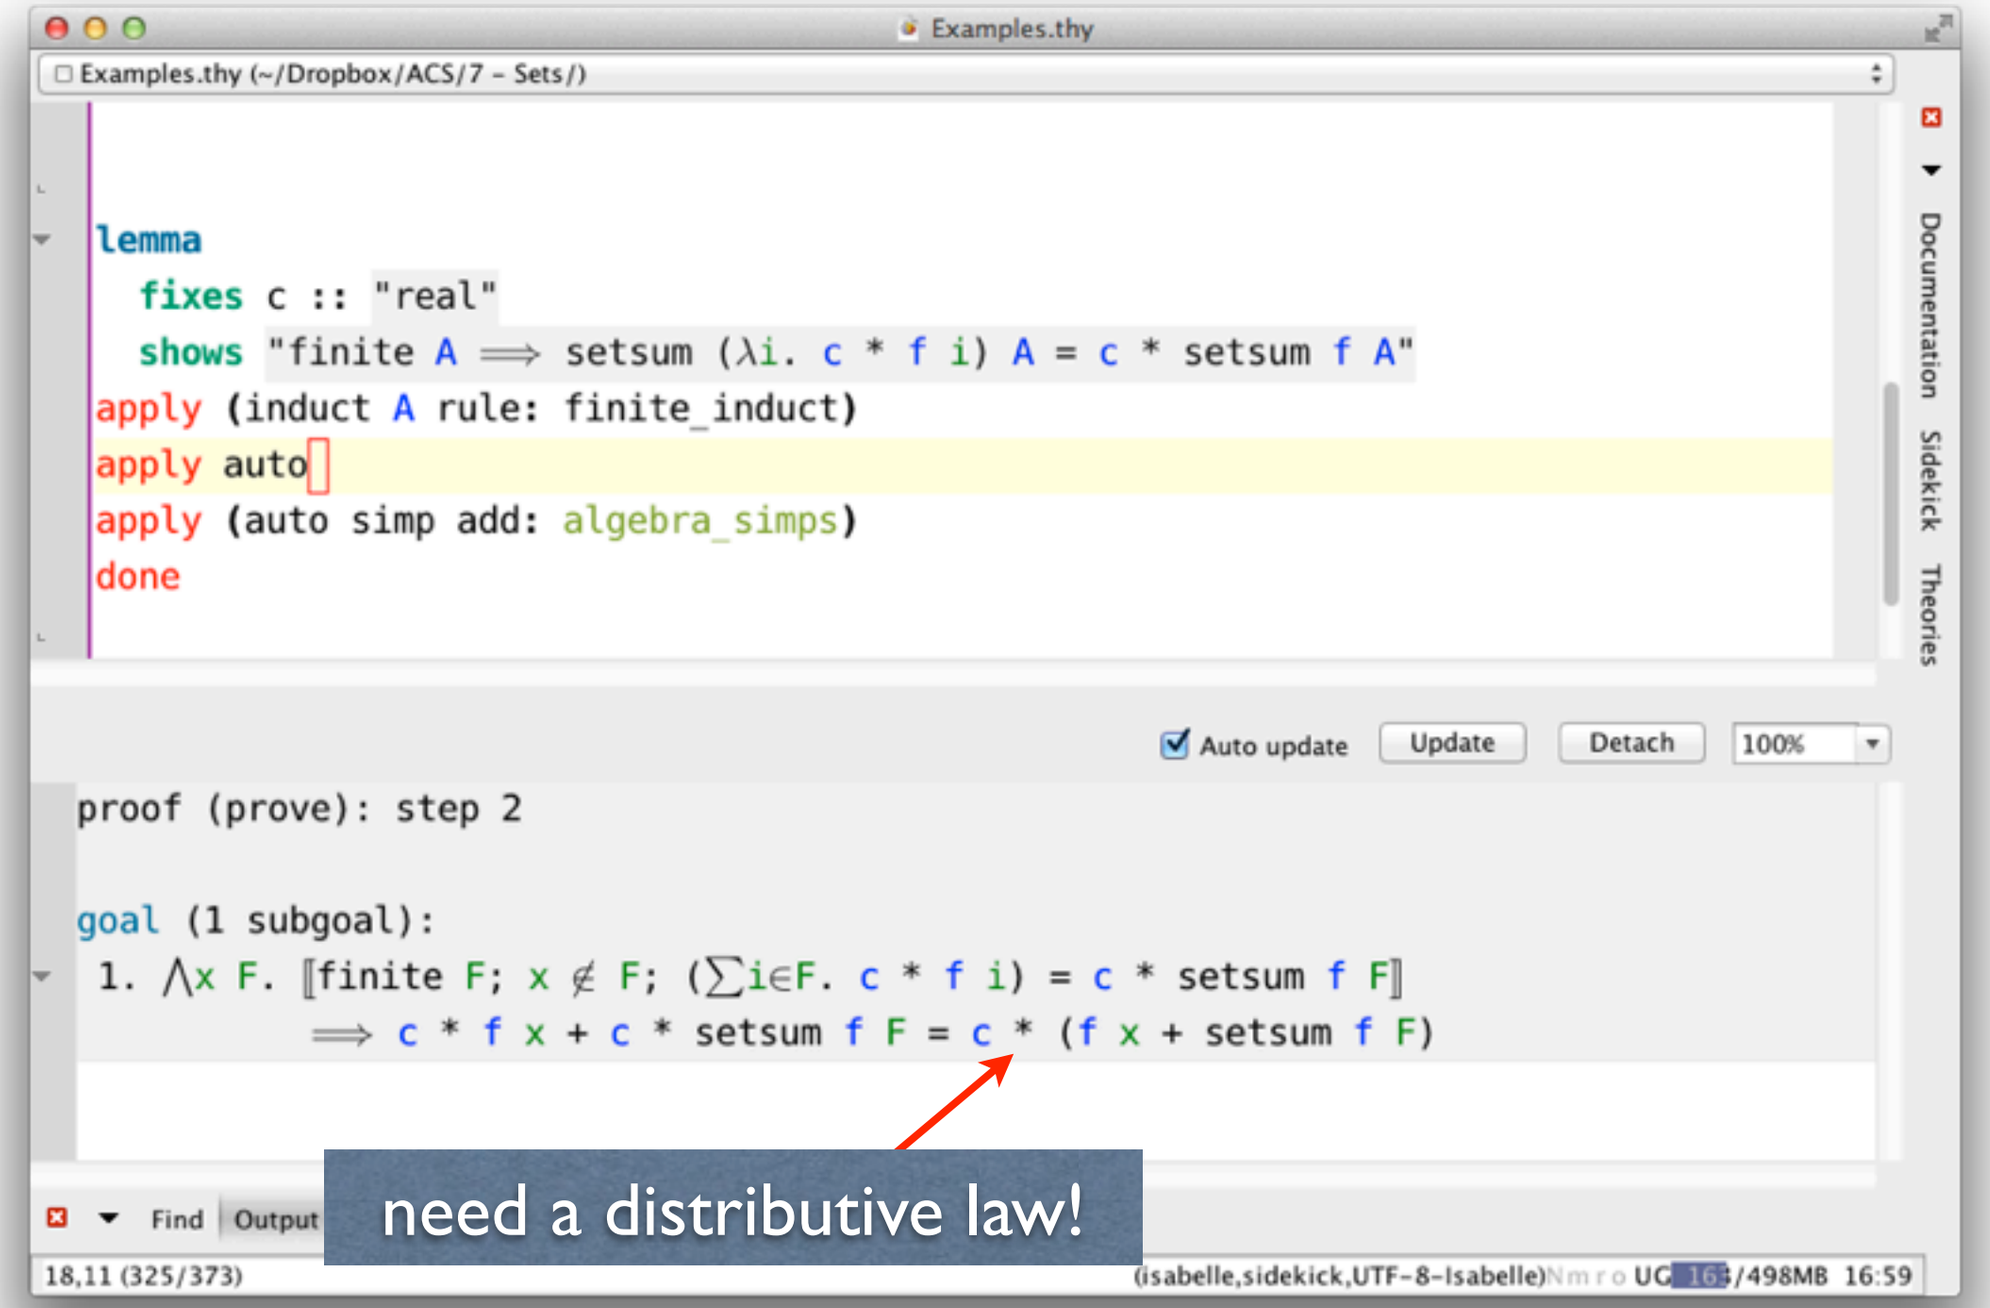

Finished!

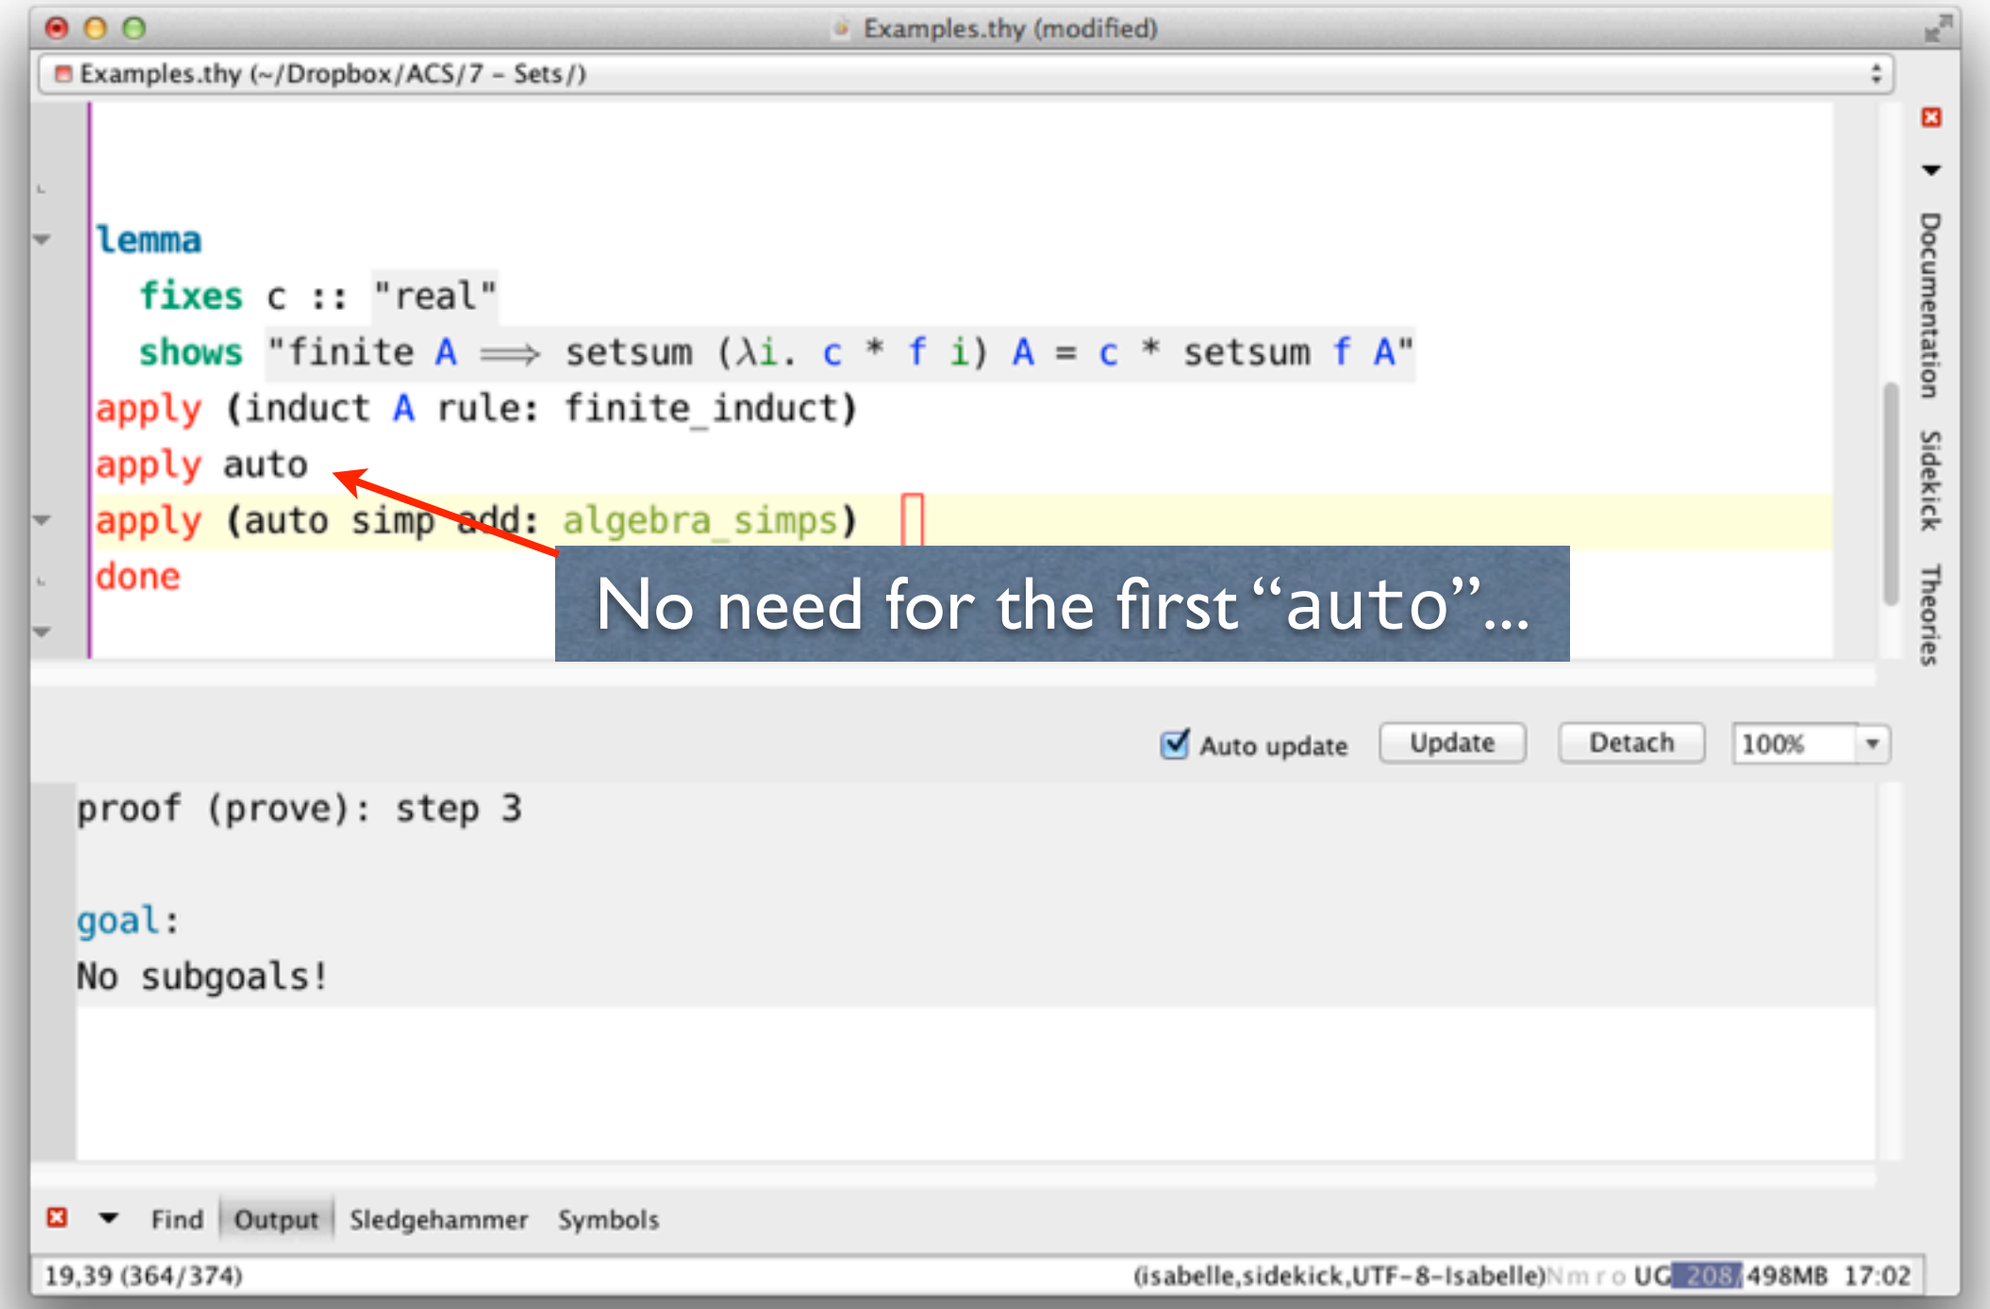

Recall that `algebra_simps` is a list of simplification rules for multiplying out algebraic expressions.

Counterexample Finding

- Don't waste time trying to prove false statements!
- Isabelle can find counterexamples quickly…
  - `quickcheck`: random testing of executable specifications (broadly interpreted)
  - `nitpick`: a general, SAT-based disprover
  - `try`: calls both of those (and sledgehammer)
- Type these commands right in the document.

Quickcheck Example

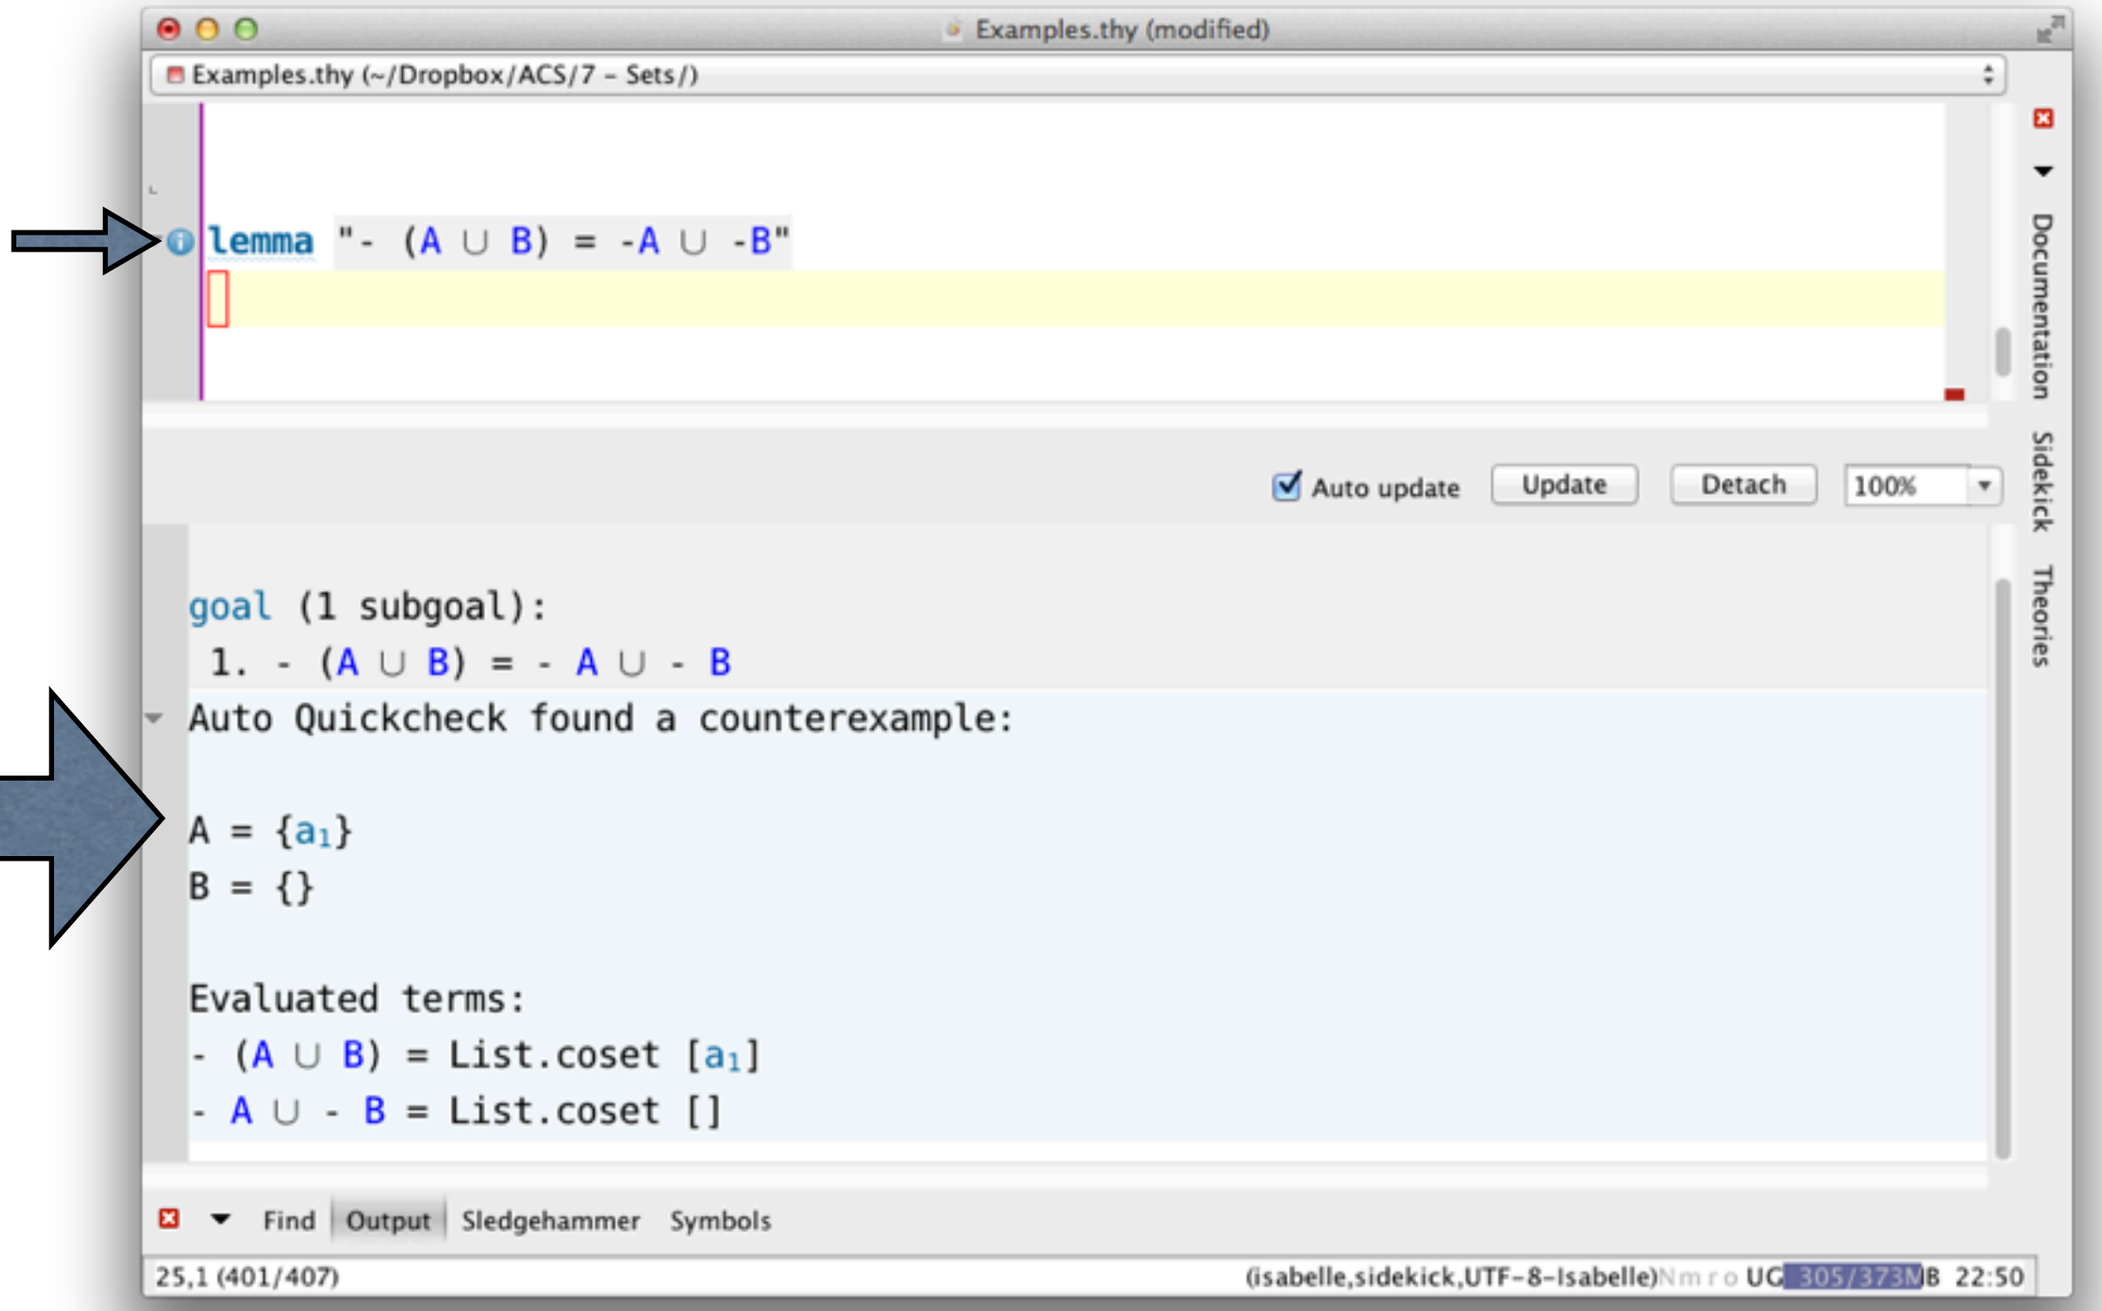

A minimal calls to quickcheck is performed automatically. Auto nitpick and even auto sledgehammer can be configured in the plugin options.

They work especially well for functional programs, but work in other domains, as we see here.

Proving Theorems about Sets

- It is not practical to learn all the built-in lemmas.
- Instead, try an automatic proof method:
  - `auto`
  - `force`
  - `blast`
- Each uses the built-in library, comprising hundreds of facts, with powerful heuristics.

Finding Theorems about Sets

See the *Tutorial*, section **3.1.11 Finding Theorems**. Virtually all theorems loaded within Isabelle can be located using this function. Unfortunately, it does not locate theorems that are proved in external libraries.

Finding Theorems about Sets

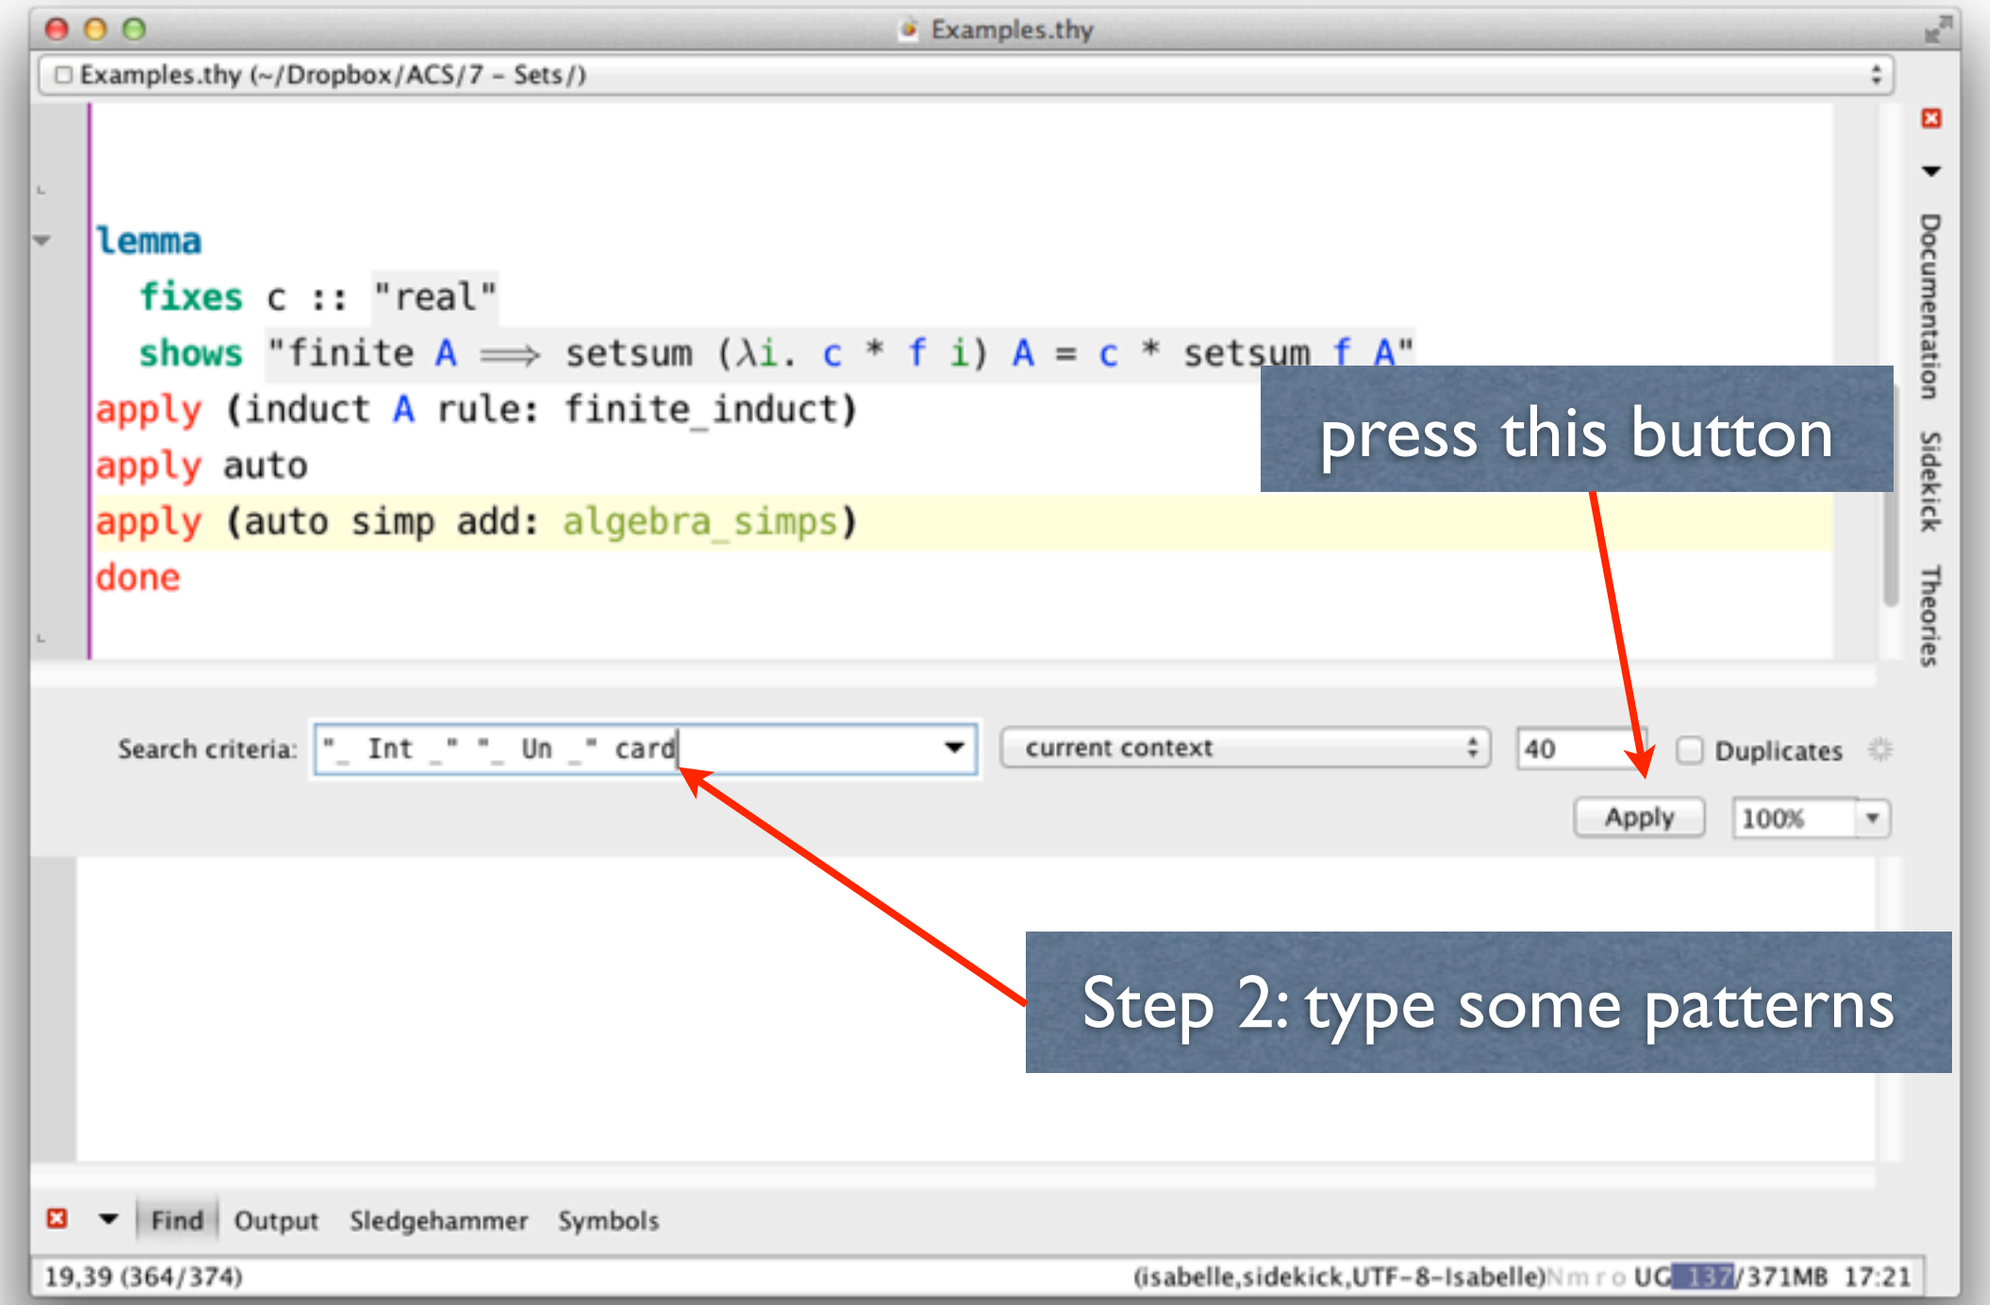

The easiest way to refer to infix operators is by entering small patterns, as shown above. More complex patterns are also permitted. The constraints are treated conjunctively: use additional constraints if you get too many results, and fewer constraints if you get no results.

Which Theorems Were Found?

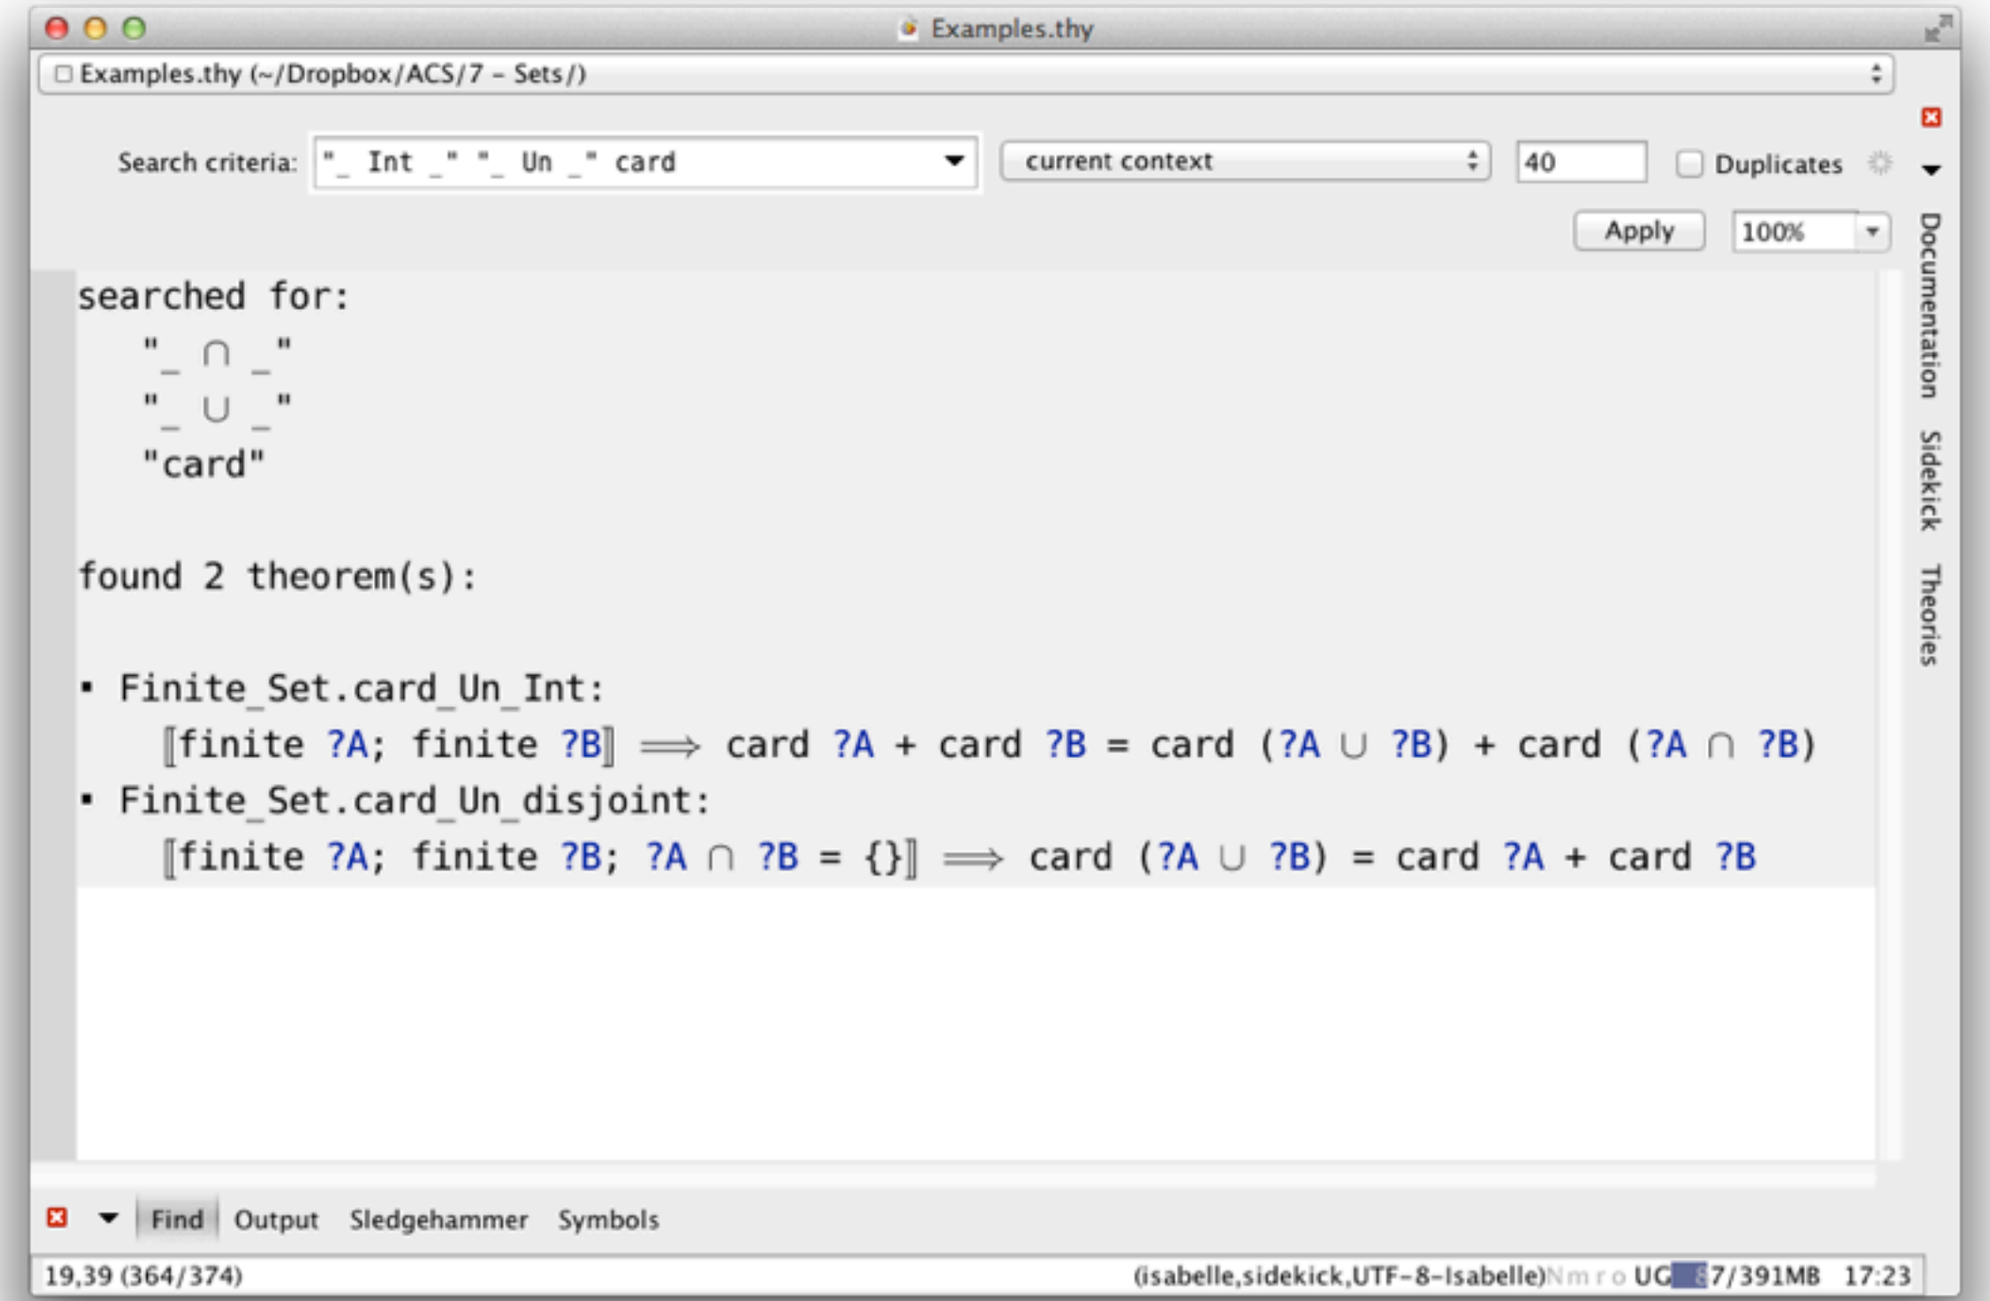

The Find panel, like all the other panels, can be detached or docked in various places so that it is always available.

Interactive Formal Verification
8: Inductive Definitions

Lawrence C Paulson
Computer Laboratory
University of Cambridge

Overview

- An introduction to inductive definitions
- A demonstration of their use in reasoning about finite sets.
- Demonstrating more automation: the `arith` proof method and the *sledgehammer* proof tool.

Defining a Set Inductively

- The set of even numbers is the least set such that
  - 0 is even.
  - If n is even, then $n+2$ is even.
- These can be viewed as *introduction rules*.
- We get an *induction principle* to express that no other numbers are even.
- Induction is used throughout mathematics, and to express the semantics of programming languages.

See the *Tutorial*, Chaper 7. **Inductively Defined Sets**.

Inductive Definitions in Isabelle

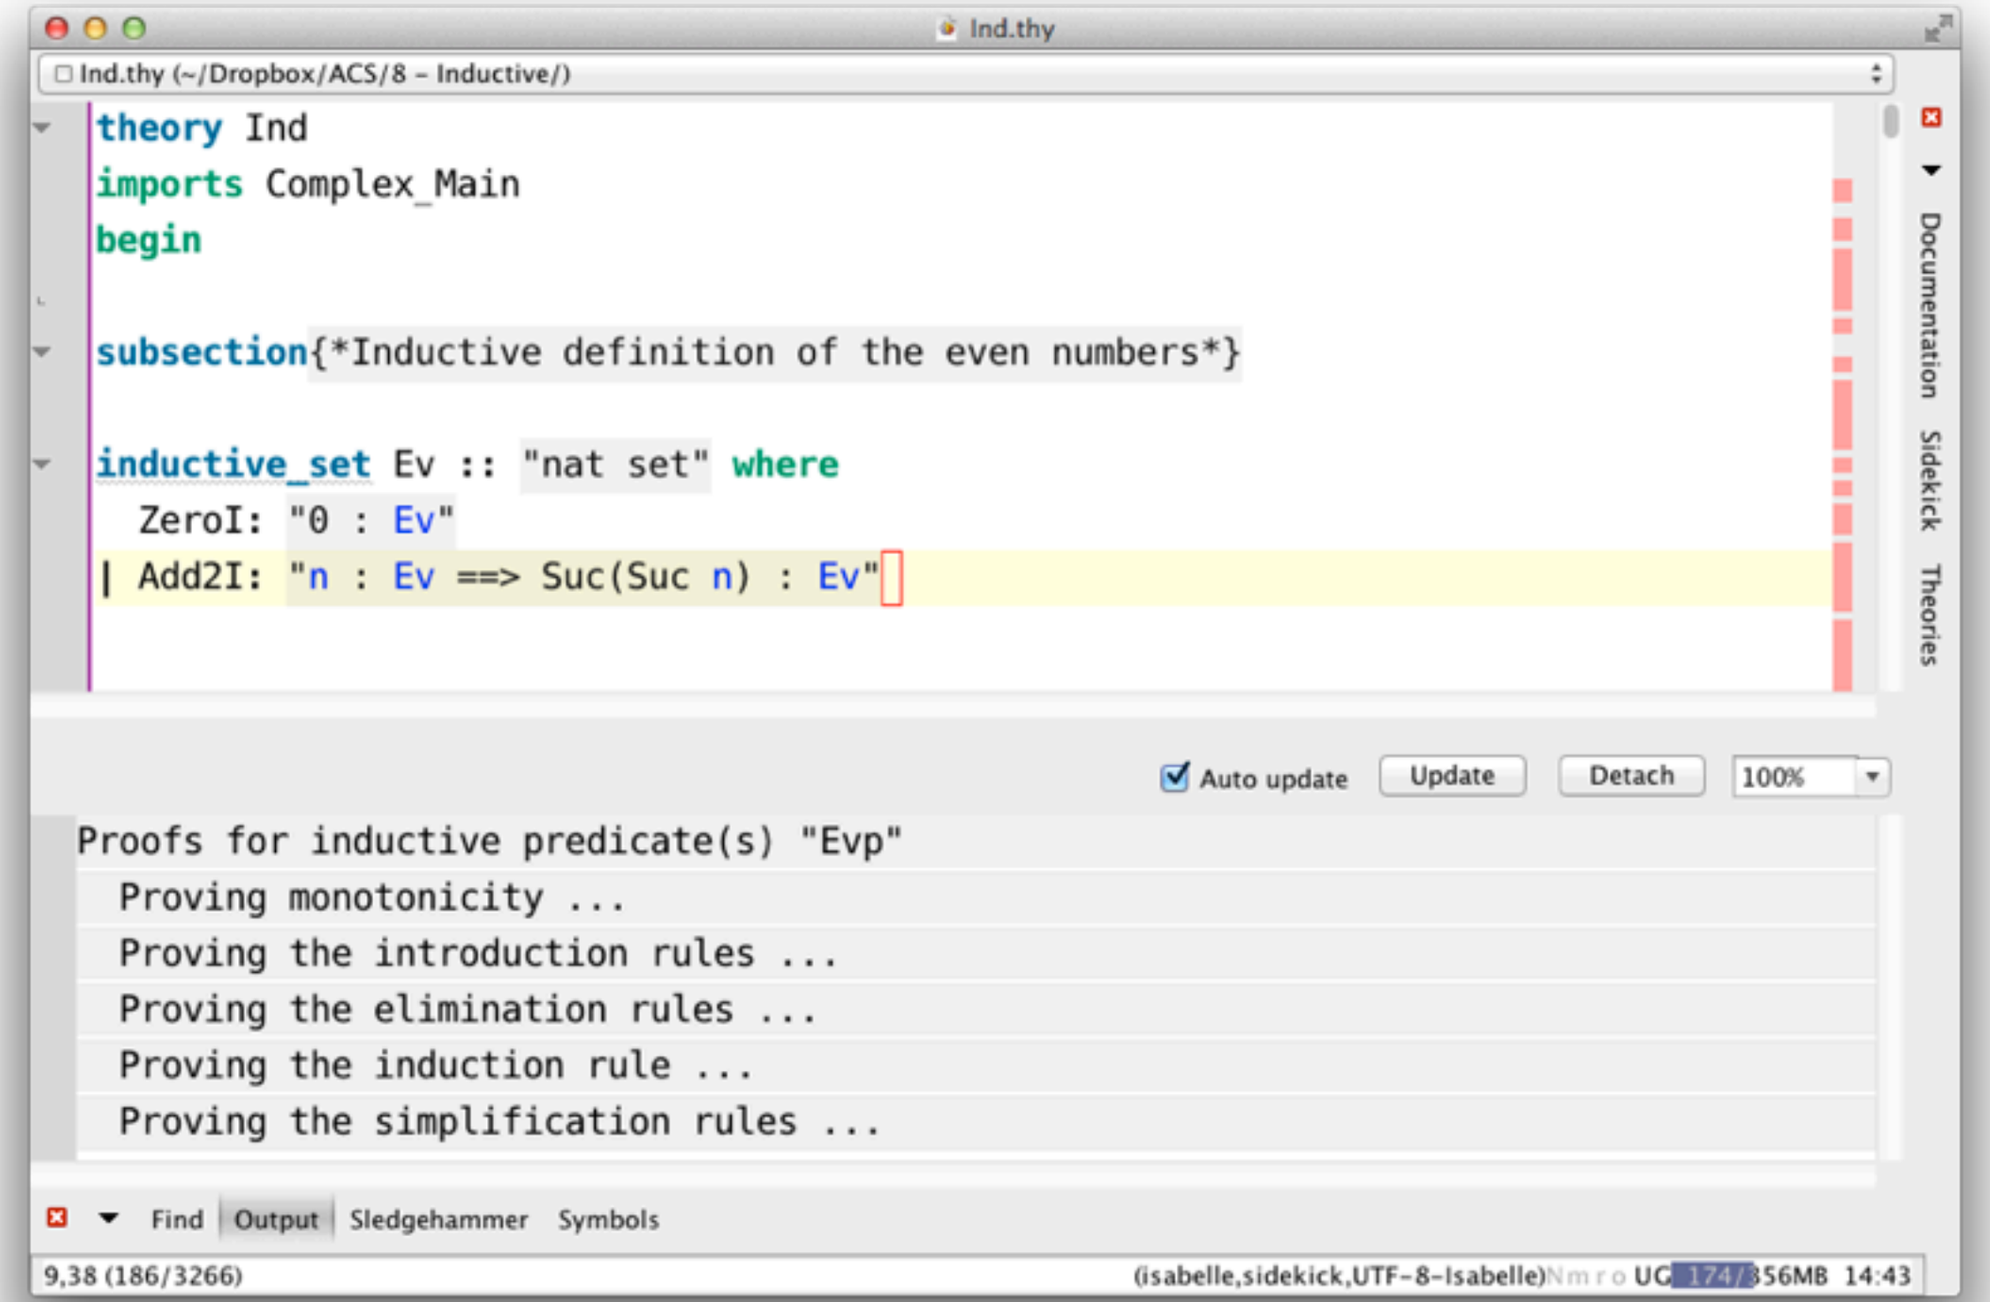

The *Tutorial* discusses precisely the same example, section 7.1.1.

Even Numbers Belong to Ev

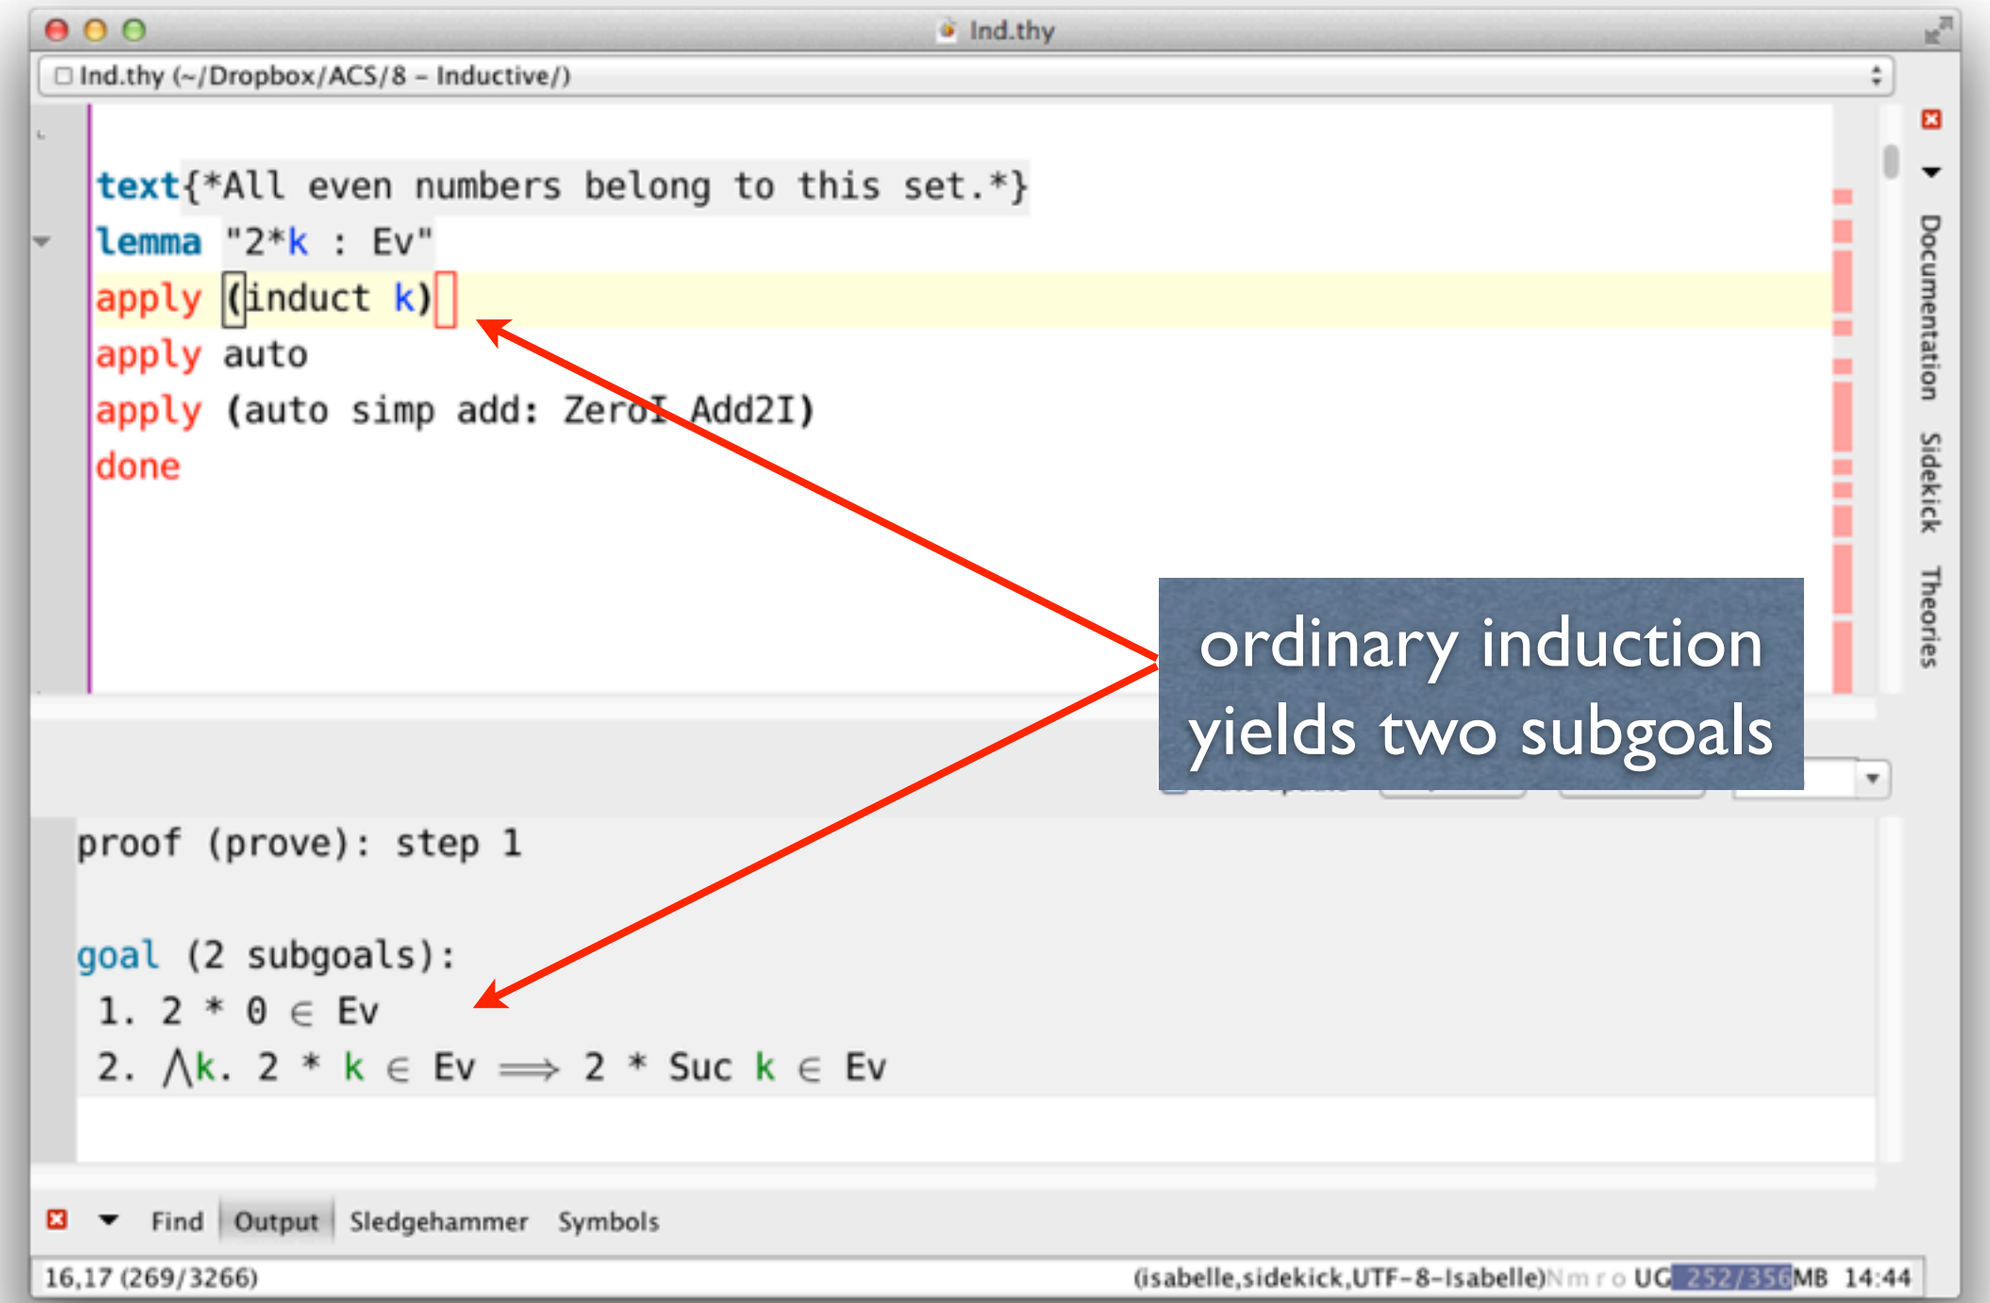

Proving Set Membership

Ind.thy

Ind.thy (~/Dropbox/ACS/8 – Inductive/)

```
text{*All even numbers belong to this set.*}
lemma "2*k : Ev"
apply (induct k)
apply auto
apply (auto simp add: ZeroI Add2I)
done
```

after simplification, the subgoals resemble the introduction rules

```
proof (prove): step 2

goal (2 subgoals):
 1. 0 ∈ Ev
 2. ⋀k. 2 * k ∈ Ev ⟹ Suc (Suc (2 * k)) ∈ Ev
```

Documentation Sidekick Theories

Find Output Sledgehammer Symbols

17,11 (280/3266) (isabelle,sidekick,UTF-8-Isabelle)Nmro UG 44/387MB 14:44

Finishing the Proof

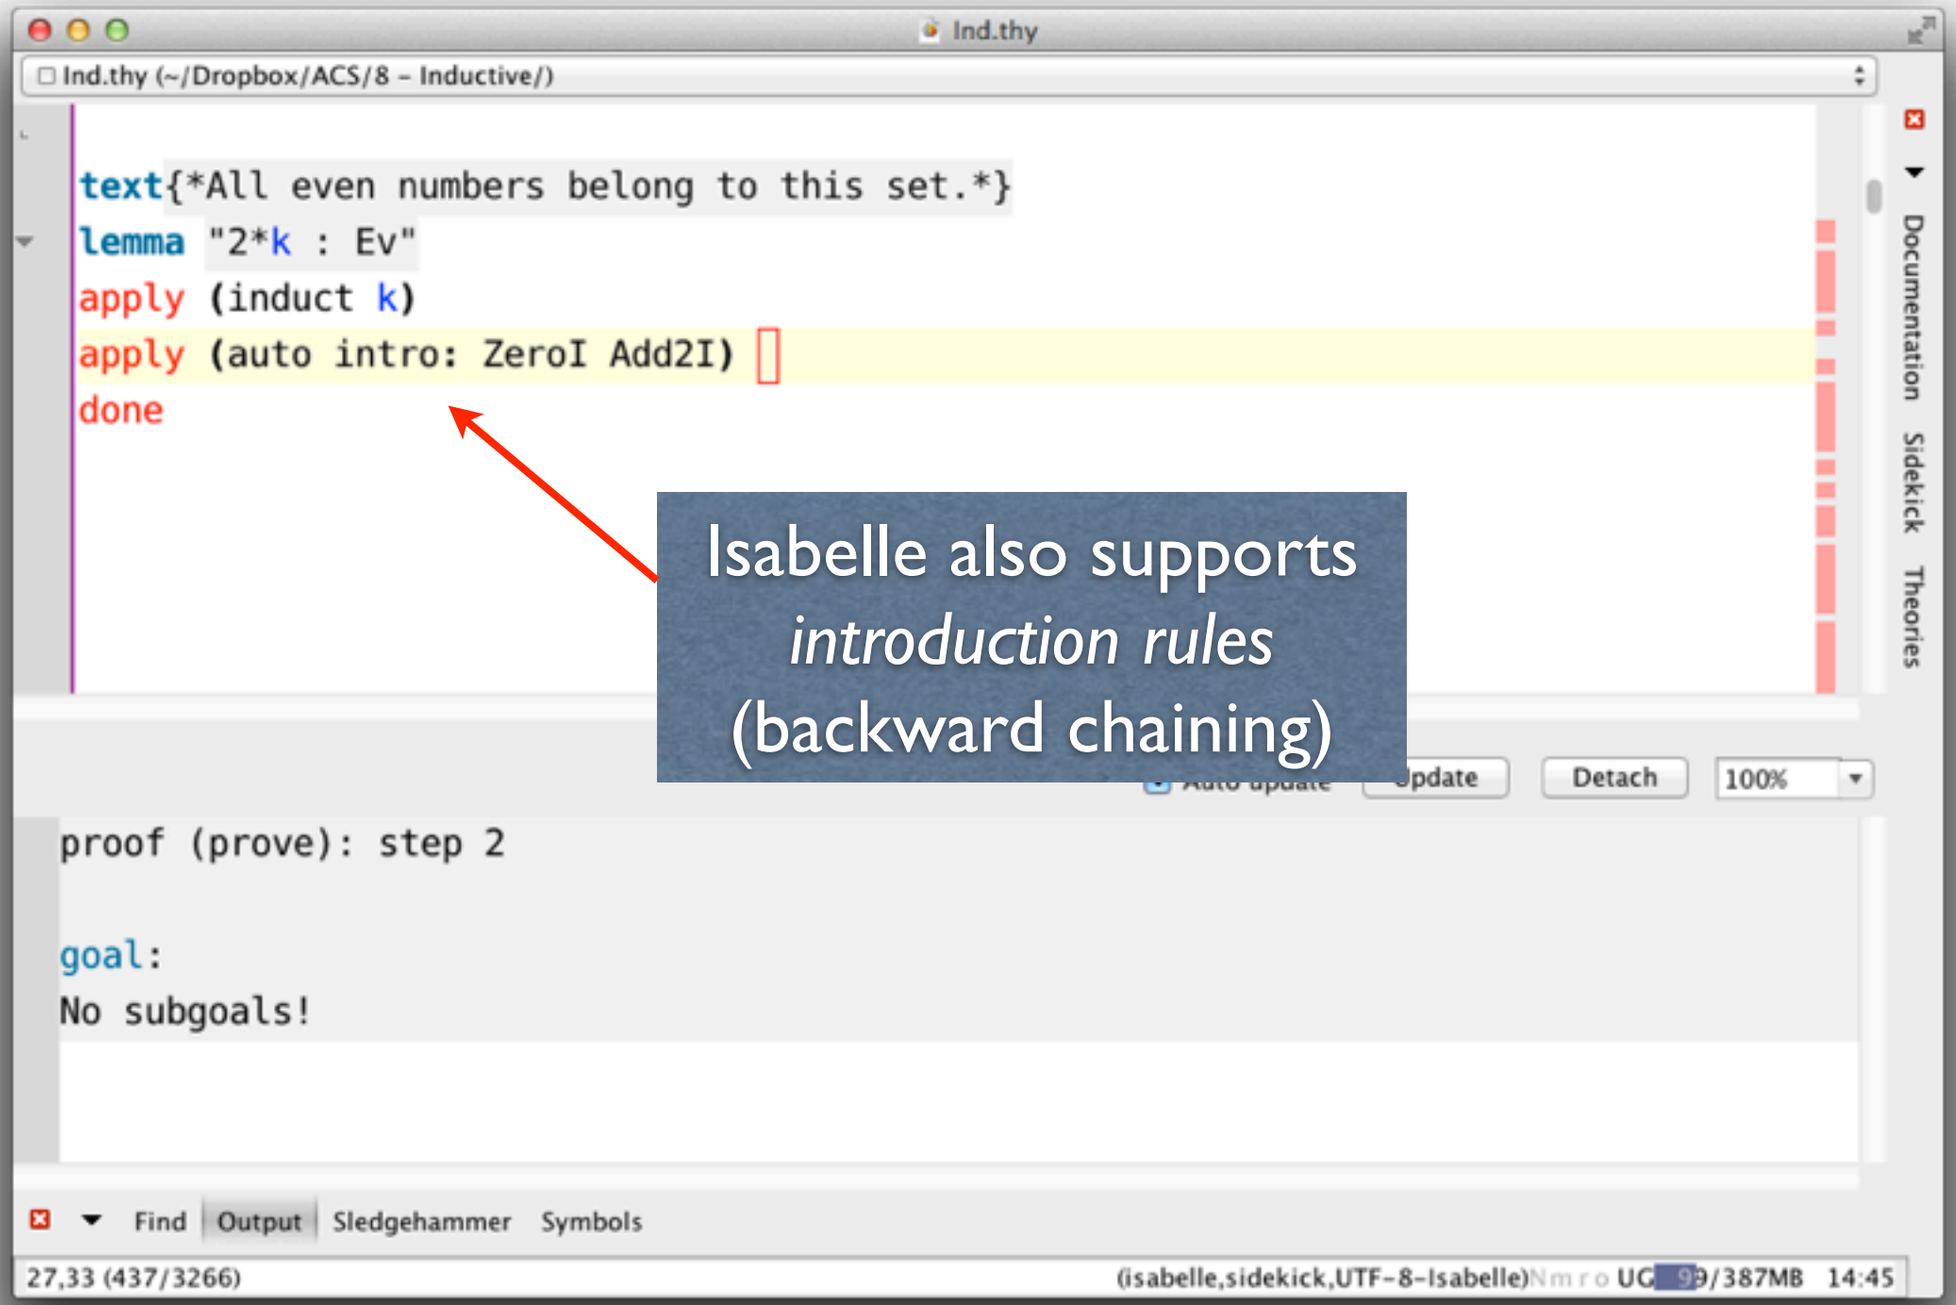

Rule Induction

- Proving something about *every* element of the set.
- It expresses that the inductive set is *minimal*.
- It is sometimes called "induction on derivations"

- There is a *base case* for every non-recursive introduction rule
- ...and an *inductive step* for the other rules.

Ev Has only Even Numbers

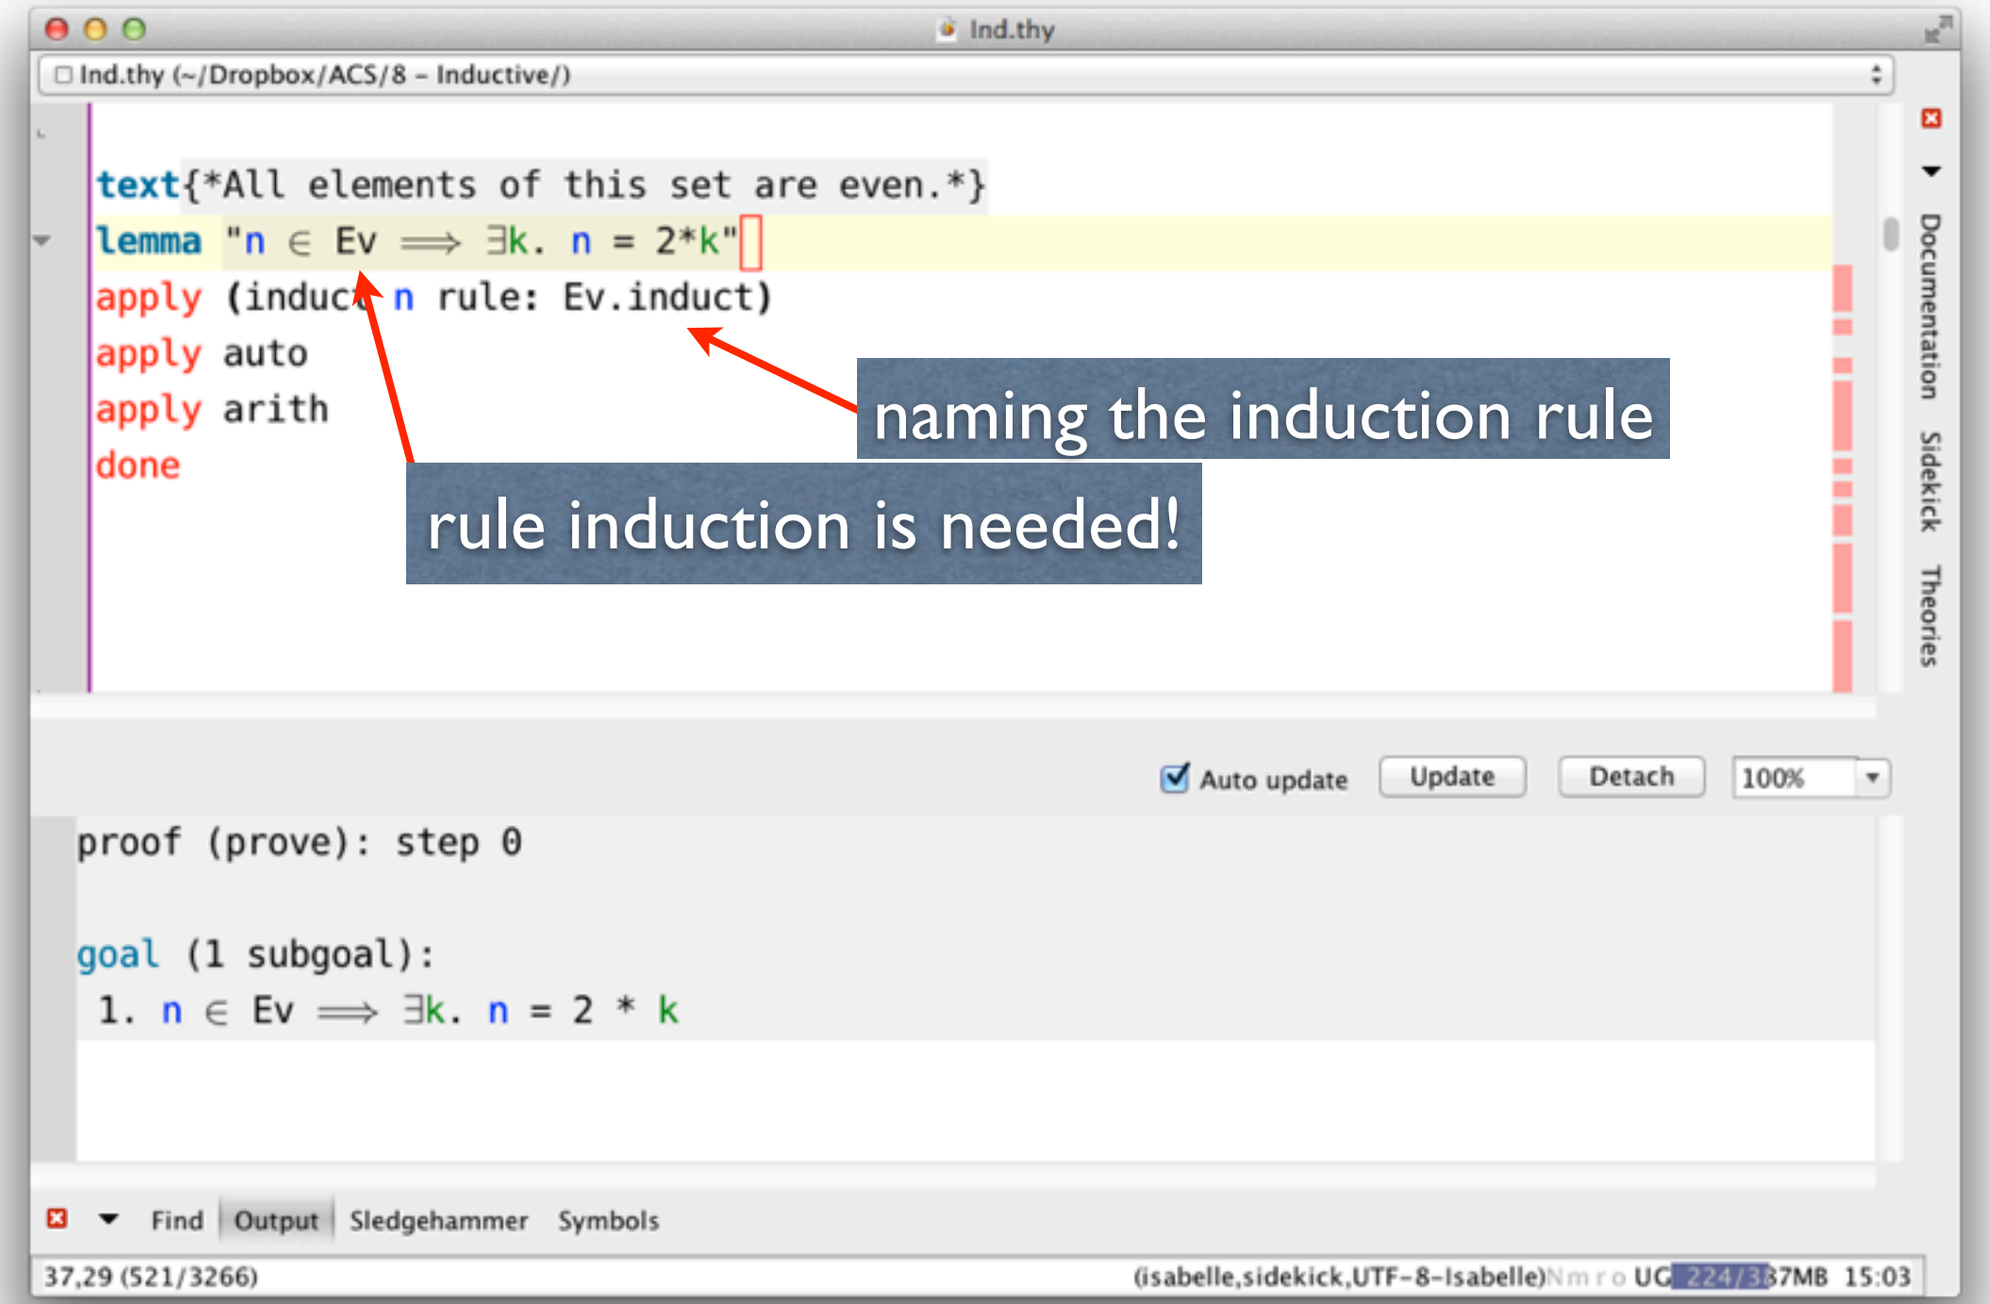

The classic sign that we need rule induction is an occurrence of the inductive set as a premise of the desired result. Of course, sometimes the theorem can be proved by referring to other facts that have been previously proved using rule induction.

An Example of Rule Induction

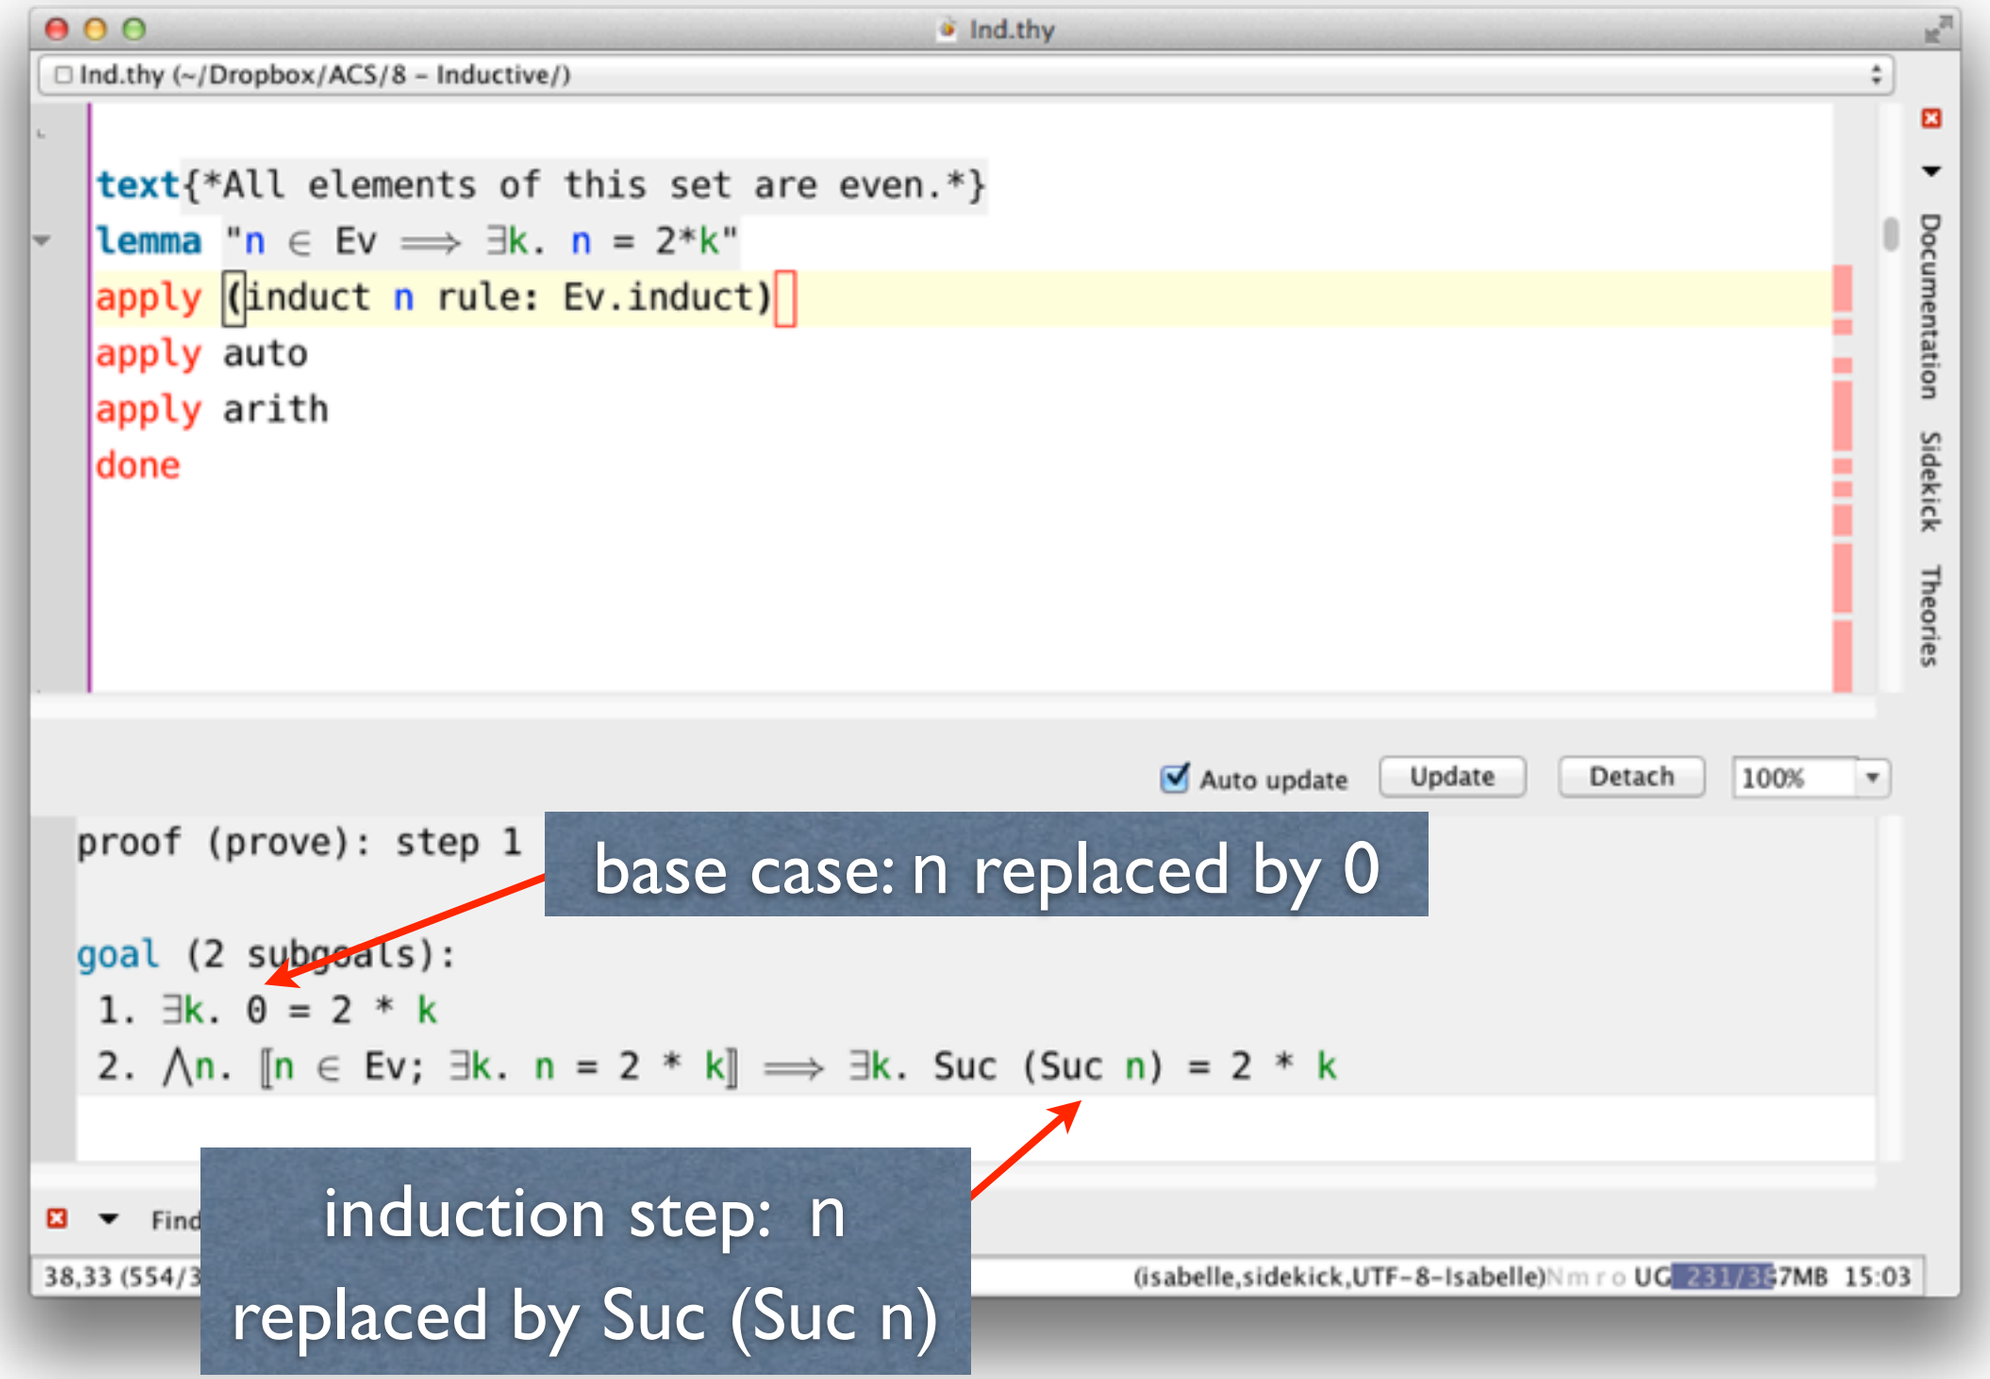

One Tricky Goal Left!

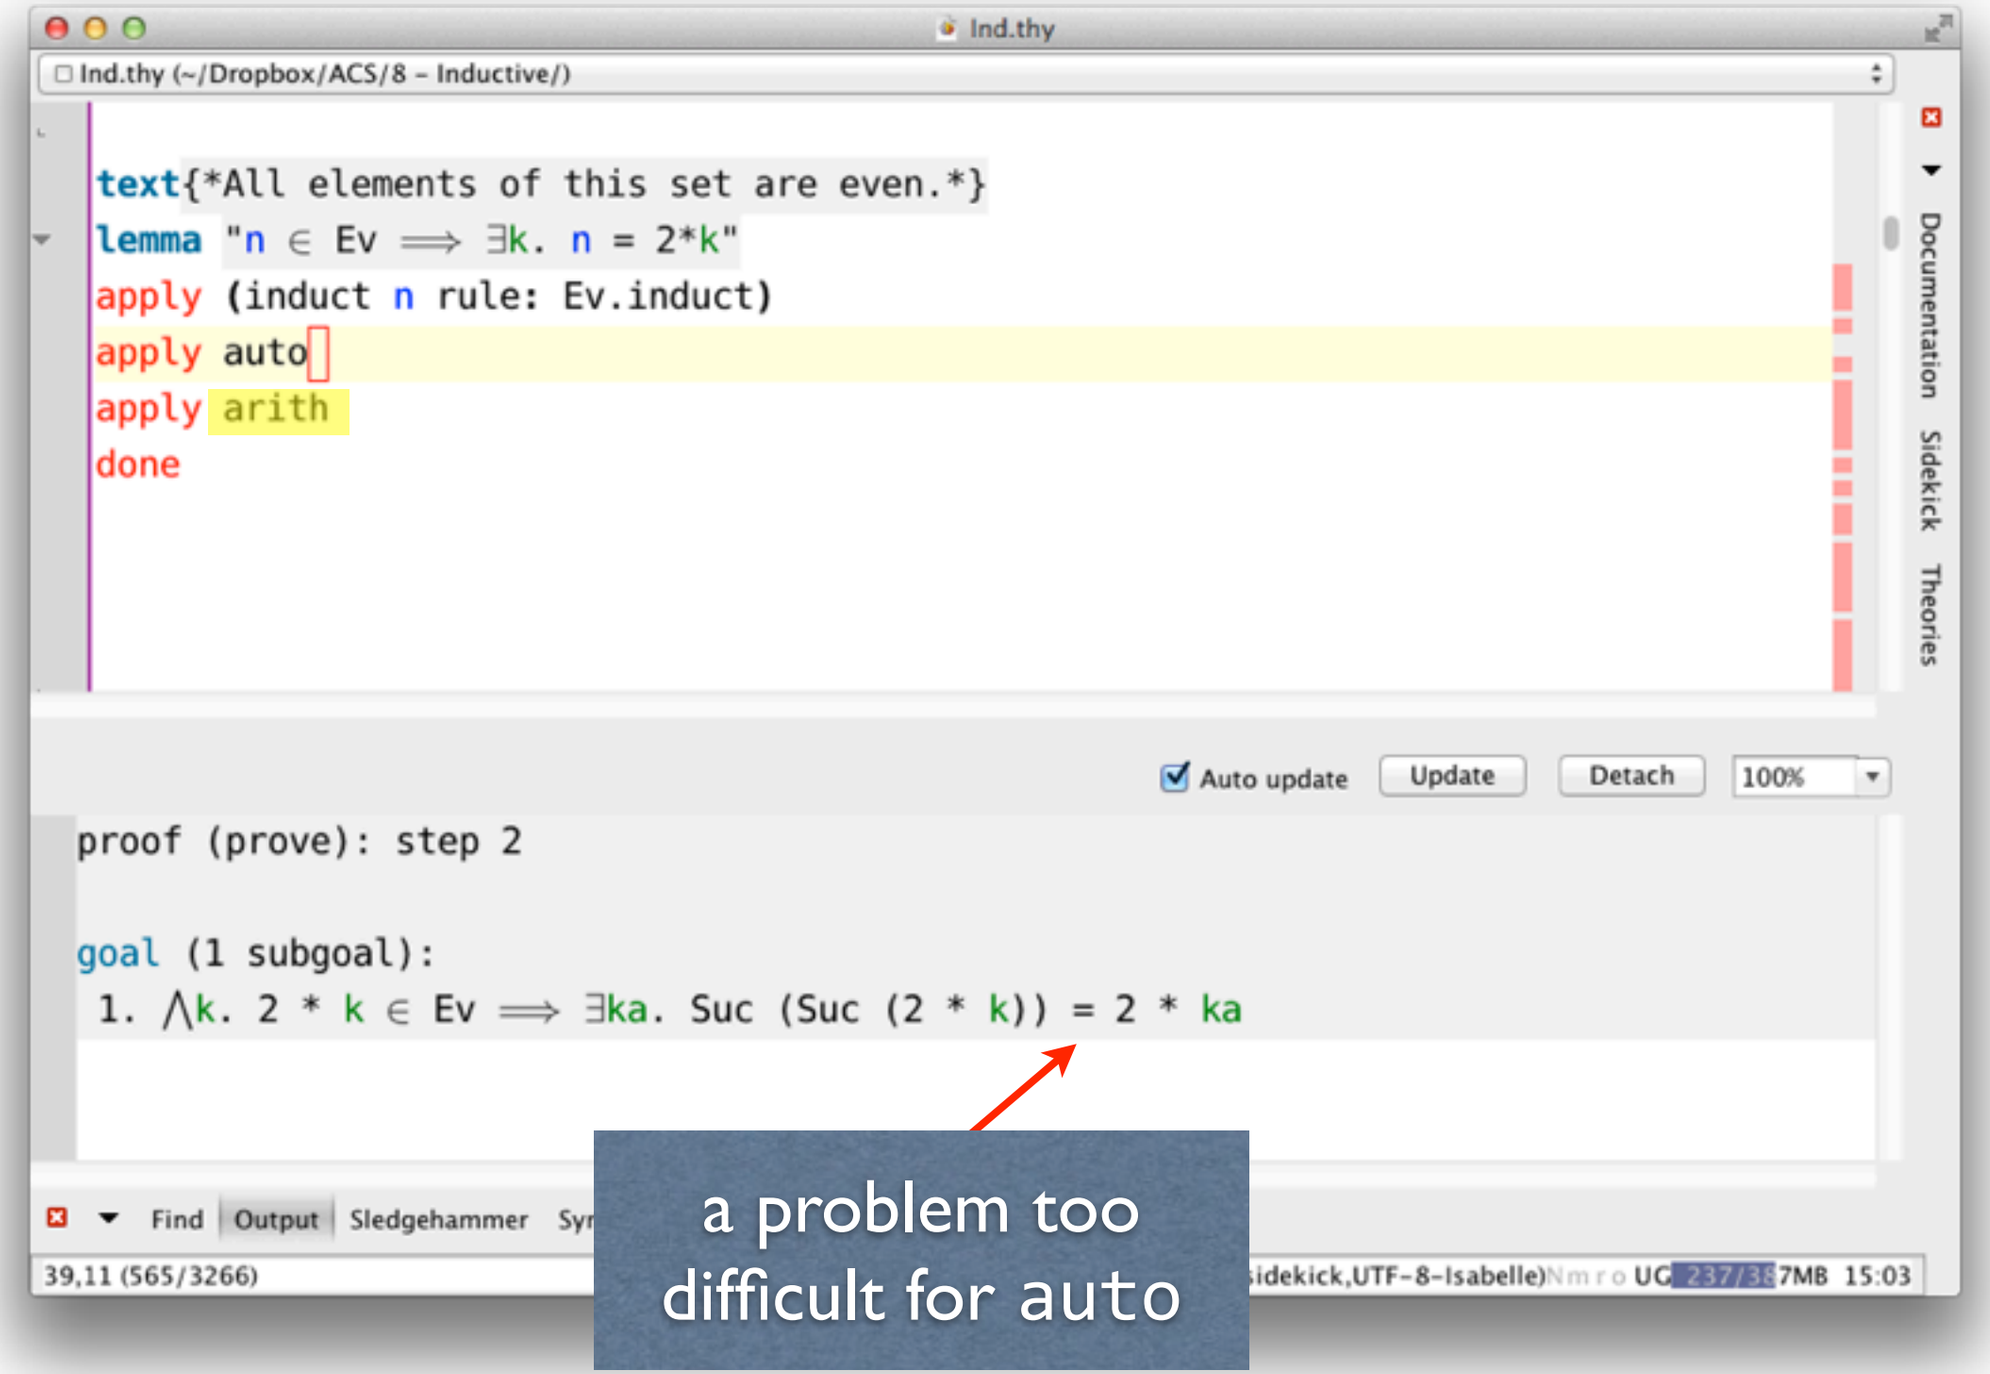

The `auto` method provides some support for arithmetic. However, complicated arithmetic arguments require specialised proof methods.

The arith Proof Method

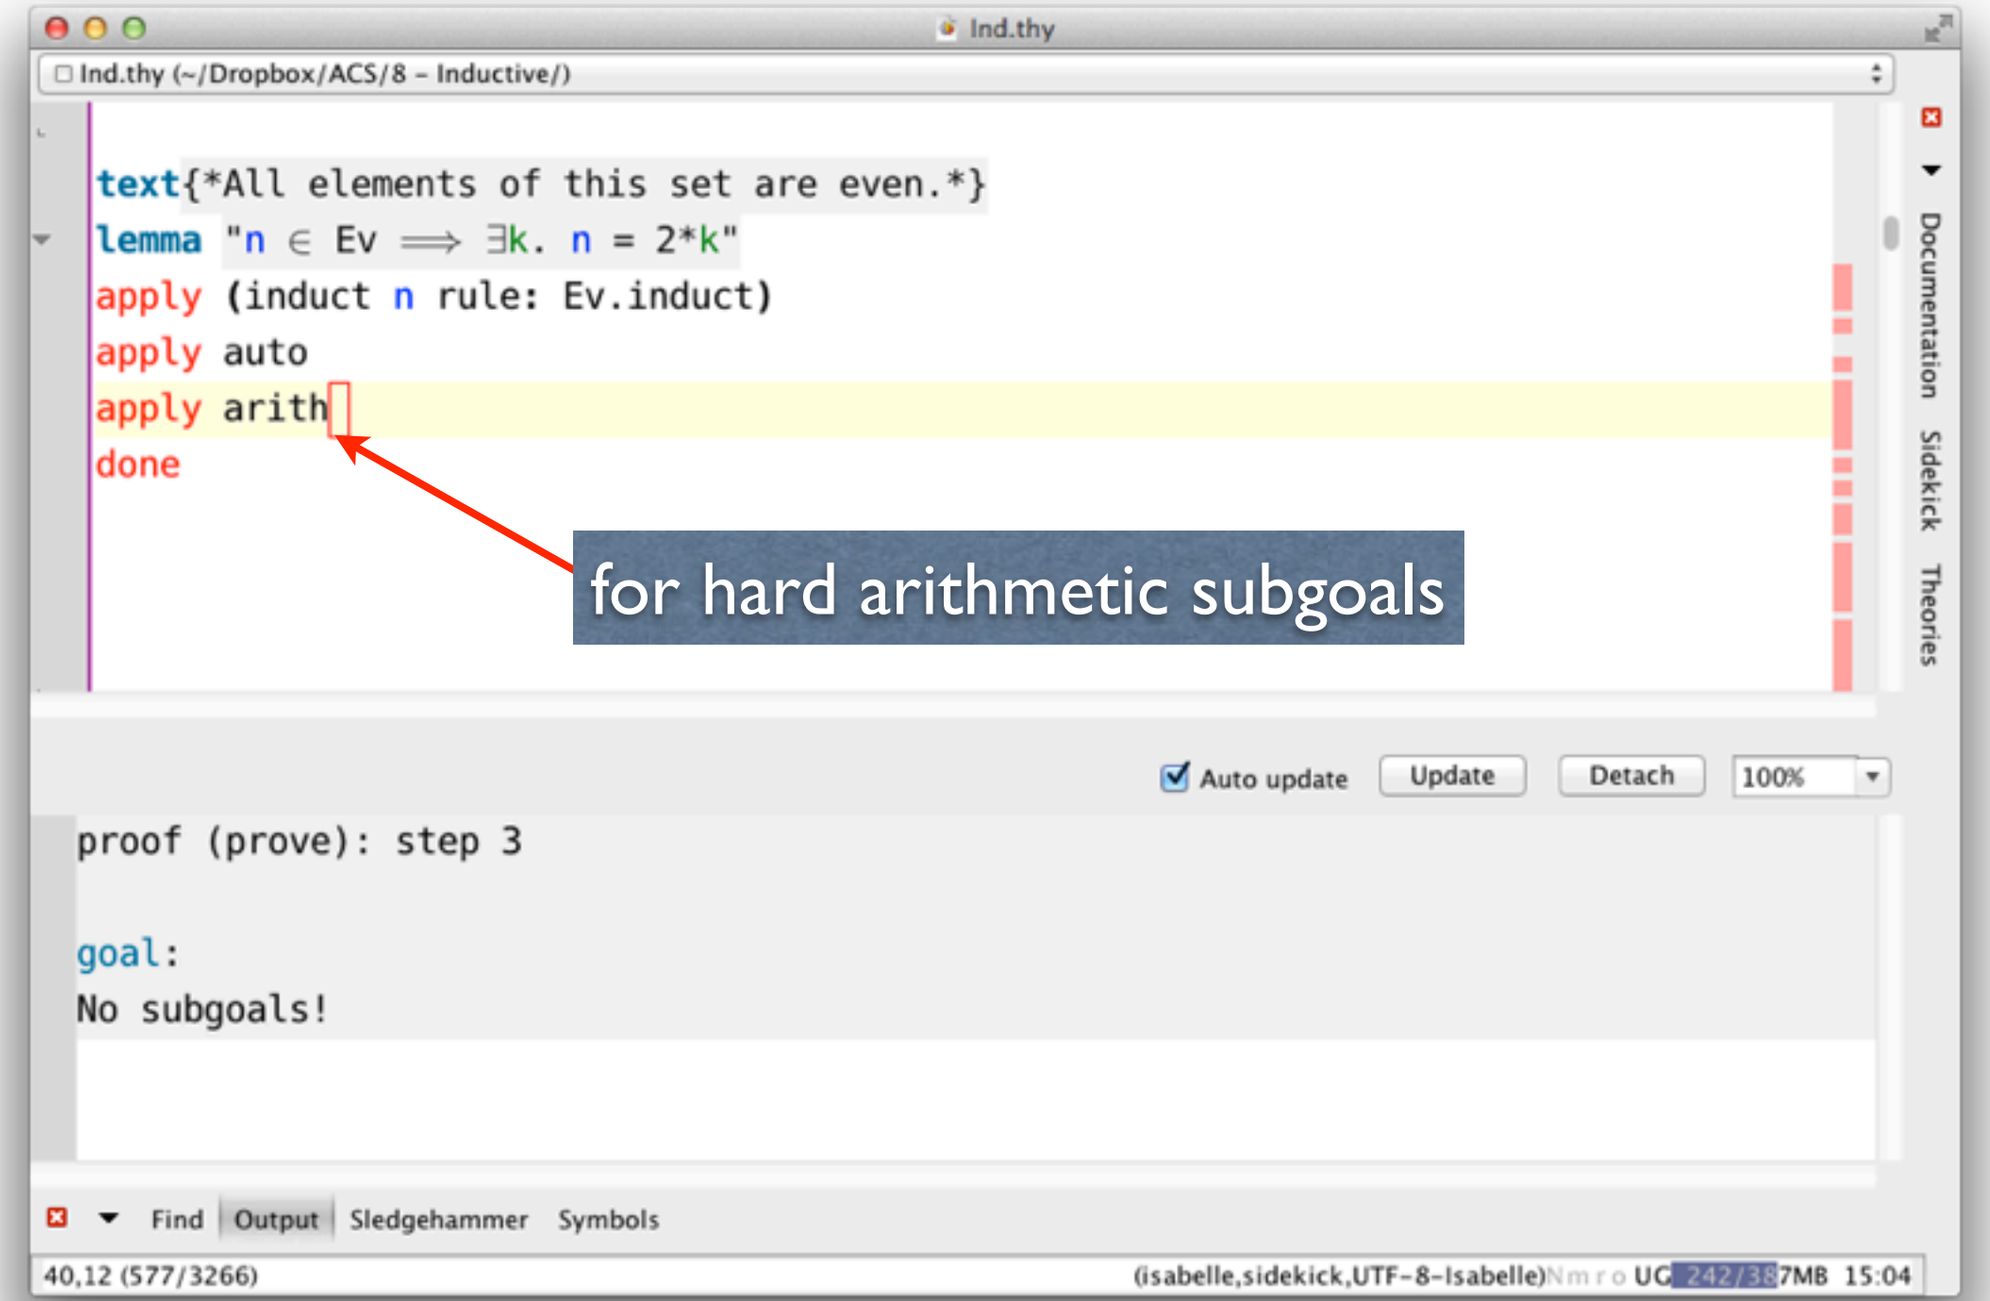

Aside: Linear Arithmetic

- A *decidable* class of formulas
  - Allowing the operators $+ - < \leq =$, and …
  - multiplication and division by *constants*: ×2, /2
- With slight variations, this class is decidable for the main arithmetic types.
- `auto` can solve simple arithmetic problems…
- `arith` handles logical connectives, quantifiers, etc.
- Decision procedures are necessary: proving arithmetic facts from first principles is too tedious.

Defining Finiteness

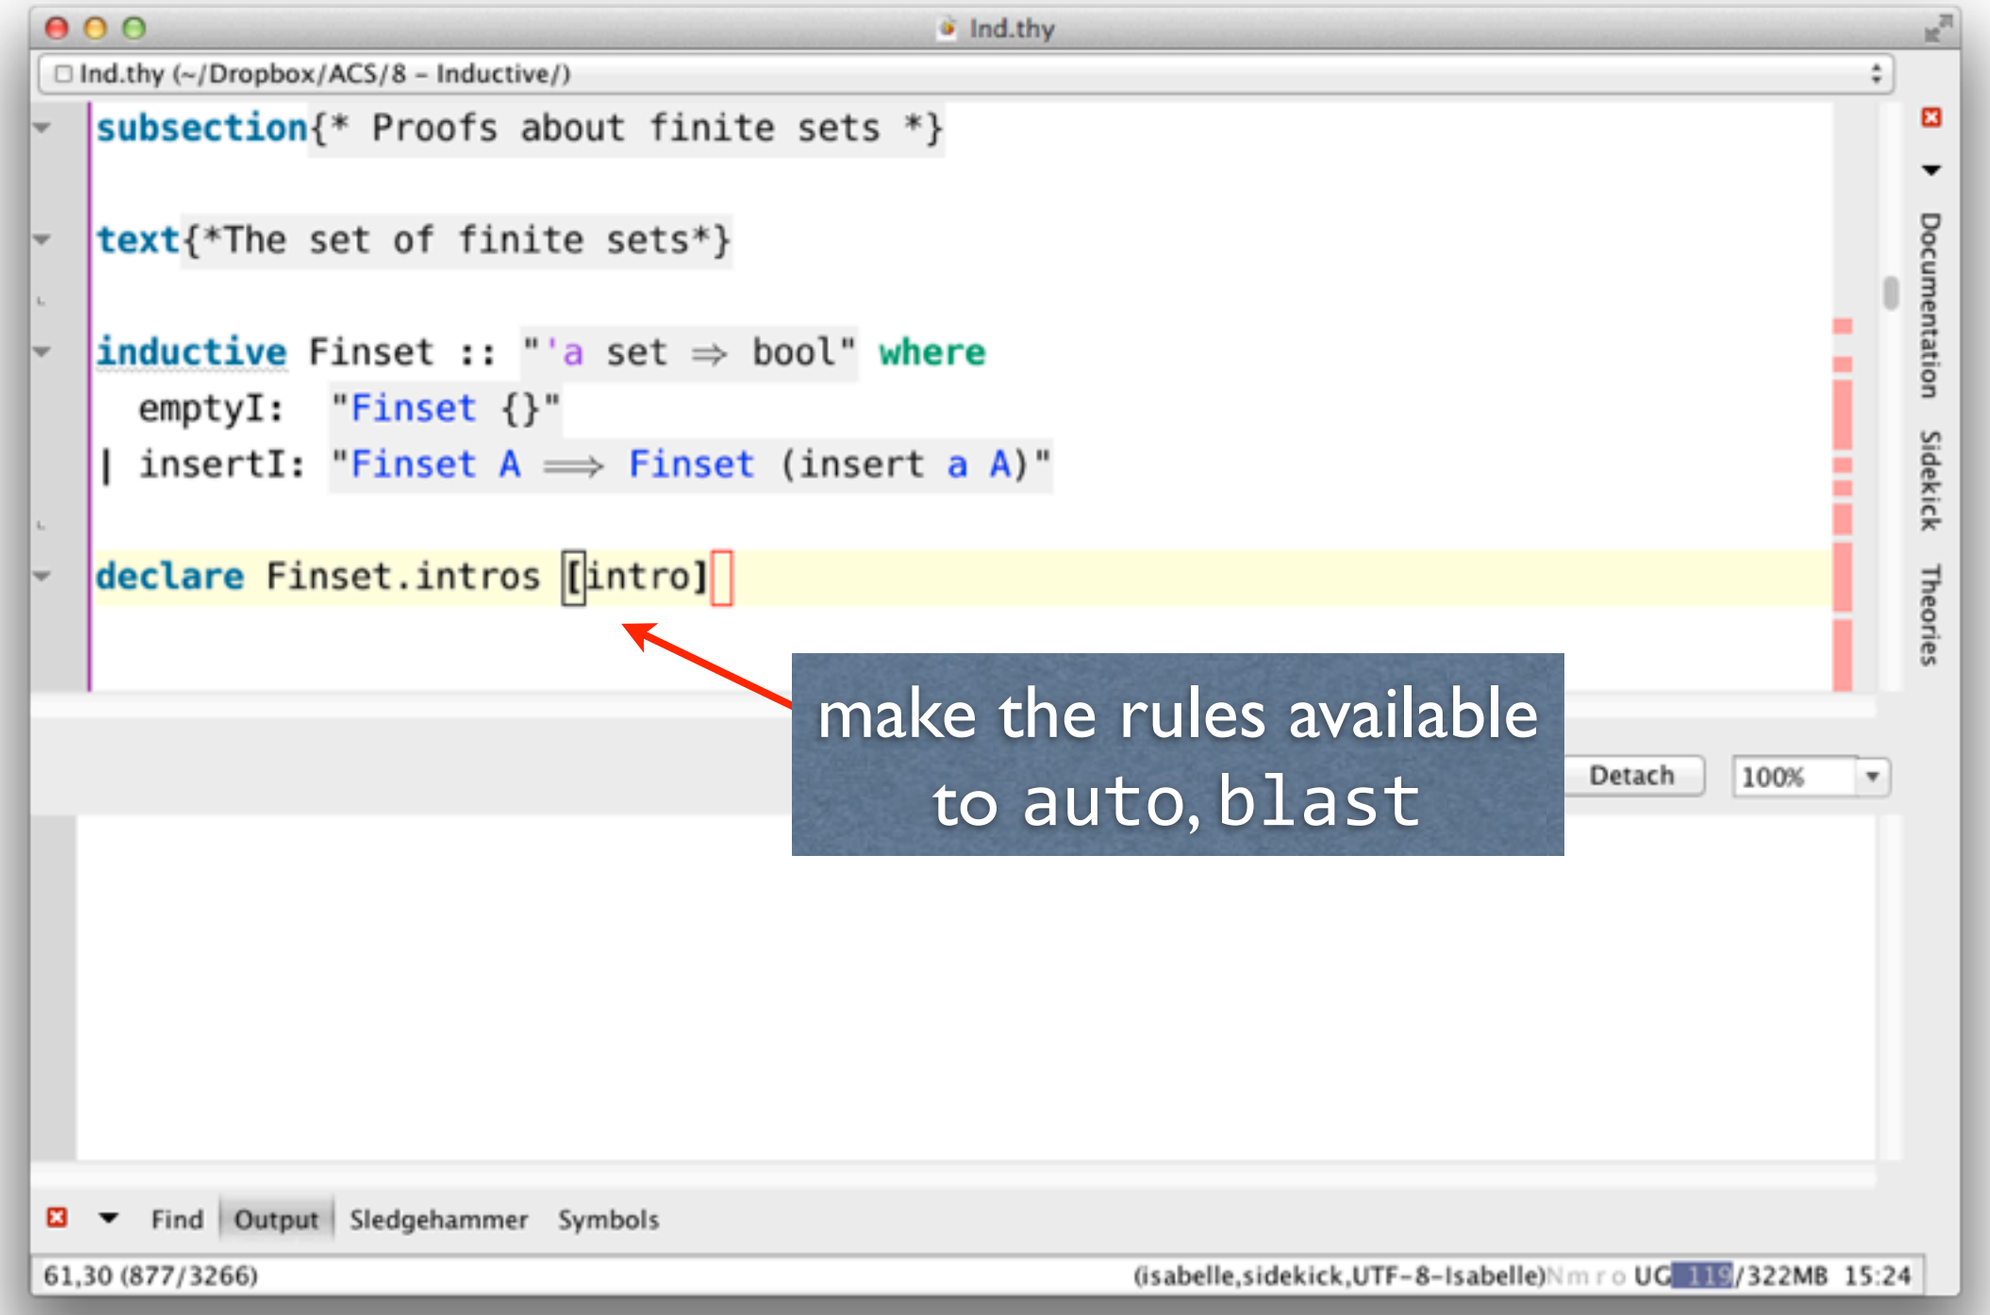

The empty set is finite. Adding one element to a finite set yields another finite set.

The Union of Two Finite Sets

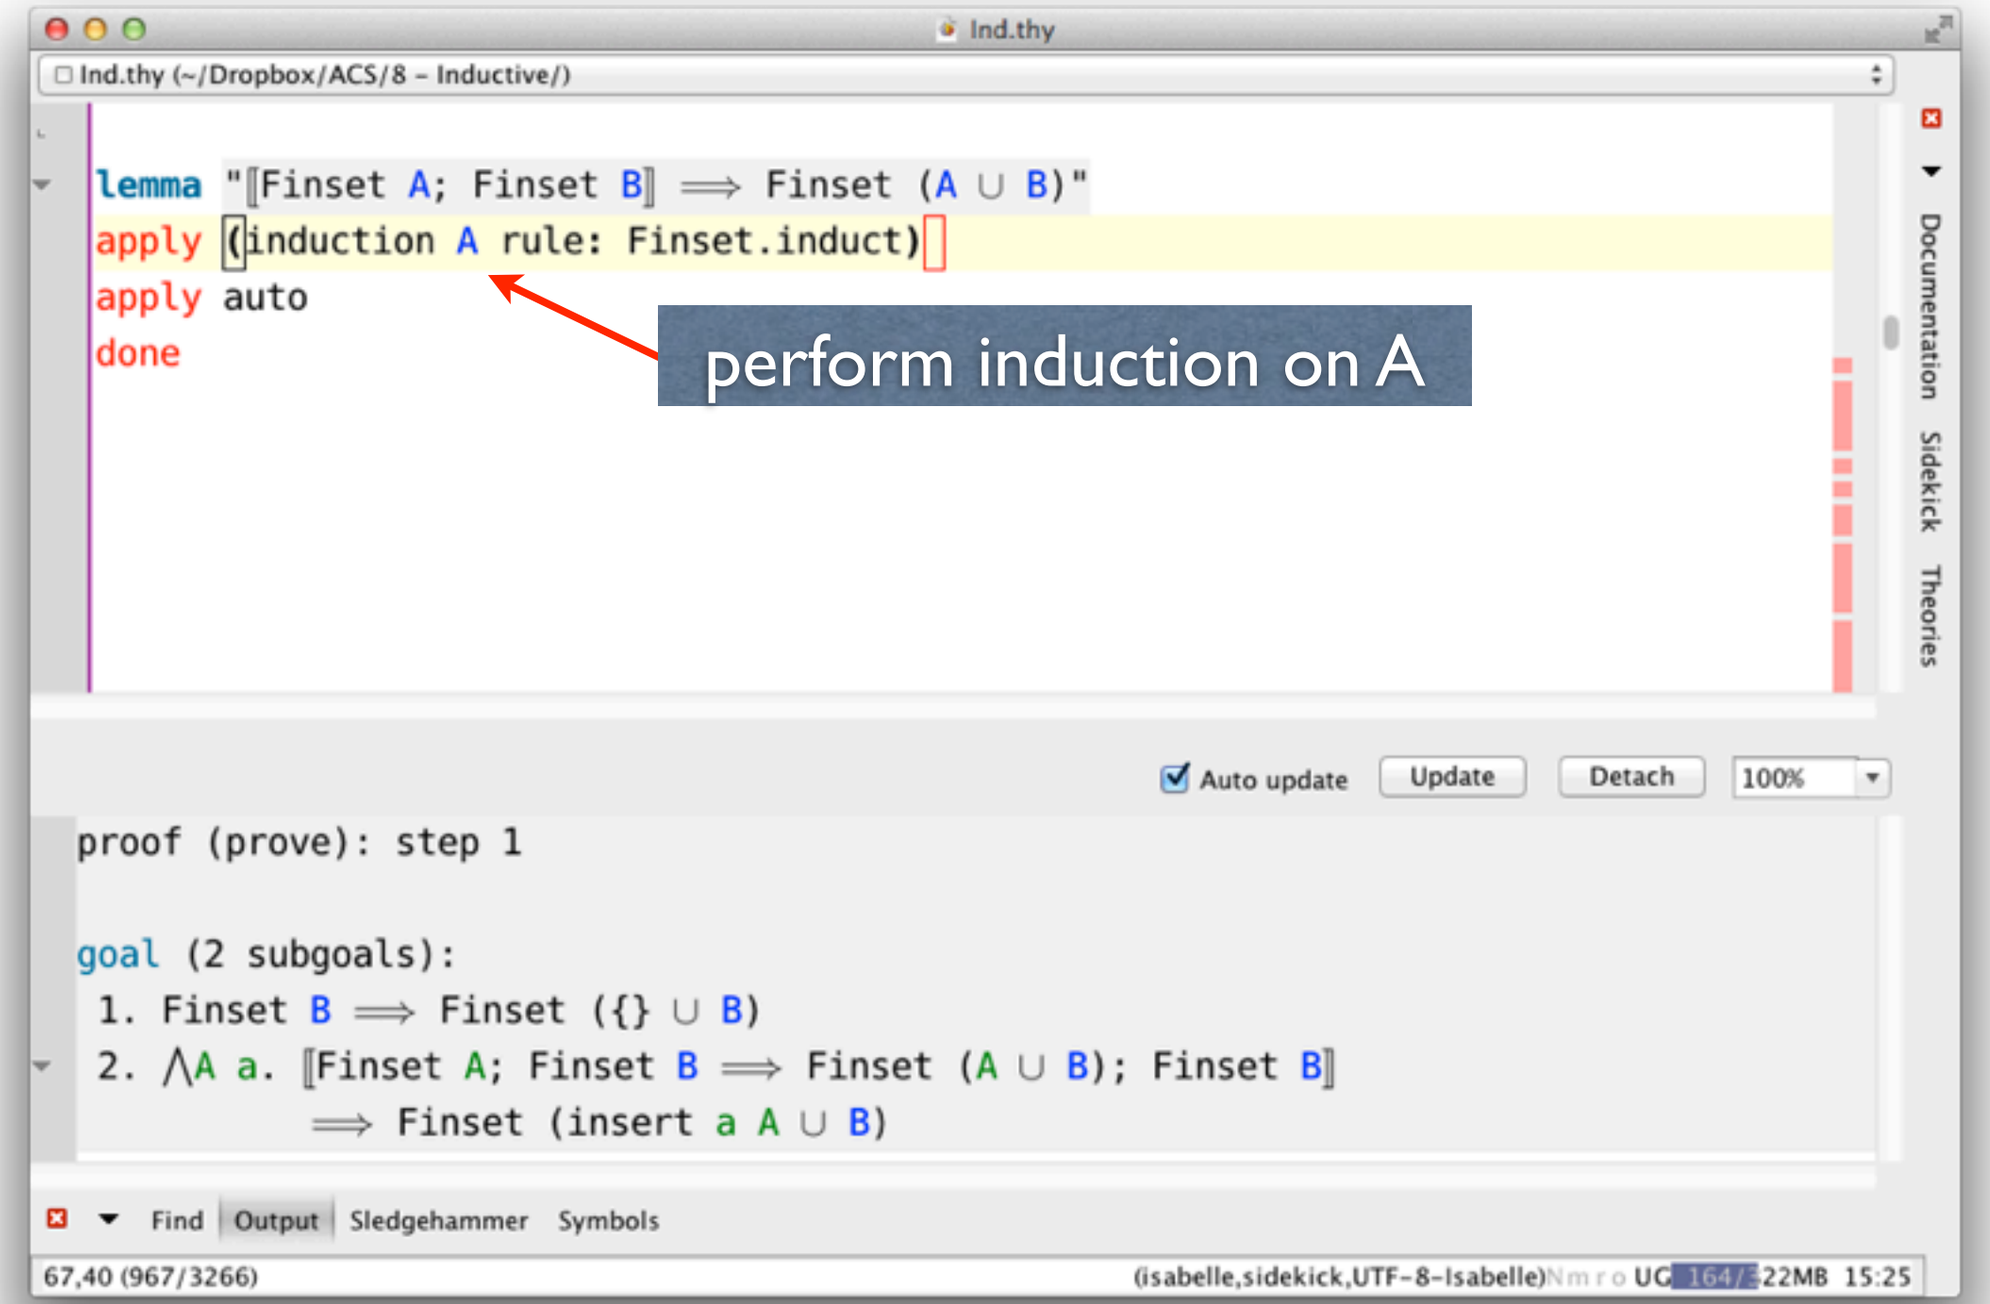

The goals are easily proved by the properties of sets and the introduction rules.

A Subset of a Finite Set

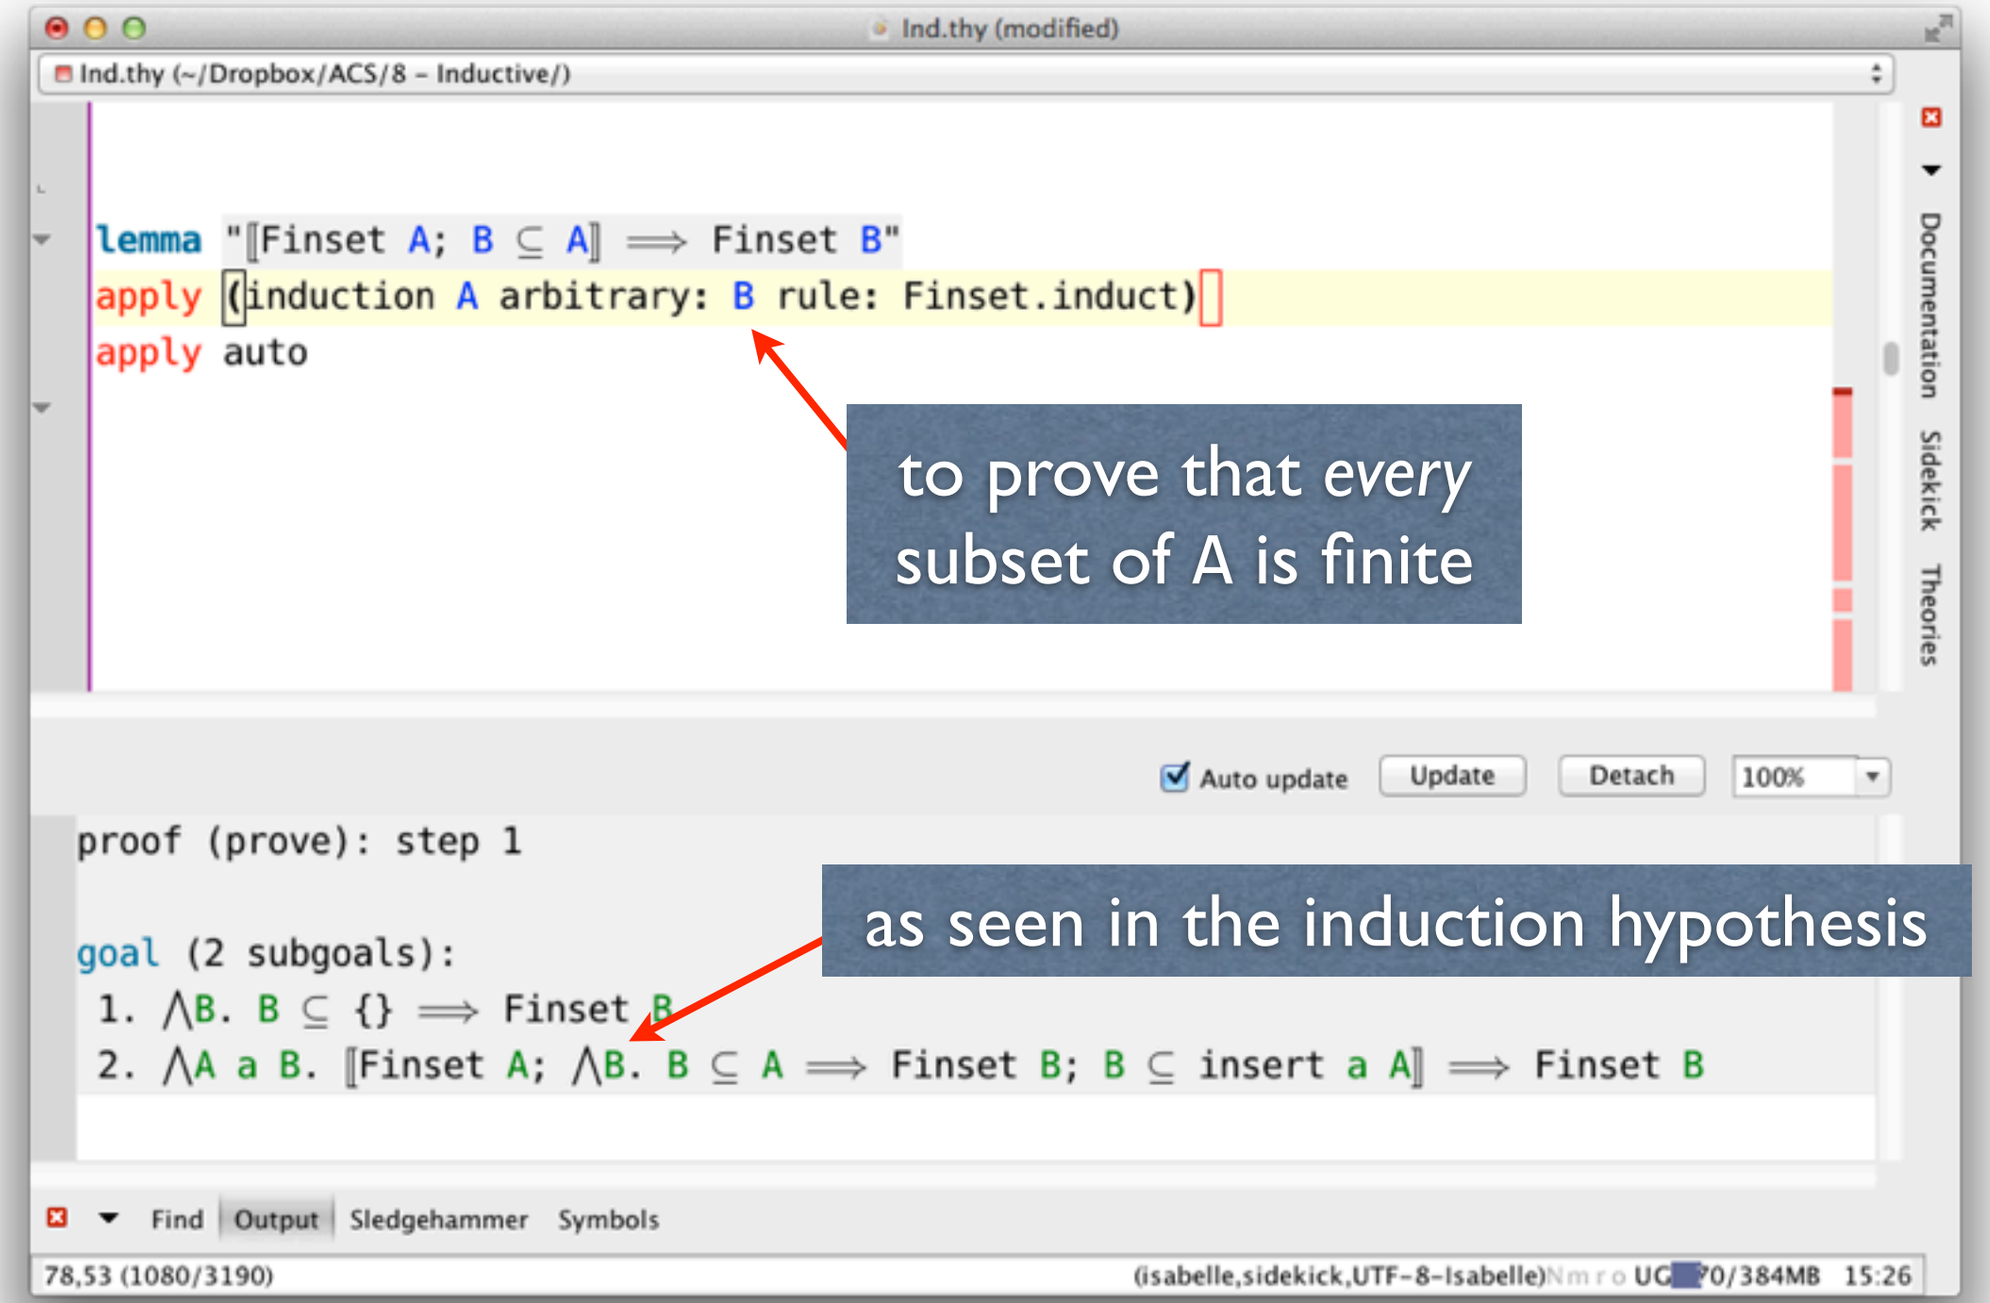

The proof is far more difficult than the preceding one, illustrating advanced techniques, in particular the sledgehammer tool.

A Critical Point in the Proof

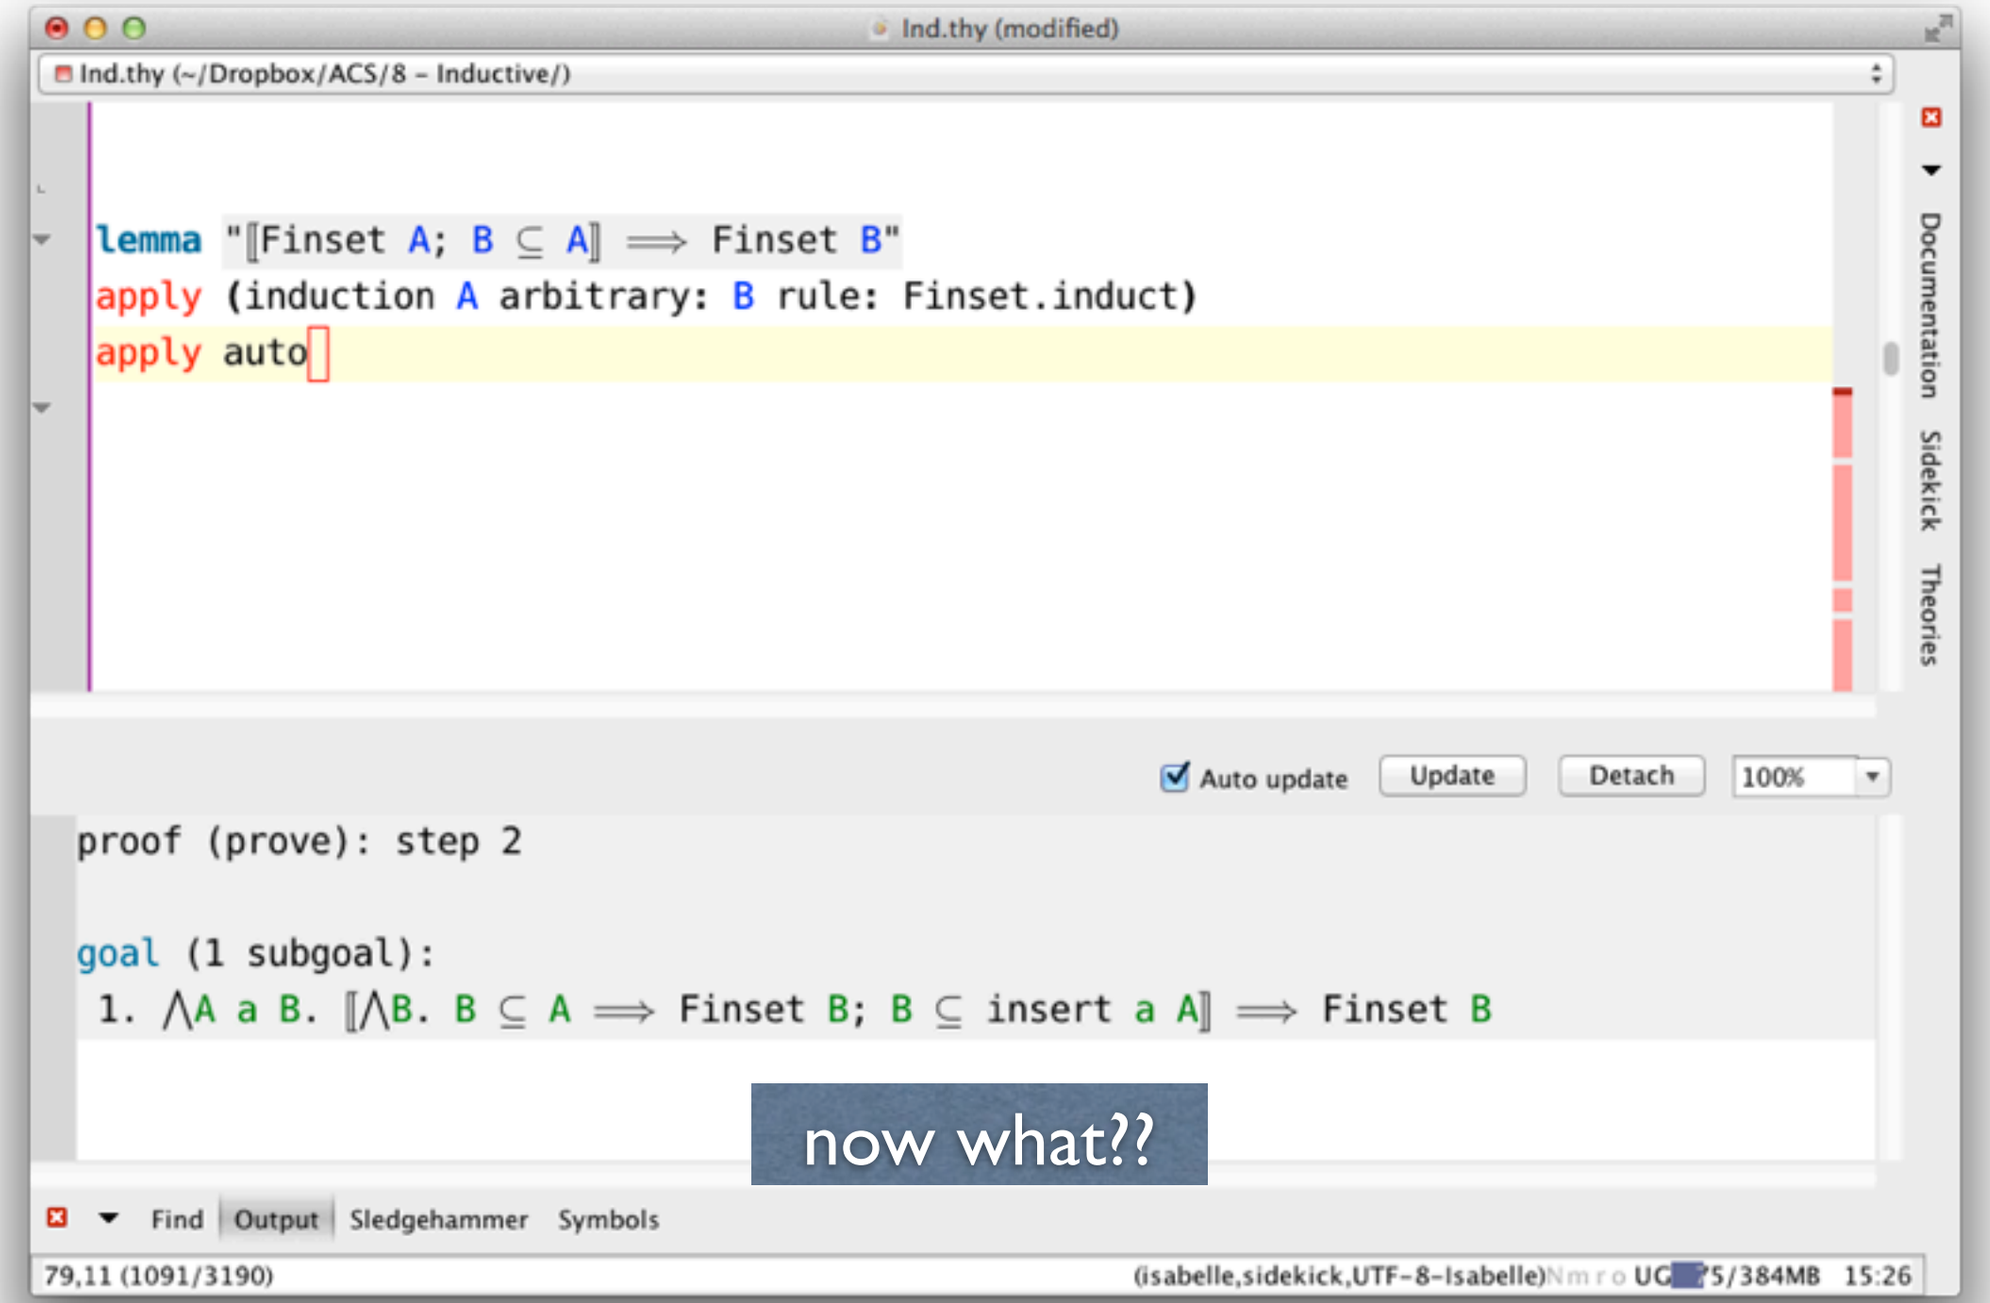

None of Isabelle's automatic proof methods (`auto`, `blast`, `force`) have any effect on this subgoal. Informally, we might consider case analysis on whether a∈B. This would require using proof tactics that have not been covered. Fortunately, Isabelle provides a general automated tool, sledgehammer.

Time to Try Sledgehammer!

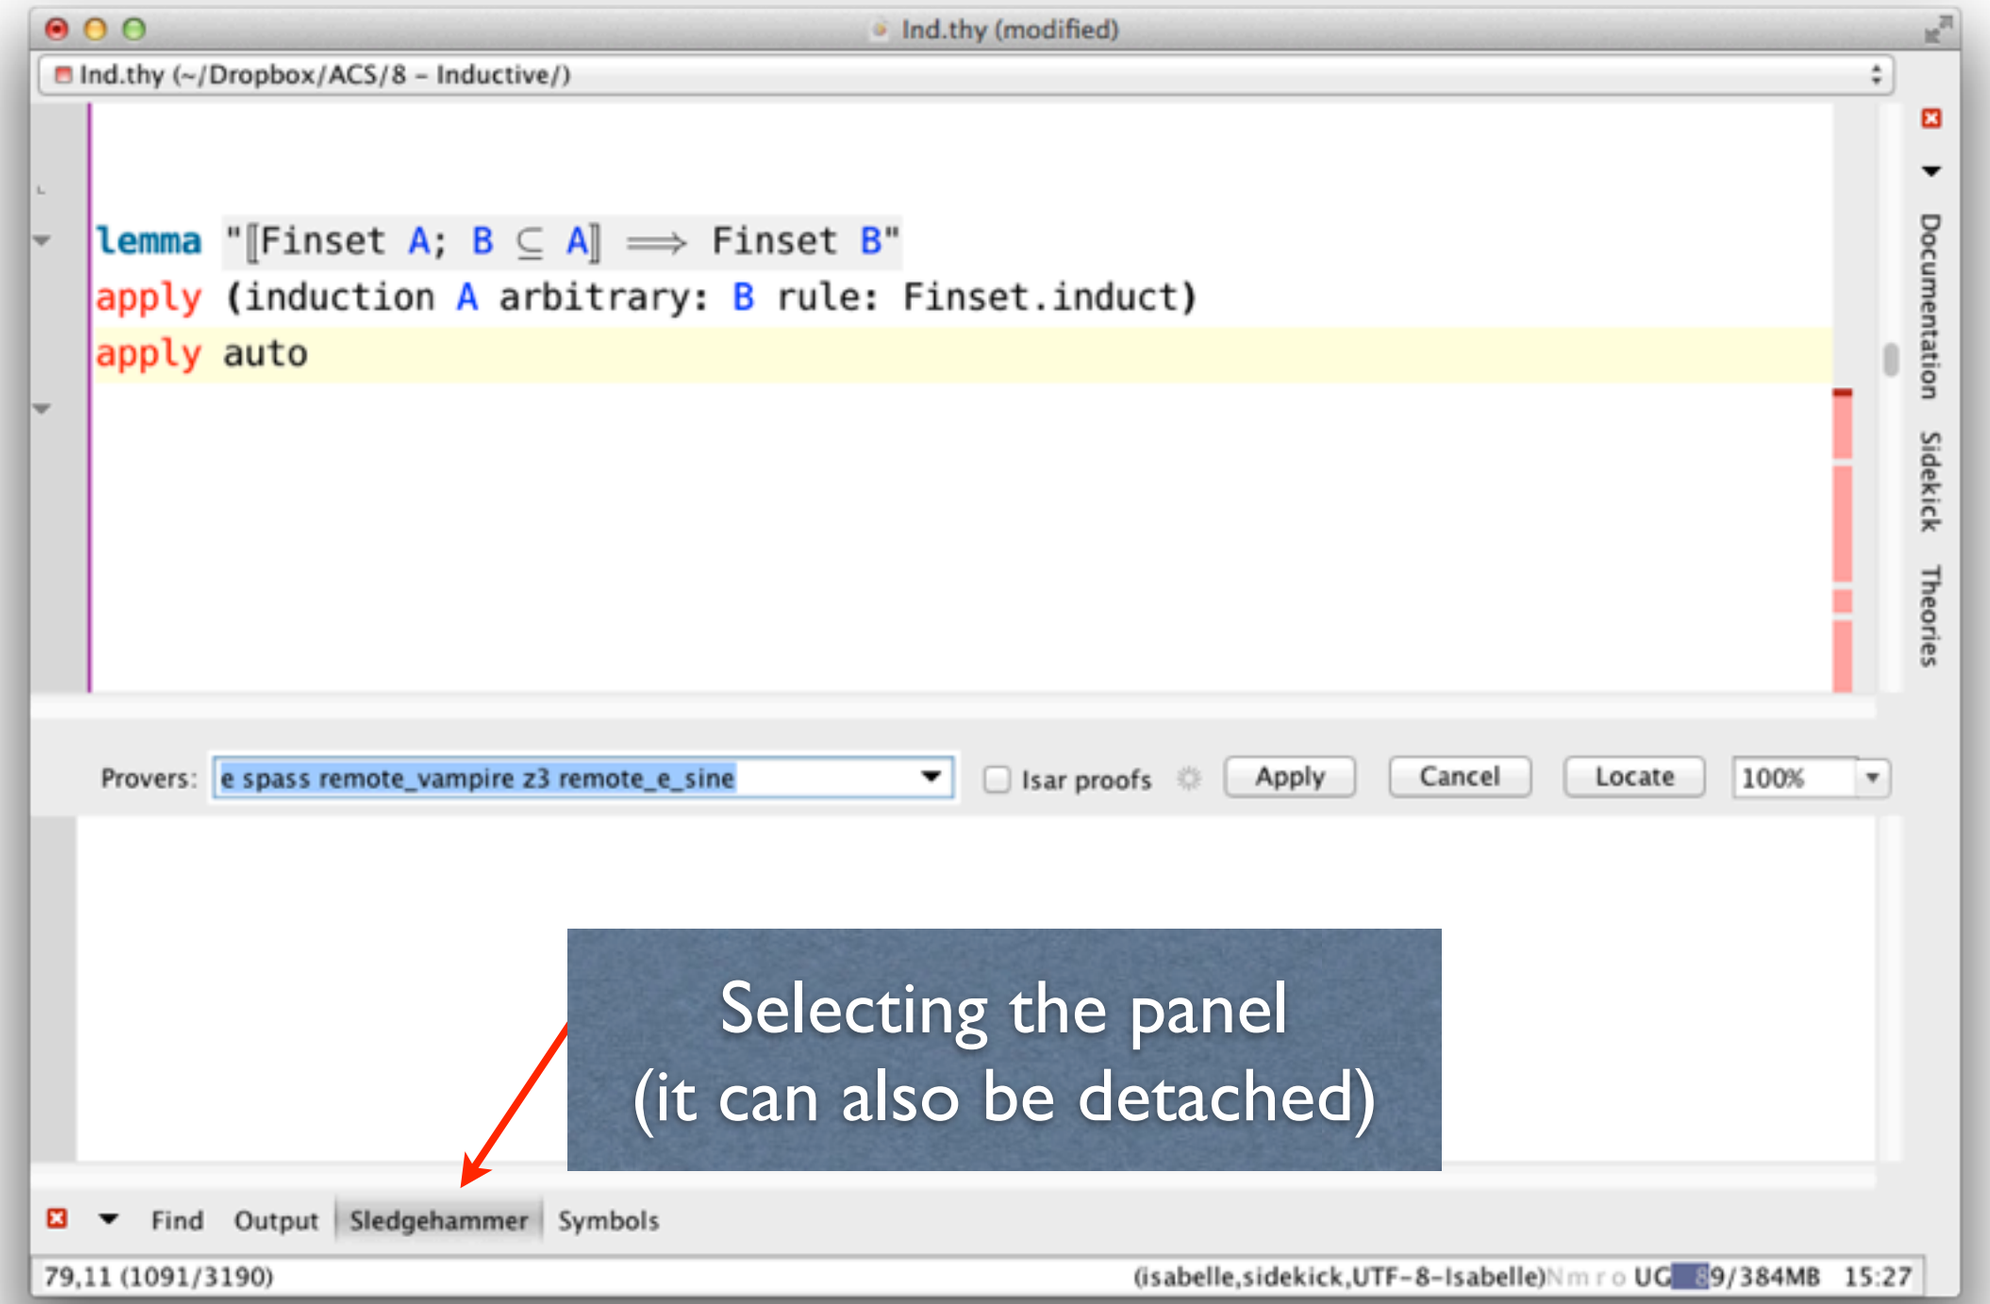

Sledgehammer calls several automated theorem provers in the background: in other words, Isabelle is still receptive to commands. You can continue to look for a proof manually.

Success!

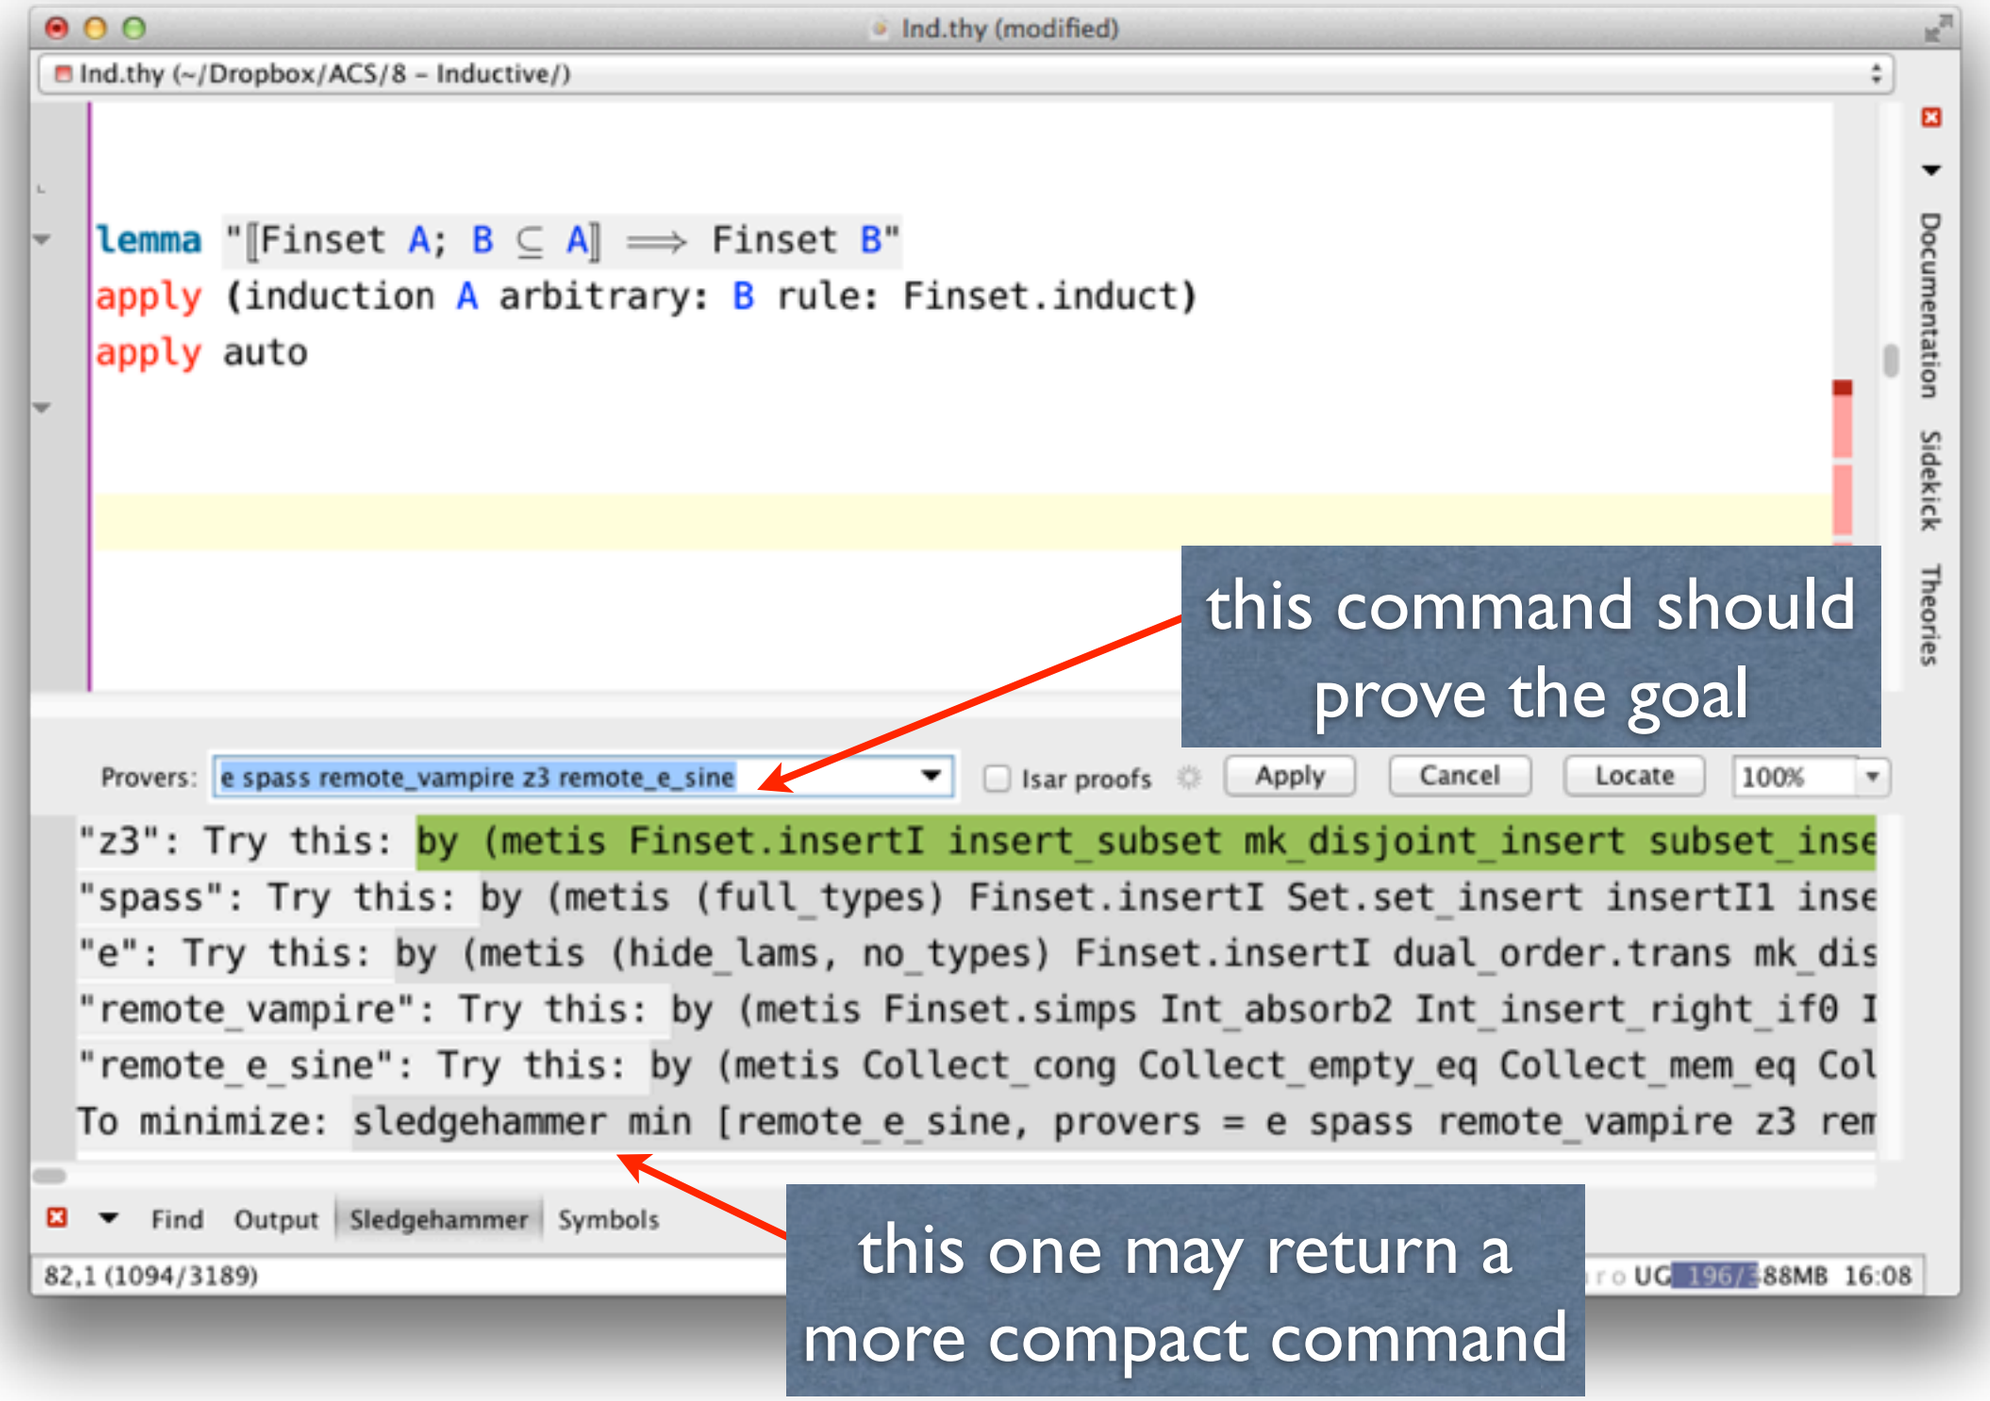

All of outputs are highlighted. They are live: clicking on either will insert that command into the proof script.

The Completed Proof

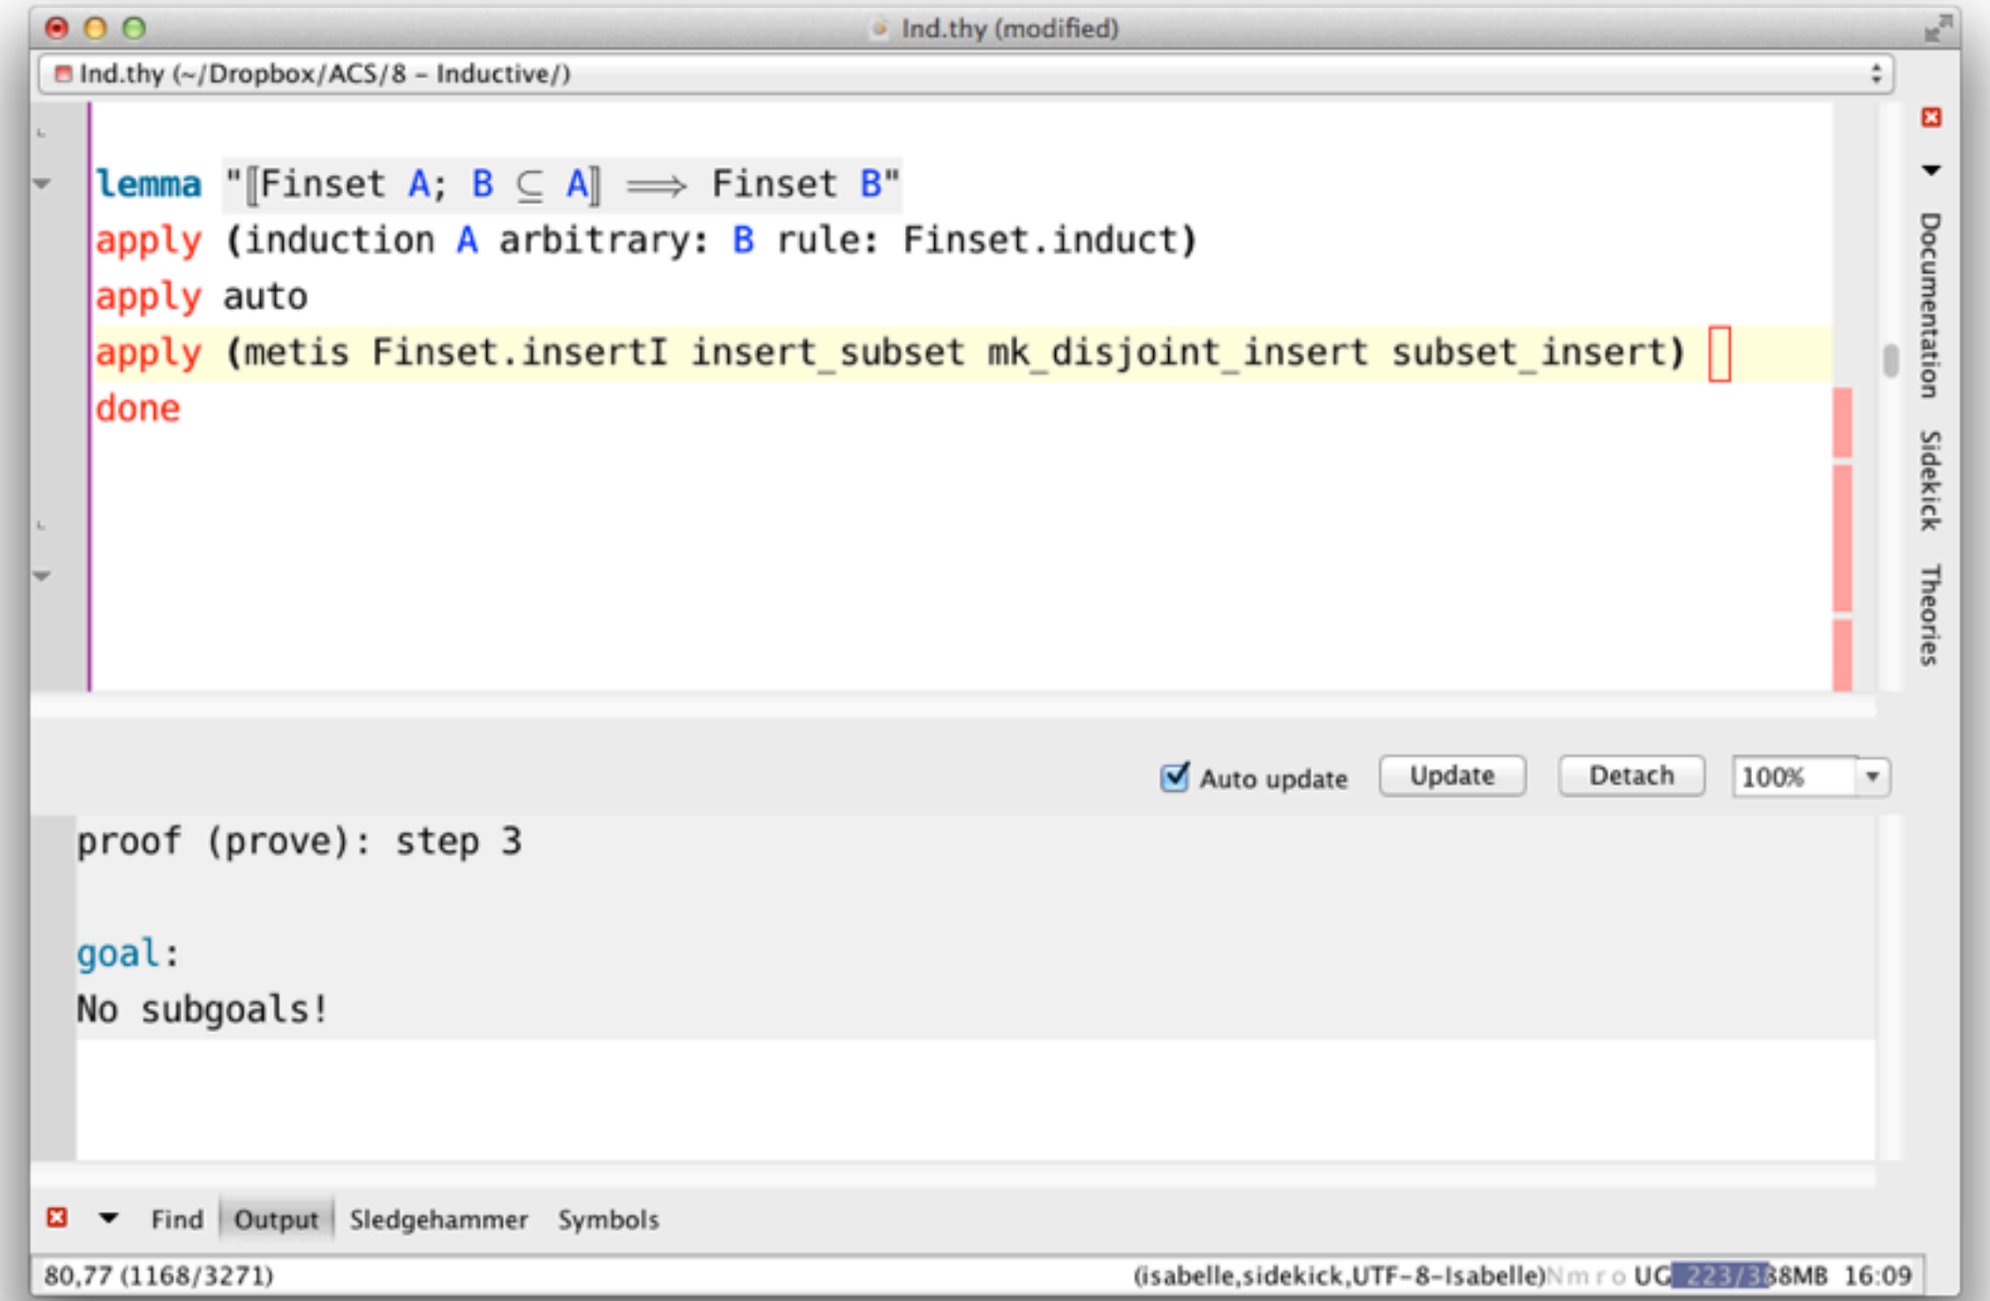

How Sledgehammer Works

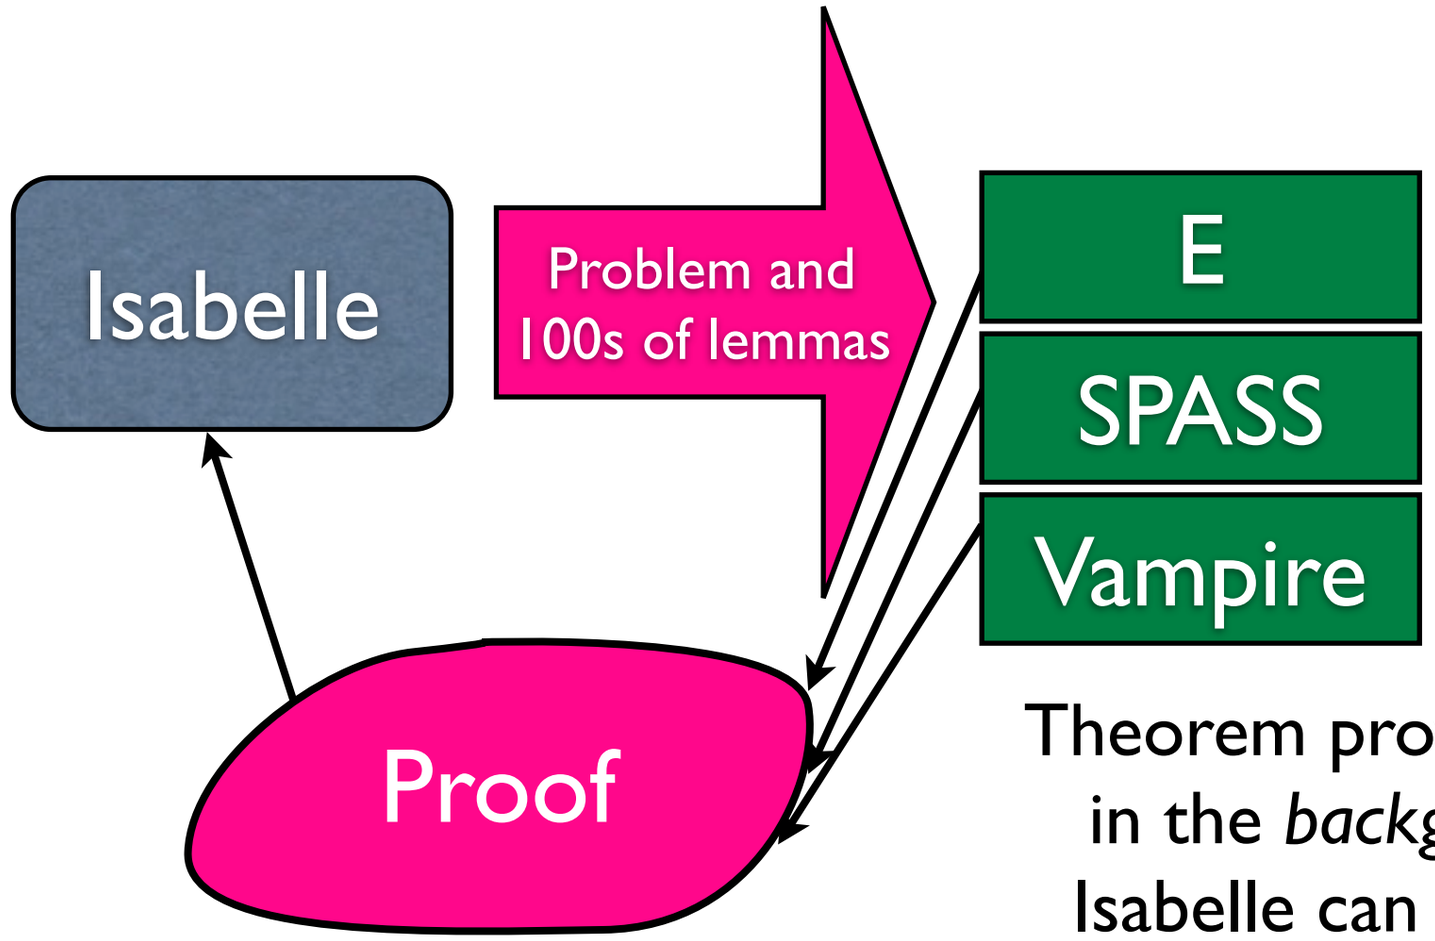

Theorem provers run in the *background*. Isabelle can still be used!

Notes on Sledgehammer

- It is always available, but it cannot work miracles.
- It does not prove the goal, but returns a call to `metis` or `smt`. This command sometimes runs slowly, and `smt` calls can be fragile.
- The minimise option removes redundant theorems, increasing the likelihood of success.
- Calling `metis` or `smt` directly is difficult unless you know exactly which lemmas are needed.

Metis is an automatic theorem prover for first order logic, written by Joe Hurd. Sledgehammer calls high-performance theorem provers, such as E and Vampire, using them as relevance filters to select from the thousands of lemmas available in Isabelle. Isabelle problems are translated for these automatic theorem provers using lightweight translations, which do not preserve soundness. For that reason, proofs found by those theorem provers may be incorrect. If that happens, the call to `metis` will generate an error message or fail to terminate. It is possible to force the use of sound translations, but sledgehammer seldom finds proofs using those.

Interactive Formal Verification
9: Operational Semantics

Lawrence C Paulson
Computer Laboratory
University of Cambridge

Overview

- The operational semantics of programming languages can be given *inductively*.
  - Type checking
  - Expression evaluation
  - Command execution, including concurrency
- Properties of the semantics are frequently proved by induction.
- Running example: an abstract language with WHILE

Language Syntax

```
typedecl loc -- "an unspecified type of locations"

type_synonym val   = nat -- "values"
type_synonym state = "loc => val"
type_synonym aexp  = "state => val"
type_synonym bexp  = "state => bool"   -- "functions on states"

datatype
  com = SKIP
      | Assign loc aexp          ("_ :== _ " 60)
      | Semi   com com           ("_; _"  [60, 60] 10)
      | Cond   bexp com com      ("IF _ THEN _ ELSE _"  60)
      | While  bexp com          ("WHILE _ DO _"  60)
```

Arithmetic & boolean expressions are *functions* over the state

For simplicity, this example does not specify arithmetic or boolean expressions in any detail. Although this approach is unrealistic, it allows us to illustrate key aspects of formalised proofs about programming language semantics.

A "Big-Step" Semantics

$$\langle \mathbf{skip}, s\rangle \to s \qquad\qquad \langle x := a, s\rangle \to s[x := a\ s]$$

$$\frac{\langle c_0, s\rangle \to s'' \qquad \langle c_1, s''\rangle \to s'}{\langle c_0; c_1, s\rangle \to s'}$$

$$\frac{b\,s \qquad \langle c_0, s\rangle \to s'}{\langle \mathbf{if}\ b\ \mathbf{then}\ c_0\ \mathbf{else}\ c_1, s\rangle \to s'} \qquad \frac{\neg\, b\,s \qquad \langle c_1, s\rangle \to s'}{\langle \mathbf{if}\ b\ \mathbf{then}\ c_0\ \mathbf{else}\ c_1, s\rangle \to s'}$$

$$\frac{\neg\, b\,s}{\langle \mathbf{while}\ b\ \mathbf{do}\ c, s\rangle \to s} \qquad \frac{b\,s \quad \langle c, s\rangle \to s'' \quad \langle \mathbf{while}\ b\ \mathbf{do}\ c, s''\rangle \to s'}{\langle \mathbf{while}\ b\ \mathbf{do}\ c, s\rangle \to s'}$$

In a big step semantics, the transition ⟨*c*,*s*⟩ → *s'* means, executing the command *c* starting in the state *s* can terminate in state *s'*.

Formalised Language Semantics

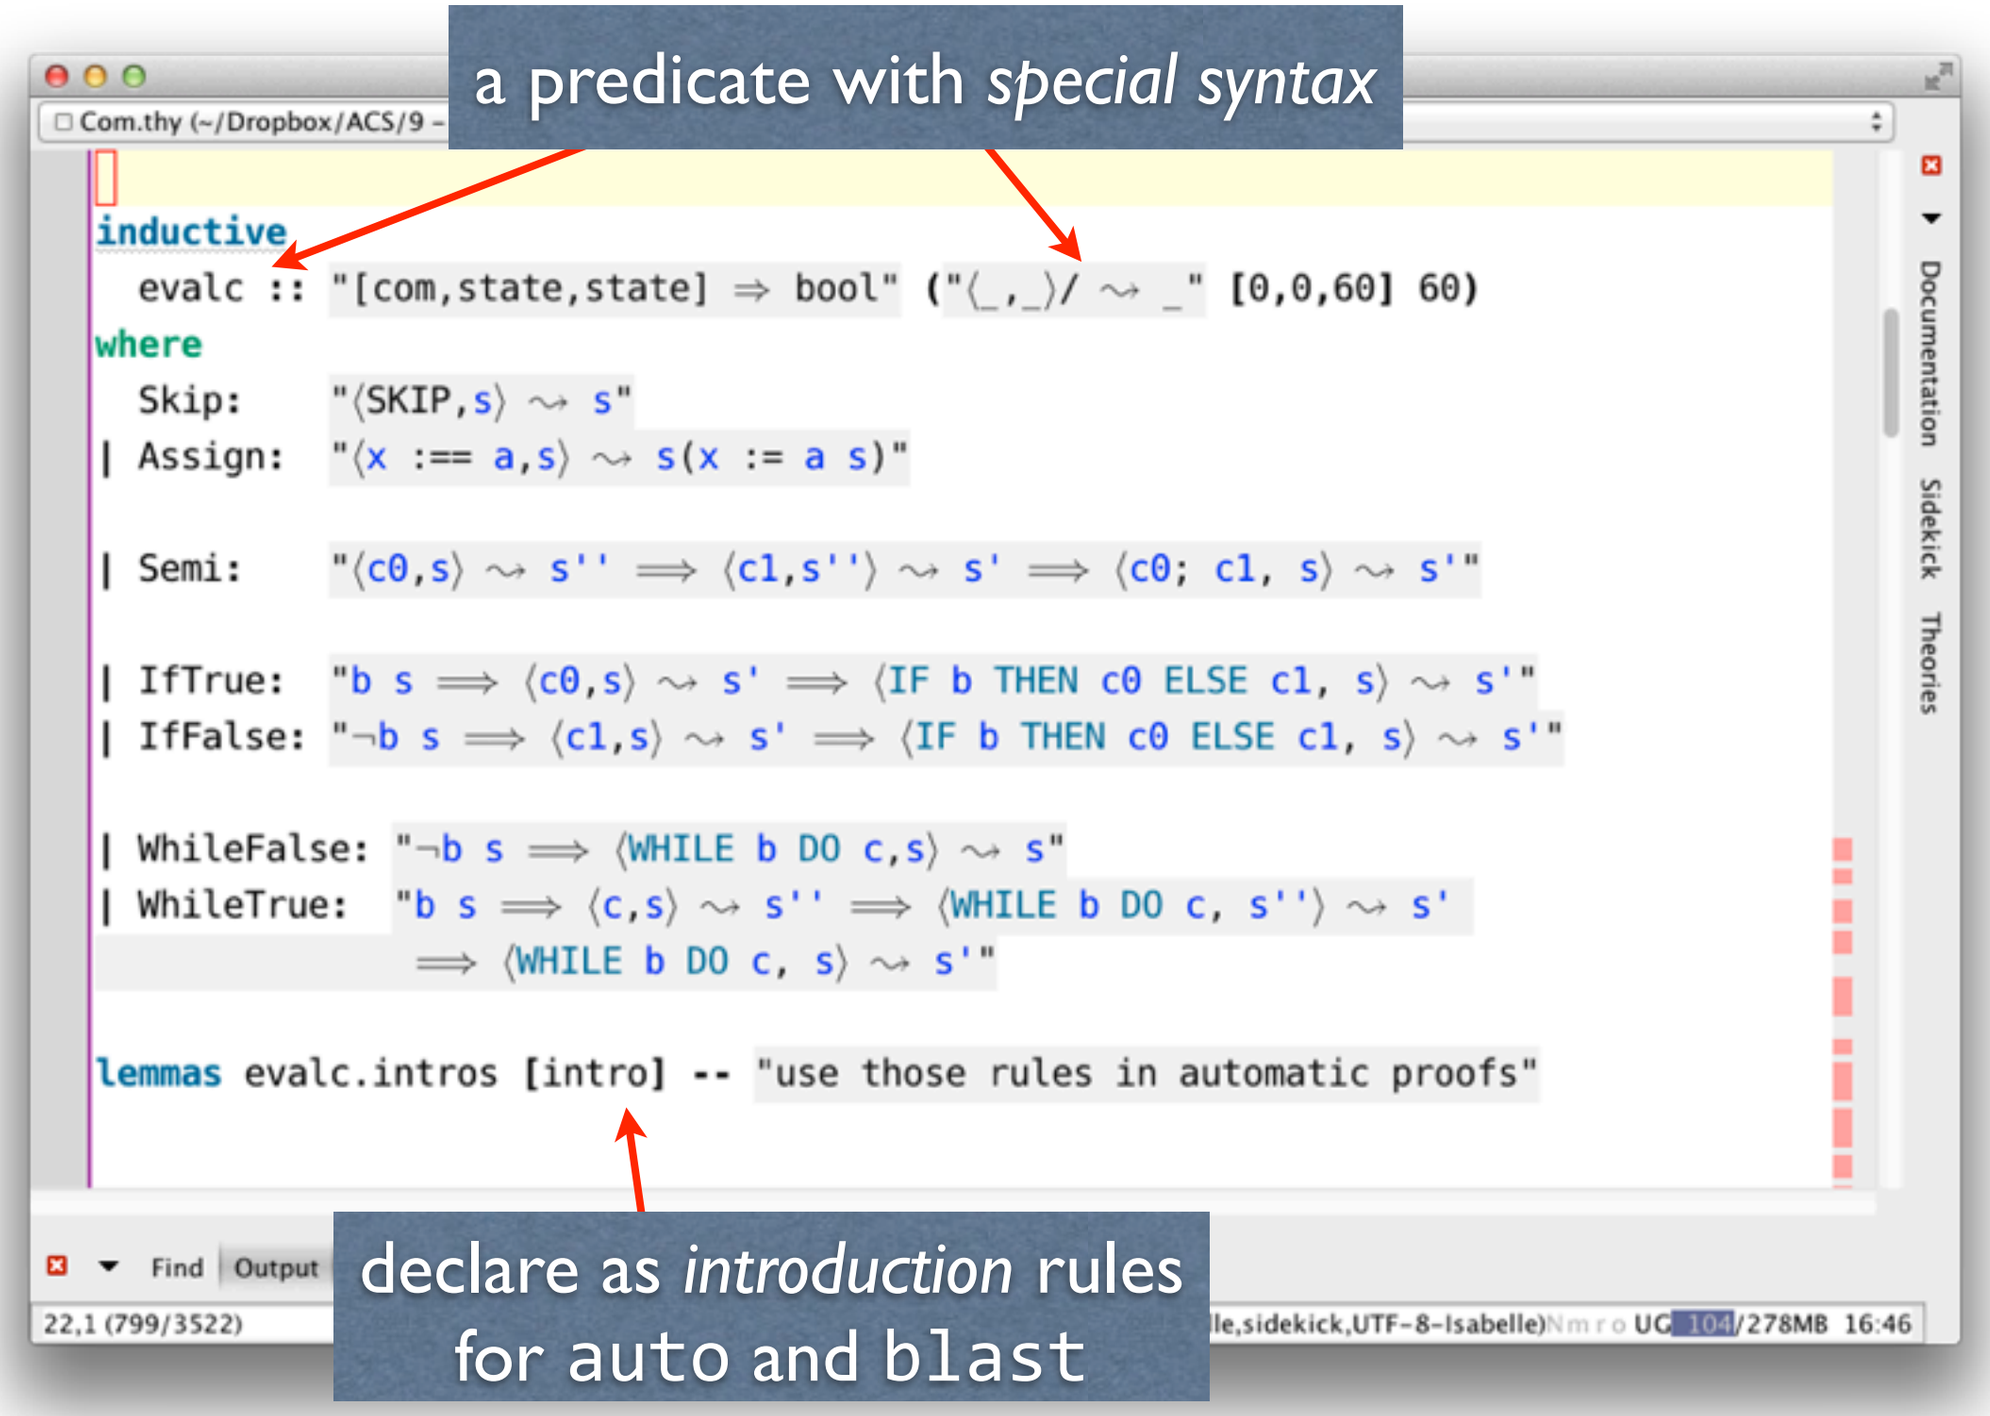

In the previous lecture, we used a related declaration, `inductive_set`. Note that there is no real difference between a set and a predicate of one argument. However, formal semantics generally requires a predicate three or four arguments, and the corresponding set of triples is a little more difficult to work with. Attaching special syntax, as shown above, also requires the use of a predicate. Therefore, formalised semantic definitions will generally use `inductive`.

Rule Inversion

- When $\langle \textbf{skip}, s \rangle \rightarrow s'$ we know $s = s'$
- When $\langle \textbf{if } b \textbf{ then } c_0 \textbf{ else } c_1, s \rangle \rightarrow s'$ we know
  - b and $\langle c_0, s \rangle \rightarrow s'$, or...
  - $\neg b$ and $\langle c_1, s \rangle \rightarrow s'$
- This sort of case analysis is easy in Isabelle.

Rule inversion refers to case analysis on the form of the induction, matching the conclusions of the introduction rules (those making up the inductive definition) with a particular pattern. It is useful when only a small percentage of the introduction rules can match the pattern. This type of reasoning is extremely common in informal proofs about operational semantics. It would not be useful in the inductive definitions covered in the previous lecture, where the conclusions of the rules had little structure.

Rule Inversion in Isabelle

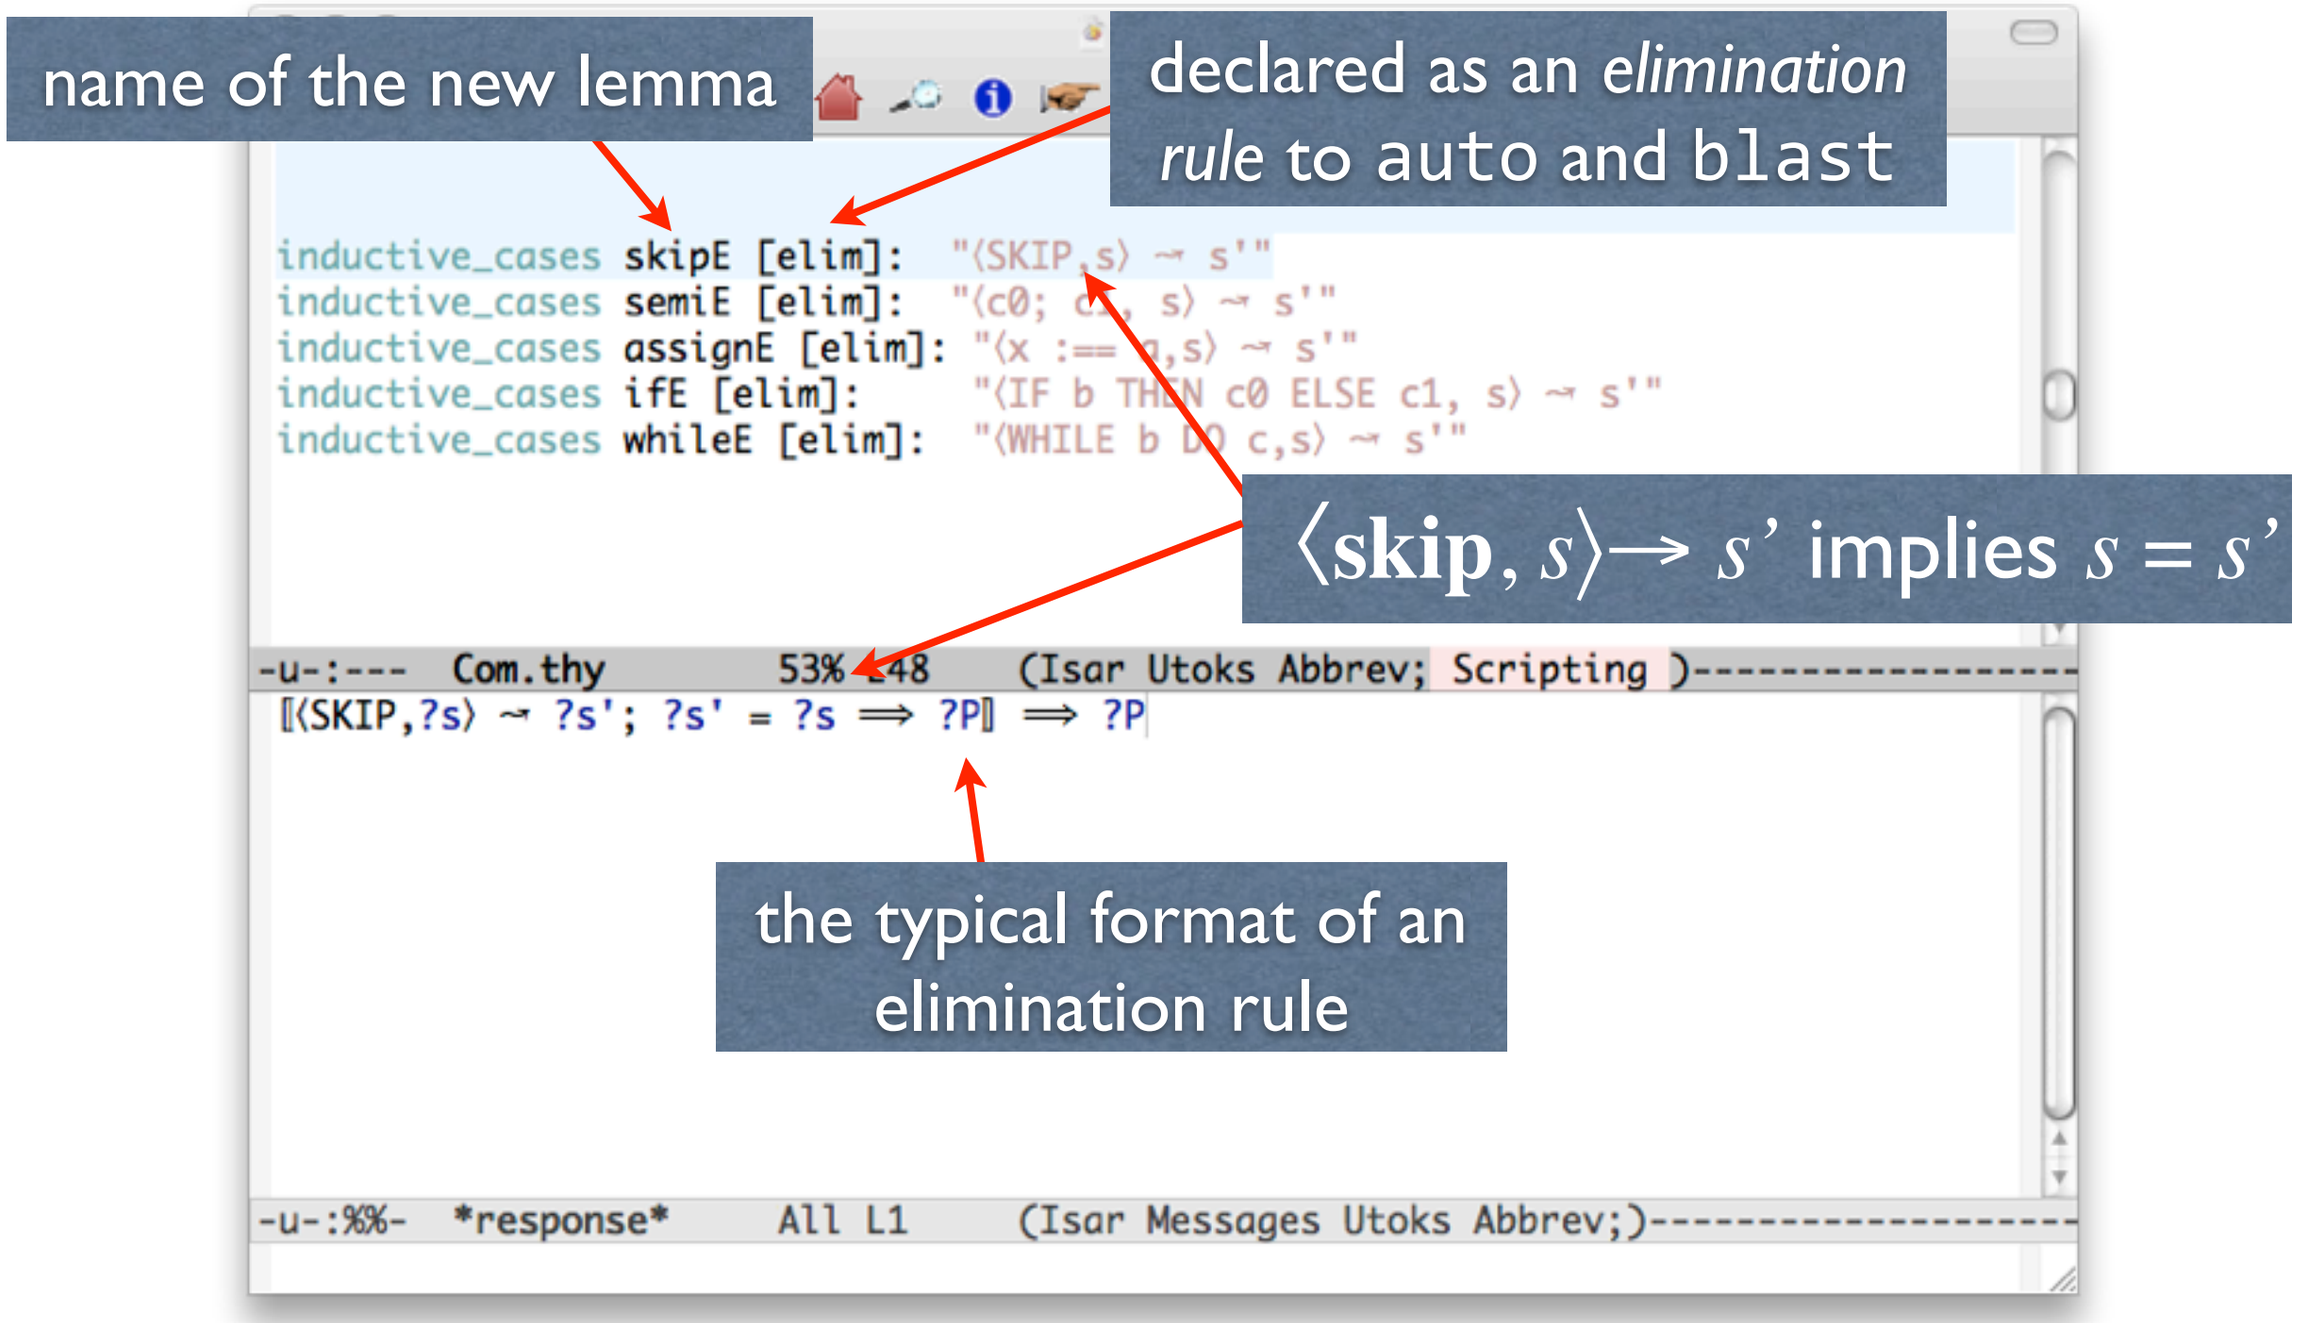

The pattern for each rule inversion lemma appears in quotation marks. Isabelle generates a theorem and gives it the name shown. Each theorem is also made available to Isabelle's automatic tools.

It is possible to write elim! rather than just elim; the exclamation mark tells Isabelle to apply the lemma aggressively. However, this must not be done with the theorem `whileE`: it expands an occurrence of ⟨**while** *b* **do** c, s⟩ → s' and generates another formula of essentially the same form, thereby running for ever.

Rule Inversion Again

Com.thy

```
inductive_cases skipE [elim]:  "⟨SKIP,s⟩ ⇝ s'"
inductive_cases semiE [elim]:  "⟨c0; c1, s⟩ ⇝ s'"
inductive_cases assignE [elim]: "⟨x :== a,s⟩ ⇝ s'"
inductive_cases ifE [elim]:     "⟨IF b THEN c0 ELSE c1, s⟩ ⇝ s'"
inductive_cases whileE [elim]:  "⟨WHILE b DO c,s⟩ ⇝ s'"
```

-u-:--- Com.thy 53% L49 (Isar Utoks Abbrev; Scripting)----------

```
⟦⟨?c0.0; ?c1.0,?s⟩ ⇝ ?s'; ⋀s''. ⟦⟨?c0.0,?s⟩ ⇝ s''; ⟨?c1.0,s''⟩ ⇝ ?s'⟧ ⟹ ?P⟧
⟹ ?P
```

expresses the existence of the *intermediate* state, s'

-u-:%%- *response* All L2 (Isar Messages Utoks Abbrev;)----------

A Non-Termination Proof

$$\langle \mathbf{while\ true\ do}\ c, s\rangle \not\rightarrow s'$$

This formula is not provable by induction!

$$\langle c, s\rangle \rightarrow s' \Longrightarrow \forall c'.\, c \neq (\mathbf{while\ true\ do}\ c')$$

The inductive version considers
all possible commands

Non-Termination in Isabelle

7 subgoals, one for each rule of the definition

Most are trivial, by distinctness

trivial for another reason

Com.thy (modified)

Com.thy (~/Dropbox/ACS/9 – Operational Semantics/)

```
lemma while_never: "⟨c, s⟩ ⇝ u ⟹ c ≠ WHILE (λs. True) DO c1"
apply (induct rule: evalc.induct)
apply auto
```

Update Detach 100%

```
1. ⋀s. SKIP ≠ WHILE λs. True DO c1
2. ⋀x a s. x :== a ≠ WHILE λs. True DO c1
3. ⋀c0 s s'' c1a s'.
      ⟦⟨c0,s⟩ ⇝ s''; c0 ≠ WHILE λs. True DO c1; ⟨c1a,s''⟩ ⇝ s'; c1a ≠ WHILE λs. True DO c1⟧
      ⟹ c0 ; c1a ≠ WHILE λs. True DO c1
4. ⋀b s c0 s' c1a.
      ⟦b s; ⟨c0,s⟩ ⇝ s'; c0 ≠ WHILE λs. True DO c1⟧
      ⟹ IF b THEN c0 ELSE c1a ≠ WHILE λs. True DO c1
5. ⋀b s c1a s' c0.
      ⟦¬ b s; ⟨c1a,s⟩ ⇝ s'; c1a ≠ WHILE λs. True DO c1⟧
      ⟹ IF b THEN c0 ELSE c1a ≠ WHILE λs. True DO c1
6. ⋀b s c. ¬ b s ⟹ WHILE b DO c ≠ WHILE λs. True DO c1
7. ⋀b s c s'' s'.
      ⟦b s; ⟨c,s⟩ ⇝ s''; c ≠ WHILE λs. True DO c1; ⟨WHILE b DO c,s''⟩ ⇝ s';
       WHILE b DO c ≠ WHILE λs. True DO c1⟧
      ⟹ WHILE b DO c ≠ WHILE λs. True DO c1
```

Documentation Sidekick Theories

Find Output Sledgehammer Symbols

58,34 (1831/3537) (isabelle,sidekick,UTF-8-Isabelle)Nmro UG 128/256MB 16:54

Done!

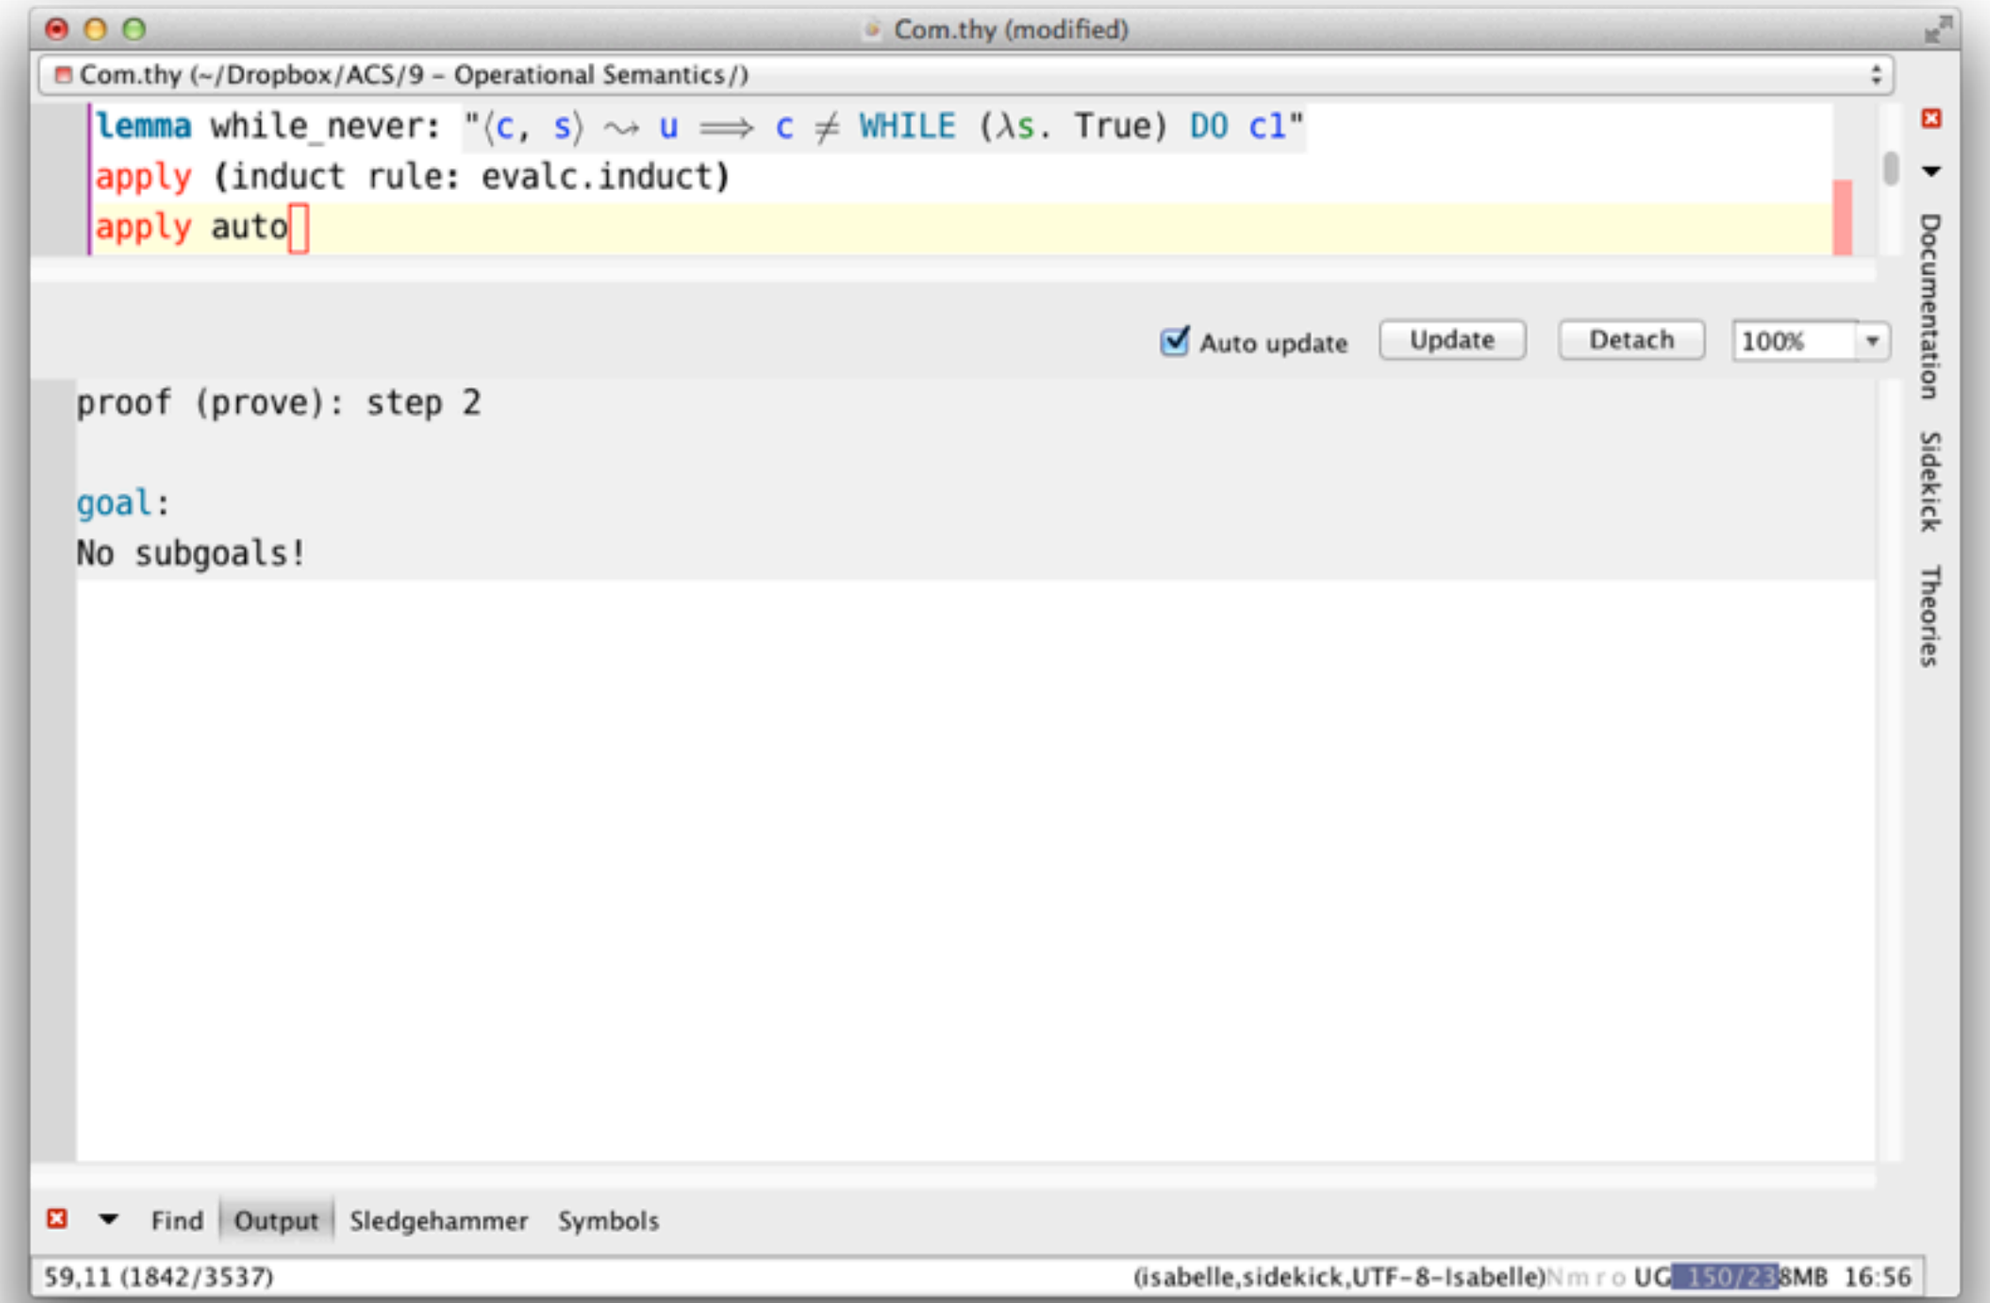

This really is a trivial proof. I timed this call to `auto` and it needed only 6 ms.

Determinacy

$$\frac{\langle c, s\rangle \to t \qquad \langle c, s\rangle \to u}{t = u}$$

If a command is executed in a given state, and it terminates, then this final state is *unique*.

Determinacy in Isabelle...

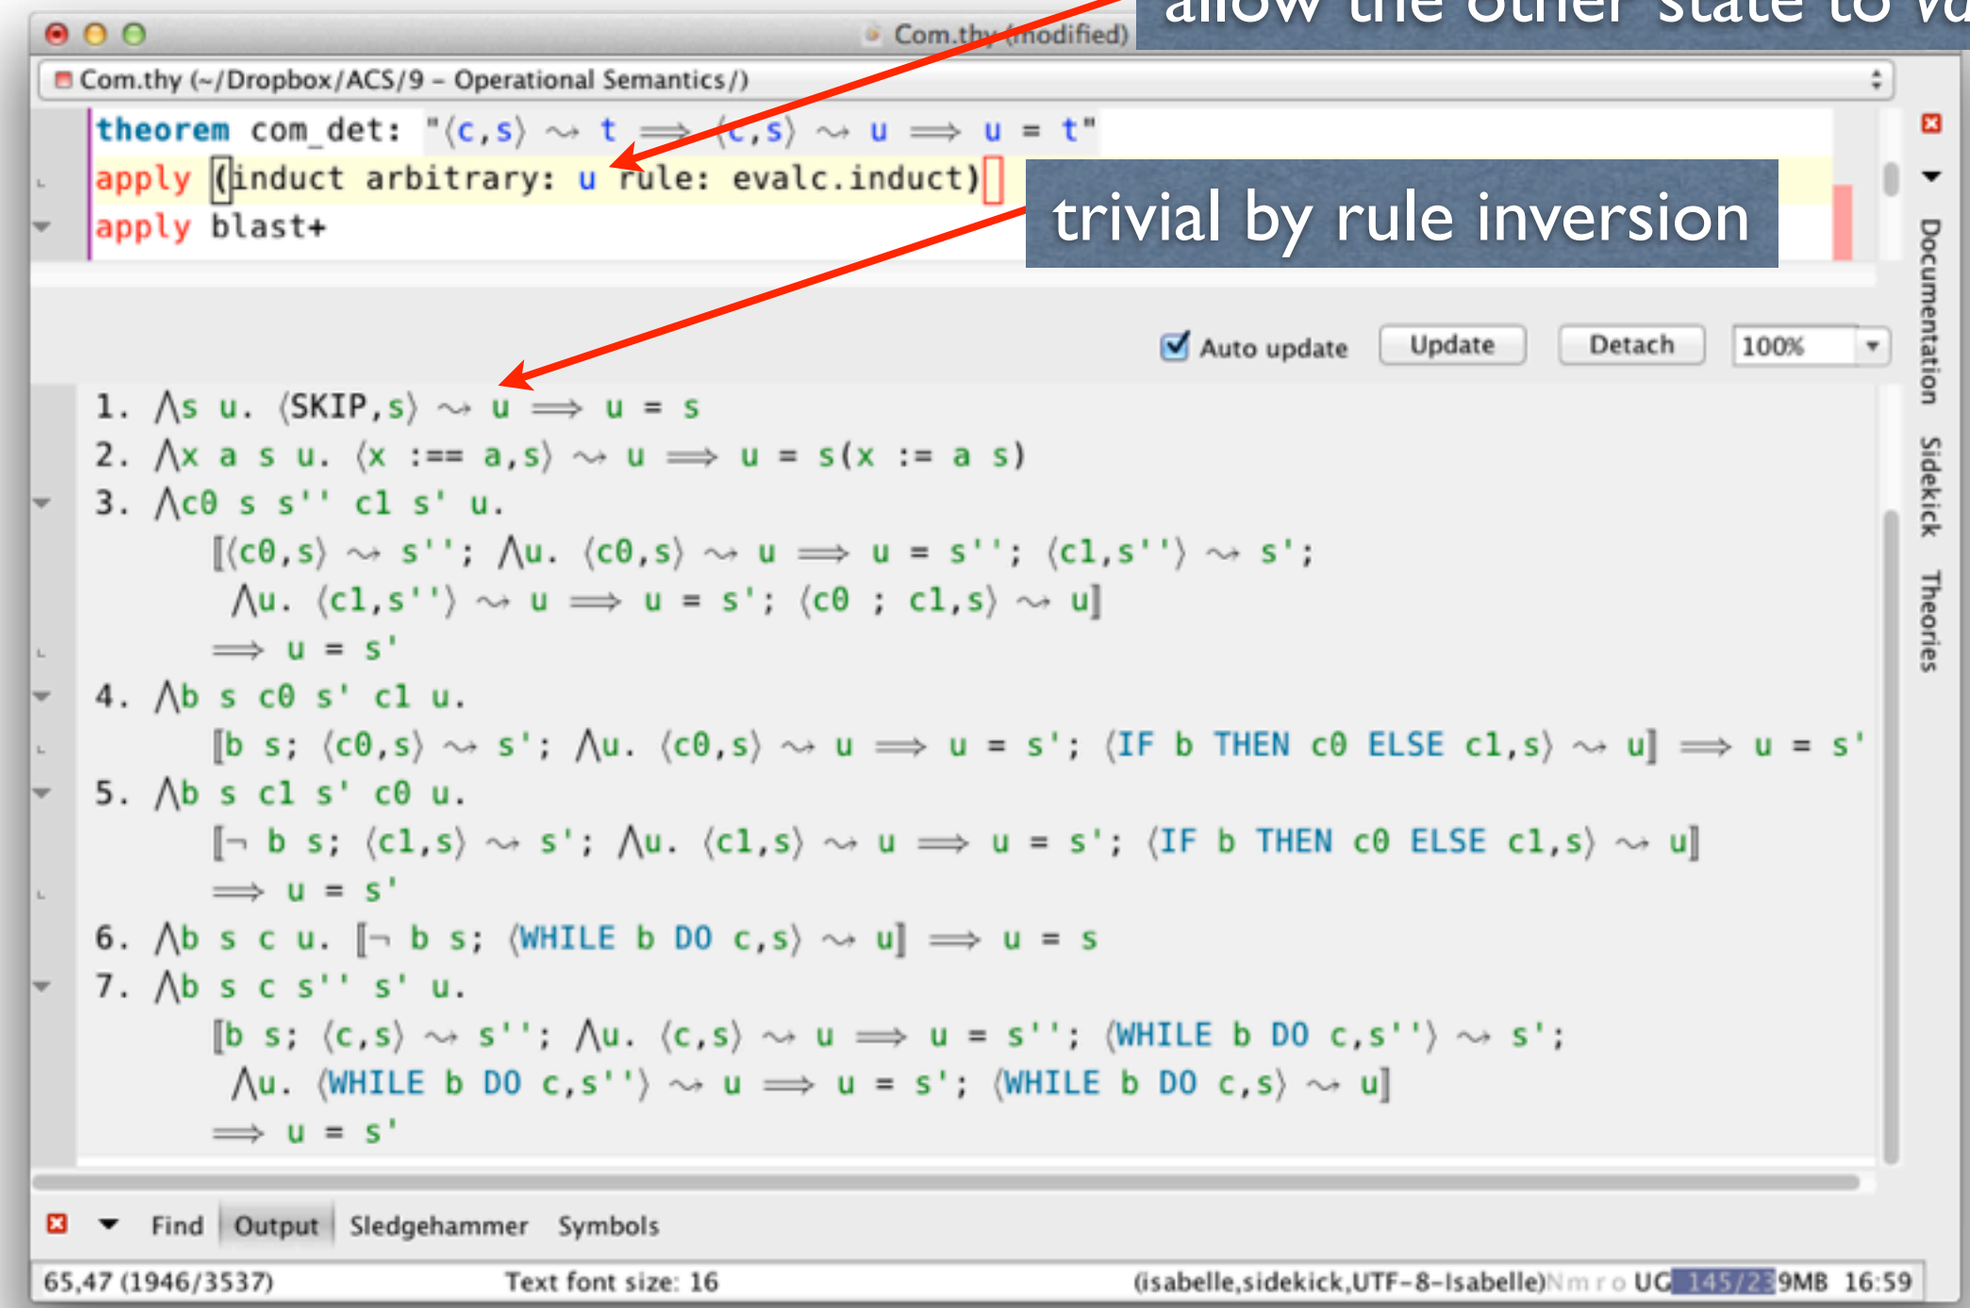

The proof method `blast` uses introduction and elimination rules, combined with powerful search heuristics. It will not terminate until it has solved the goal. Unlike `auto` and `force`, it does not perform simplification (rewriting) or arithmetic reasoning. These subgoals are mostly trivial: rule inversion, which we set up previously, expresses precisely what we need: that if the given commands have executed, then corresponding intermediate states have been reached. The induction hypothesis allow us to assume the determinacy of the sub-commands.

Proved by Rule Inversion

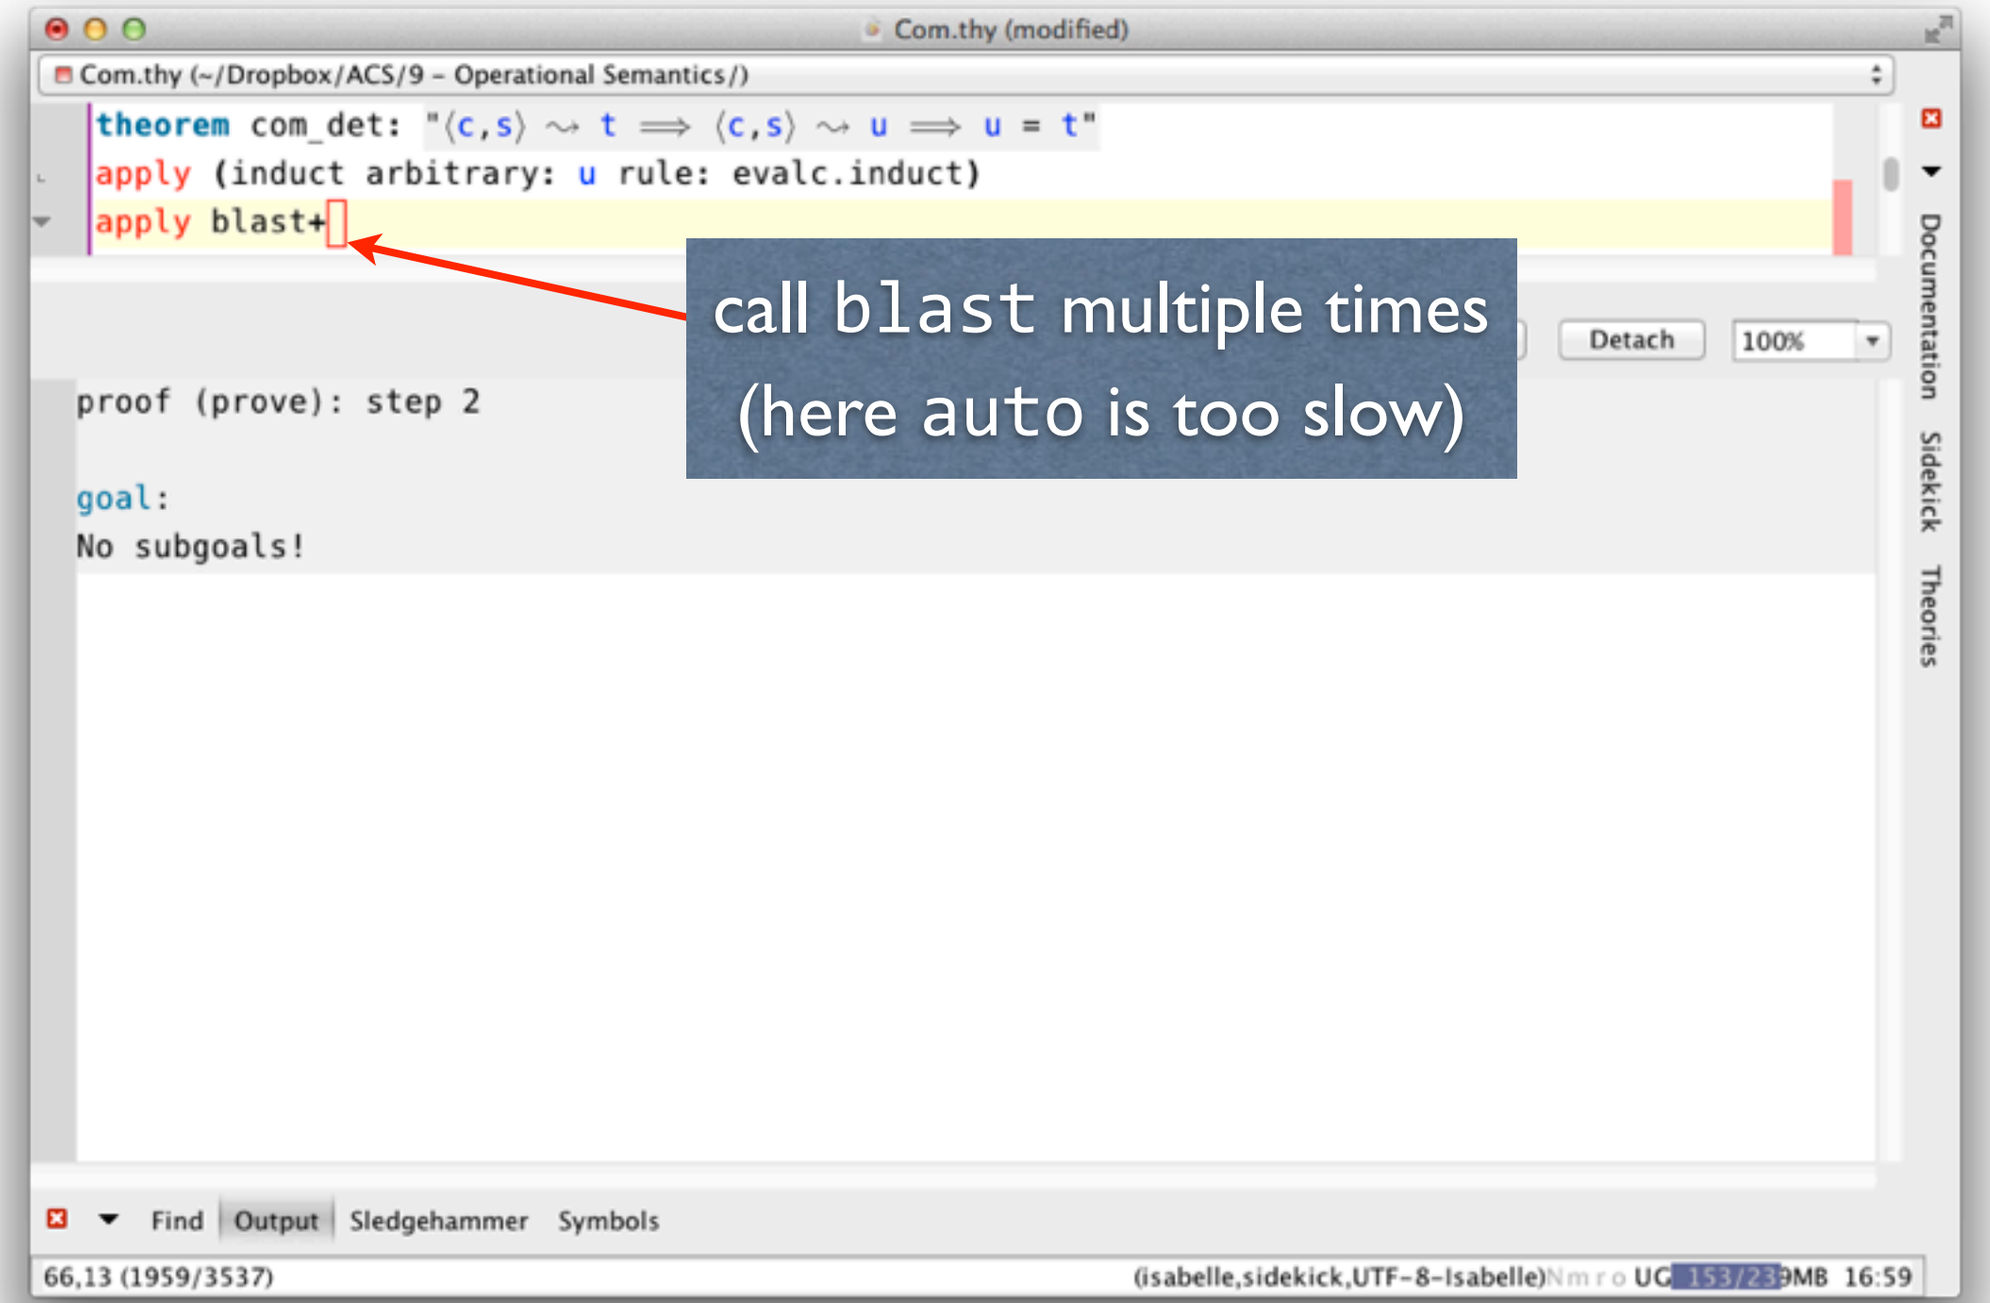

The proof involves a long, tedious and detailed series of rule inversions. Apart from its length, the proof is trivial. This proof needed only 32 ms.

Semantic Equivalence

We can even define the infix syntax

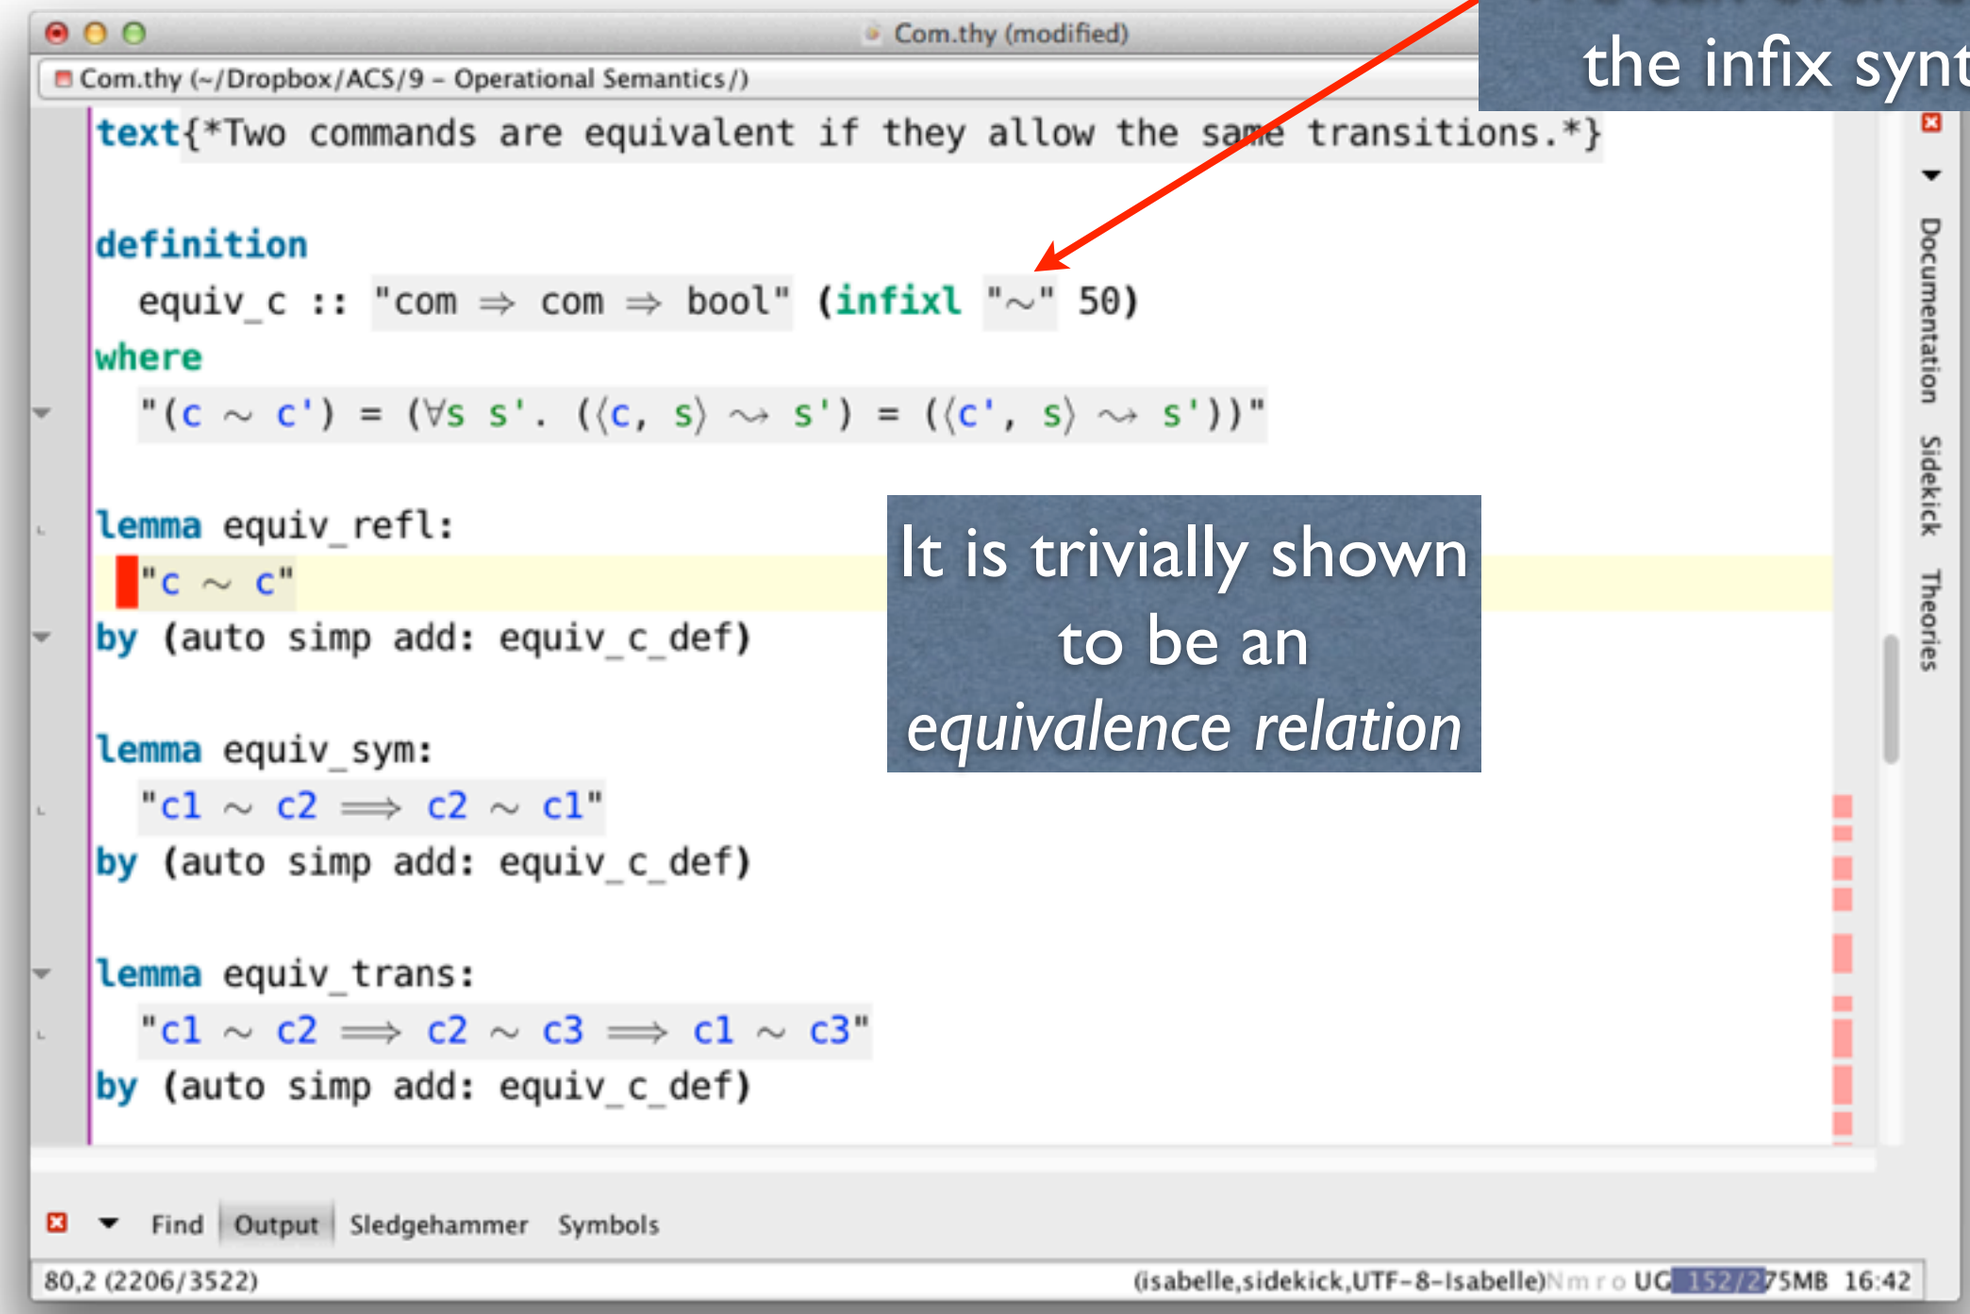

It is trivially shown to be an *equivalence relation*

The printed version of these notes does not include the actual proofs, because they are revealed during the presentation. They are reproduced below. It is necessary to unfold the definition of semantic equivalence, `equiv_c`. By default, Isabelle does not unfold nonrecursive definitions.

```
lemma equiv_refl:
  "c ~ c"
by (auto simp add: equiv_c_def)

lemma equiv_sym:
  "c1 ~ c2 ==> c2 ~ c1"
by (auto simp add: equiv_c_def)

lemma equiv_trans:
  "c1 ~ c2 ==> c2 ~ c3 ==> c1 ~ c3"
by (auto simp add: equiv_c_def)
```

More Semantic Equivalence!

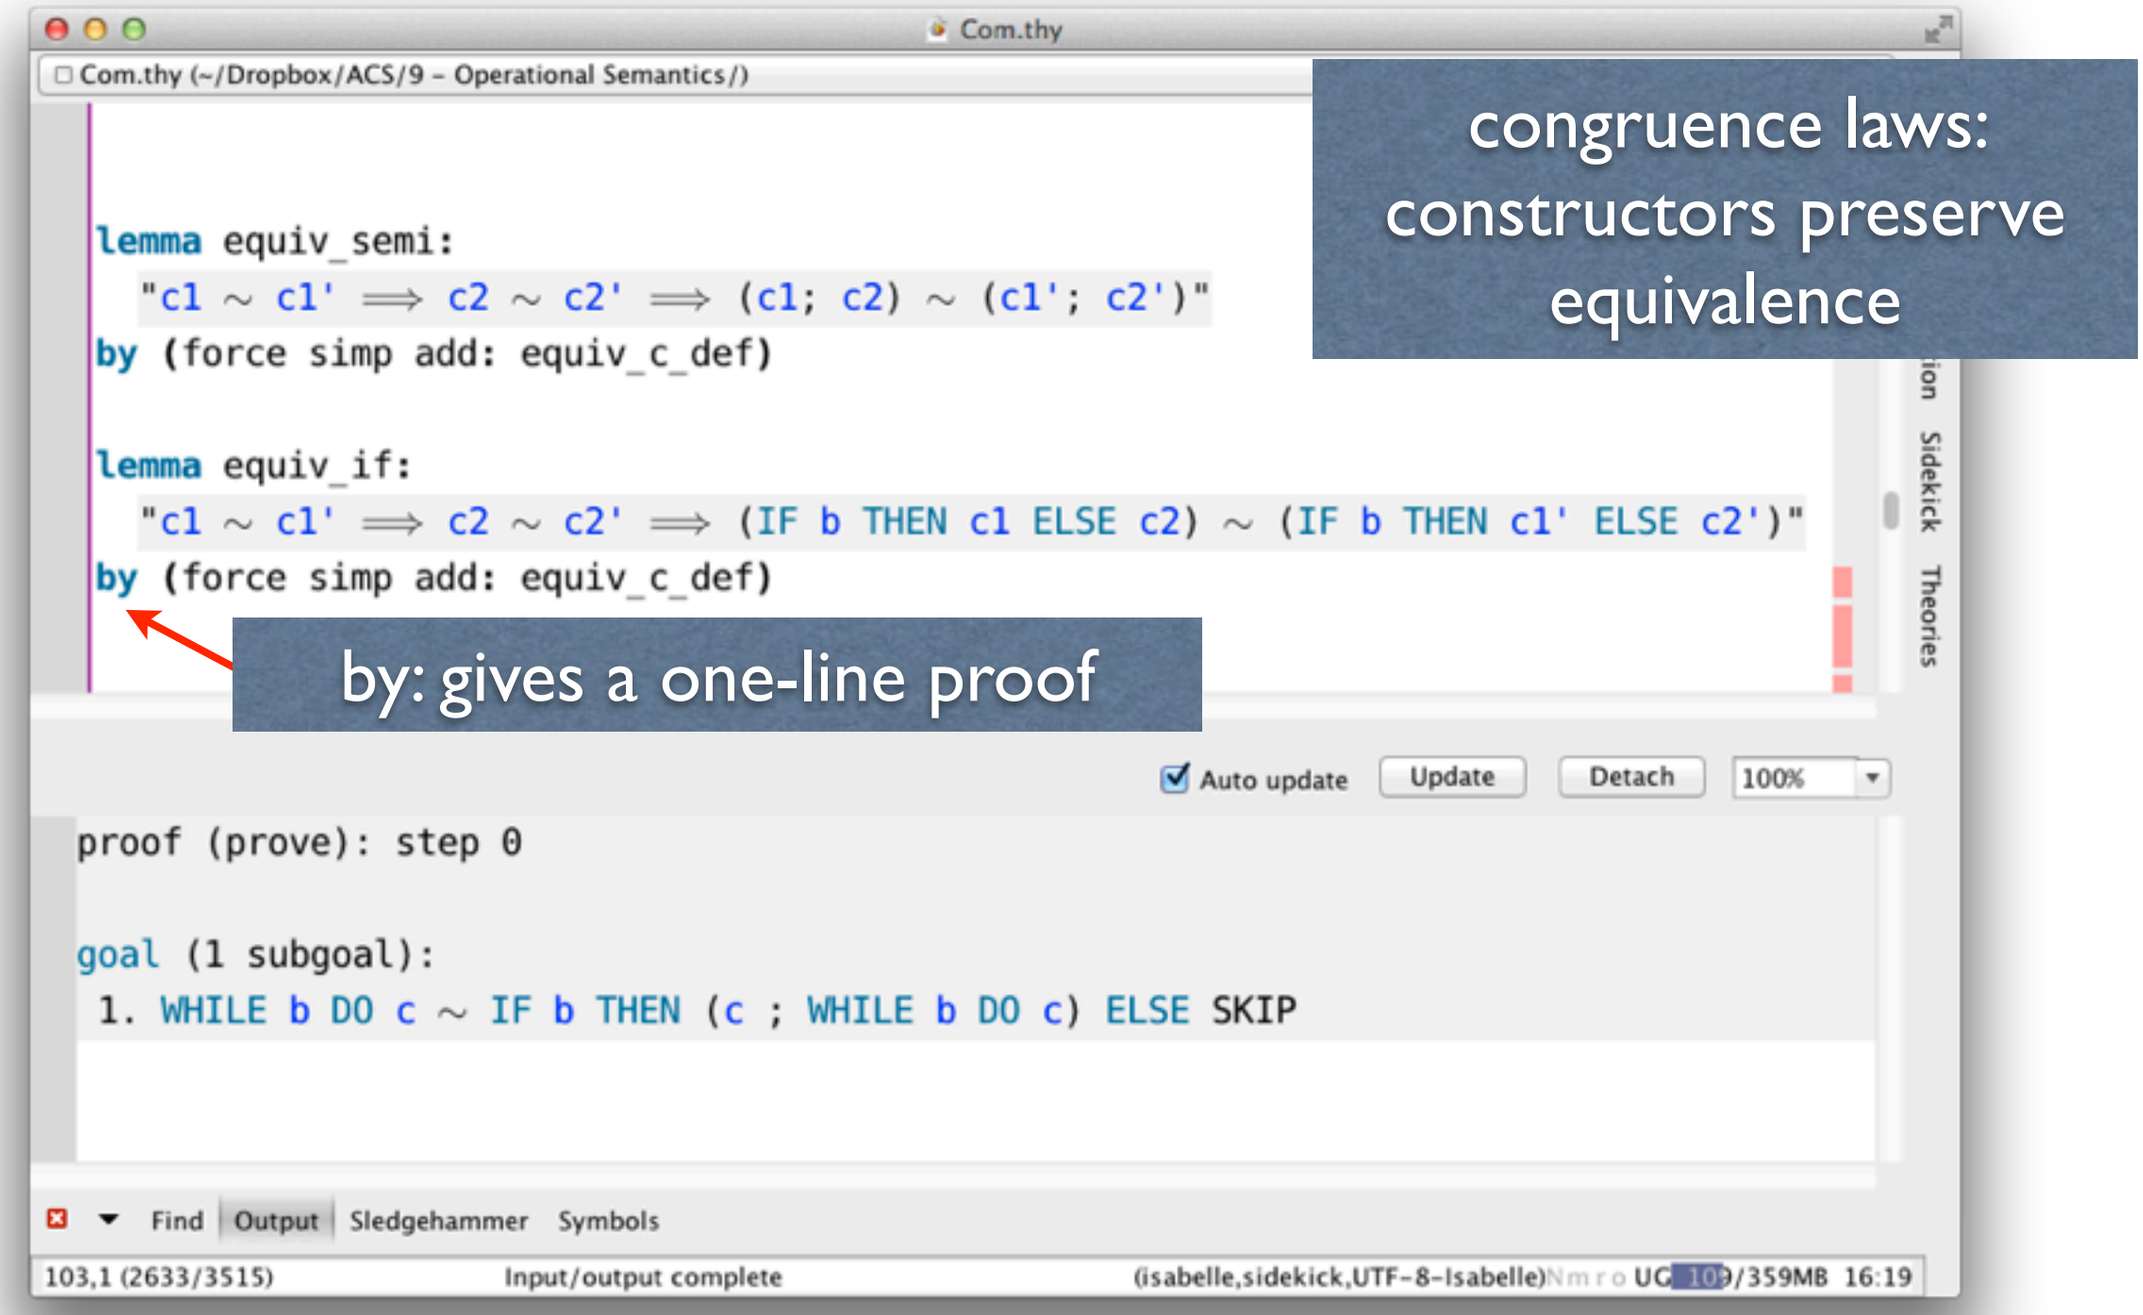

The properties shown here establish that semantic equivalence is a congruence relation with respect to the command constructors `Semi` and `Cond`. The proofs are again trivial, providing we remember to unfold the definition of semantic equivalence, `equiv_c`. Proving the analogous congruence property for While is harder, requiring rule induction with an induction formula similar to that used for another proof about `While` earlier in this lecture.

The proof method `force` is similar to `auto`, but it is more aggressive and it will not terminate until it has proved the subgoal it was applied to. In these examples, `auto` will give up too easily.

And More!!

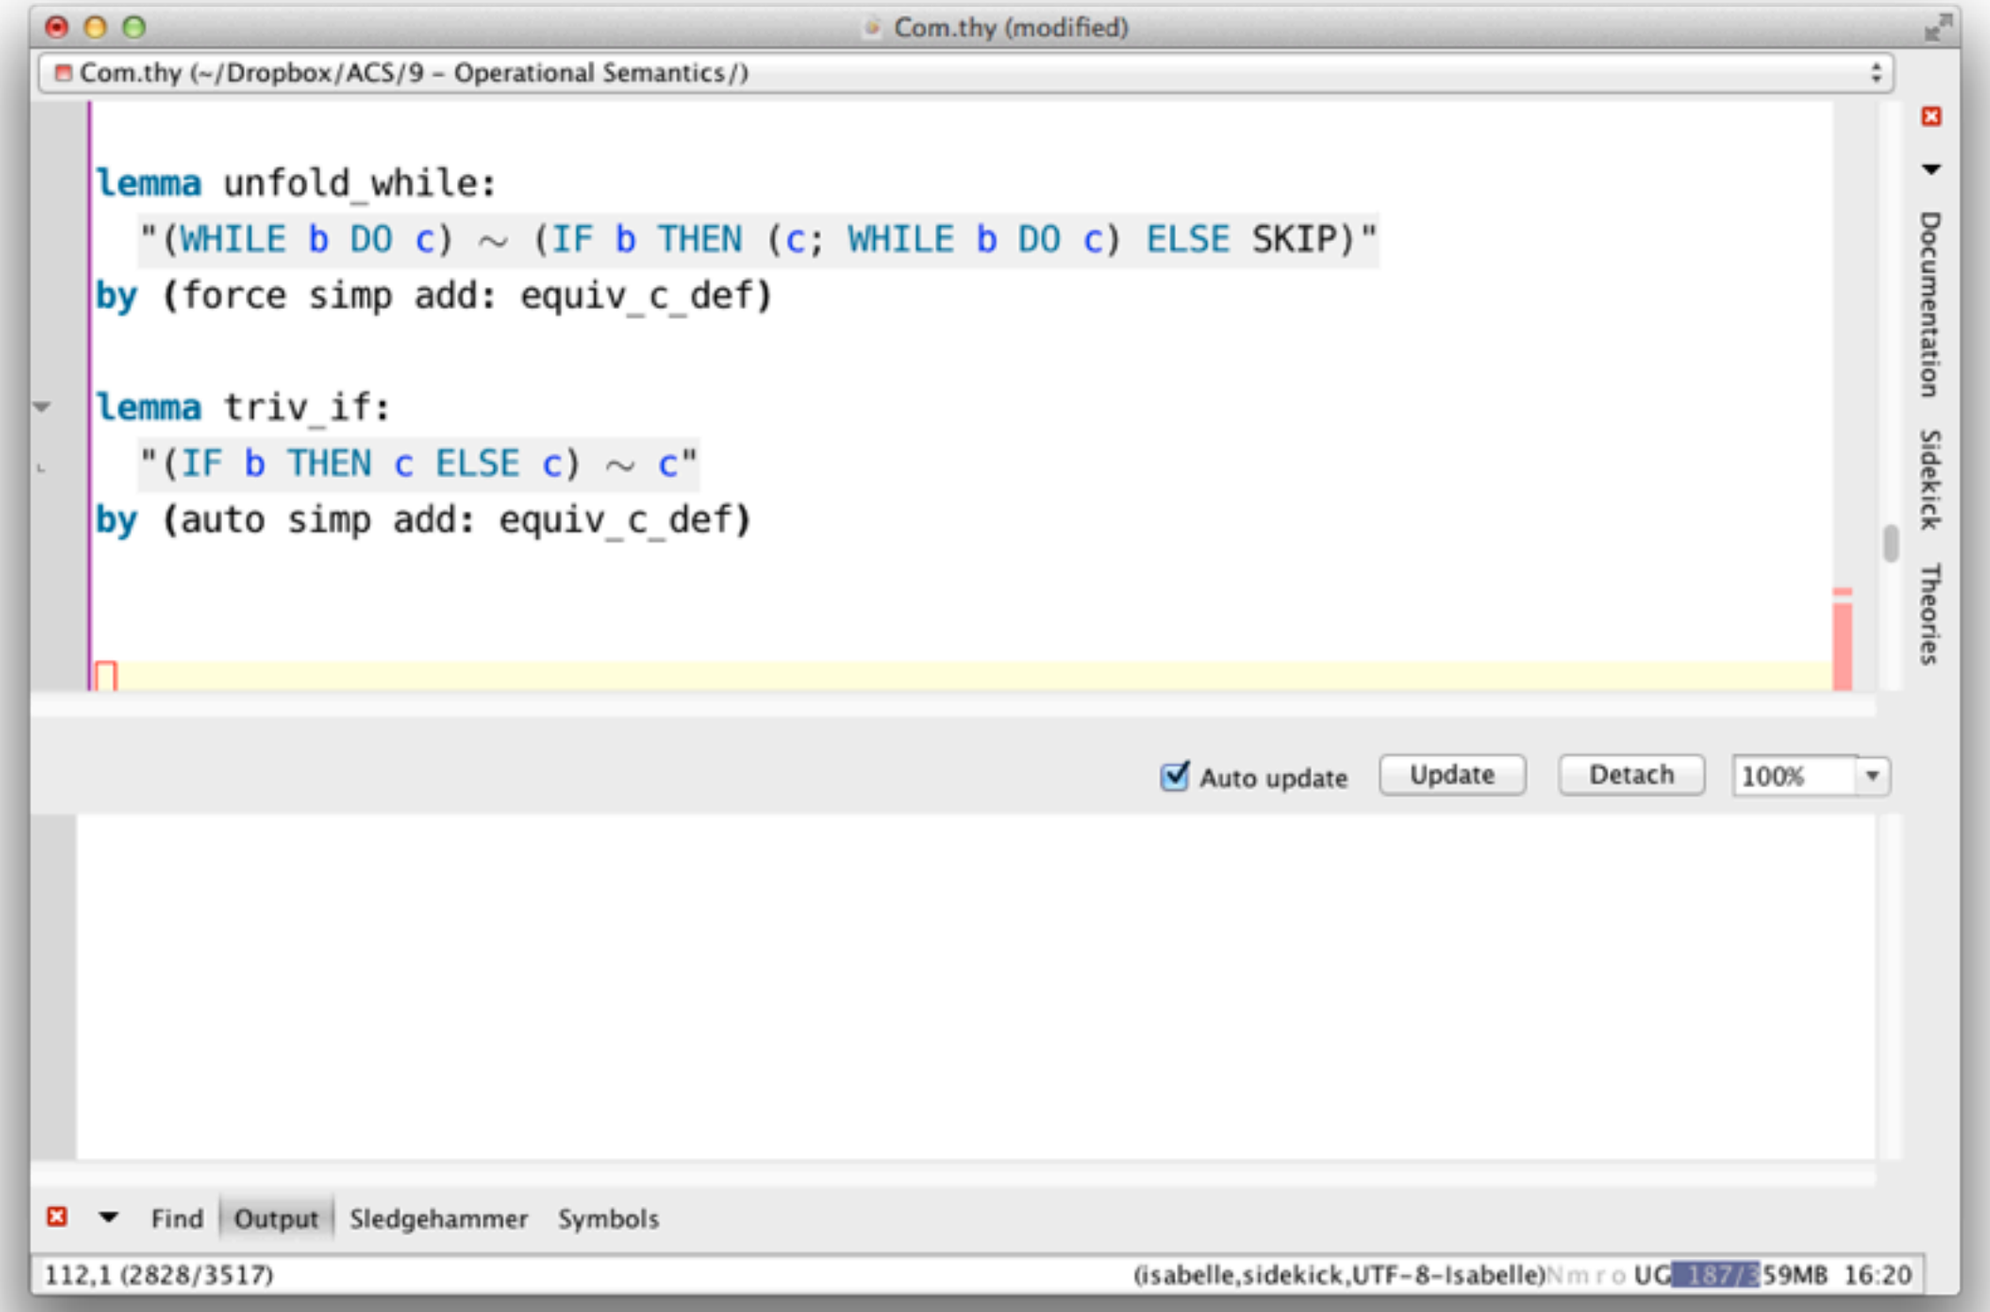

By some fluke, `force` will not solve the second of these. Sometimes you just have to try different things.

Note that a proof consisting of a single proof method can be written using the command “by”, which is more concise than writing “`apply`” followed by “`done`”. It is a small matter here, but structured proofs (which we are about to discuss) typically consist of numerous one line proofs expressed using “by”.

Intro-Rule for Equivalence

$$\frac{\langle c, s\rangle \to s' \iff \langle c', s\rangle \to s'}{c \sim c'} \quad s \text{ and } s' \text{ not free...}$$

declared like this

formalised like this

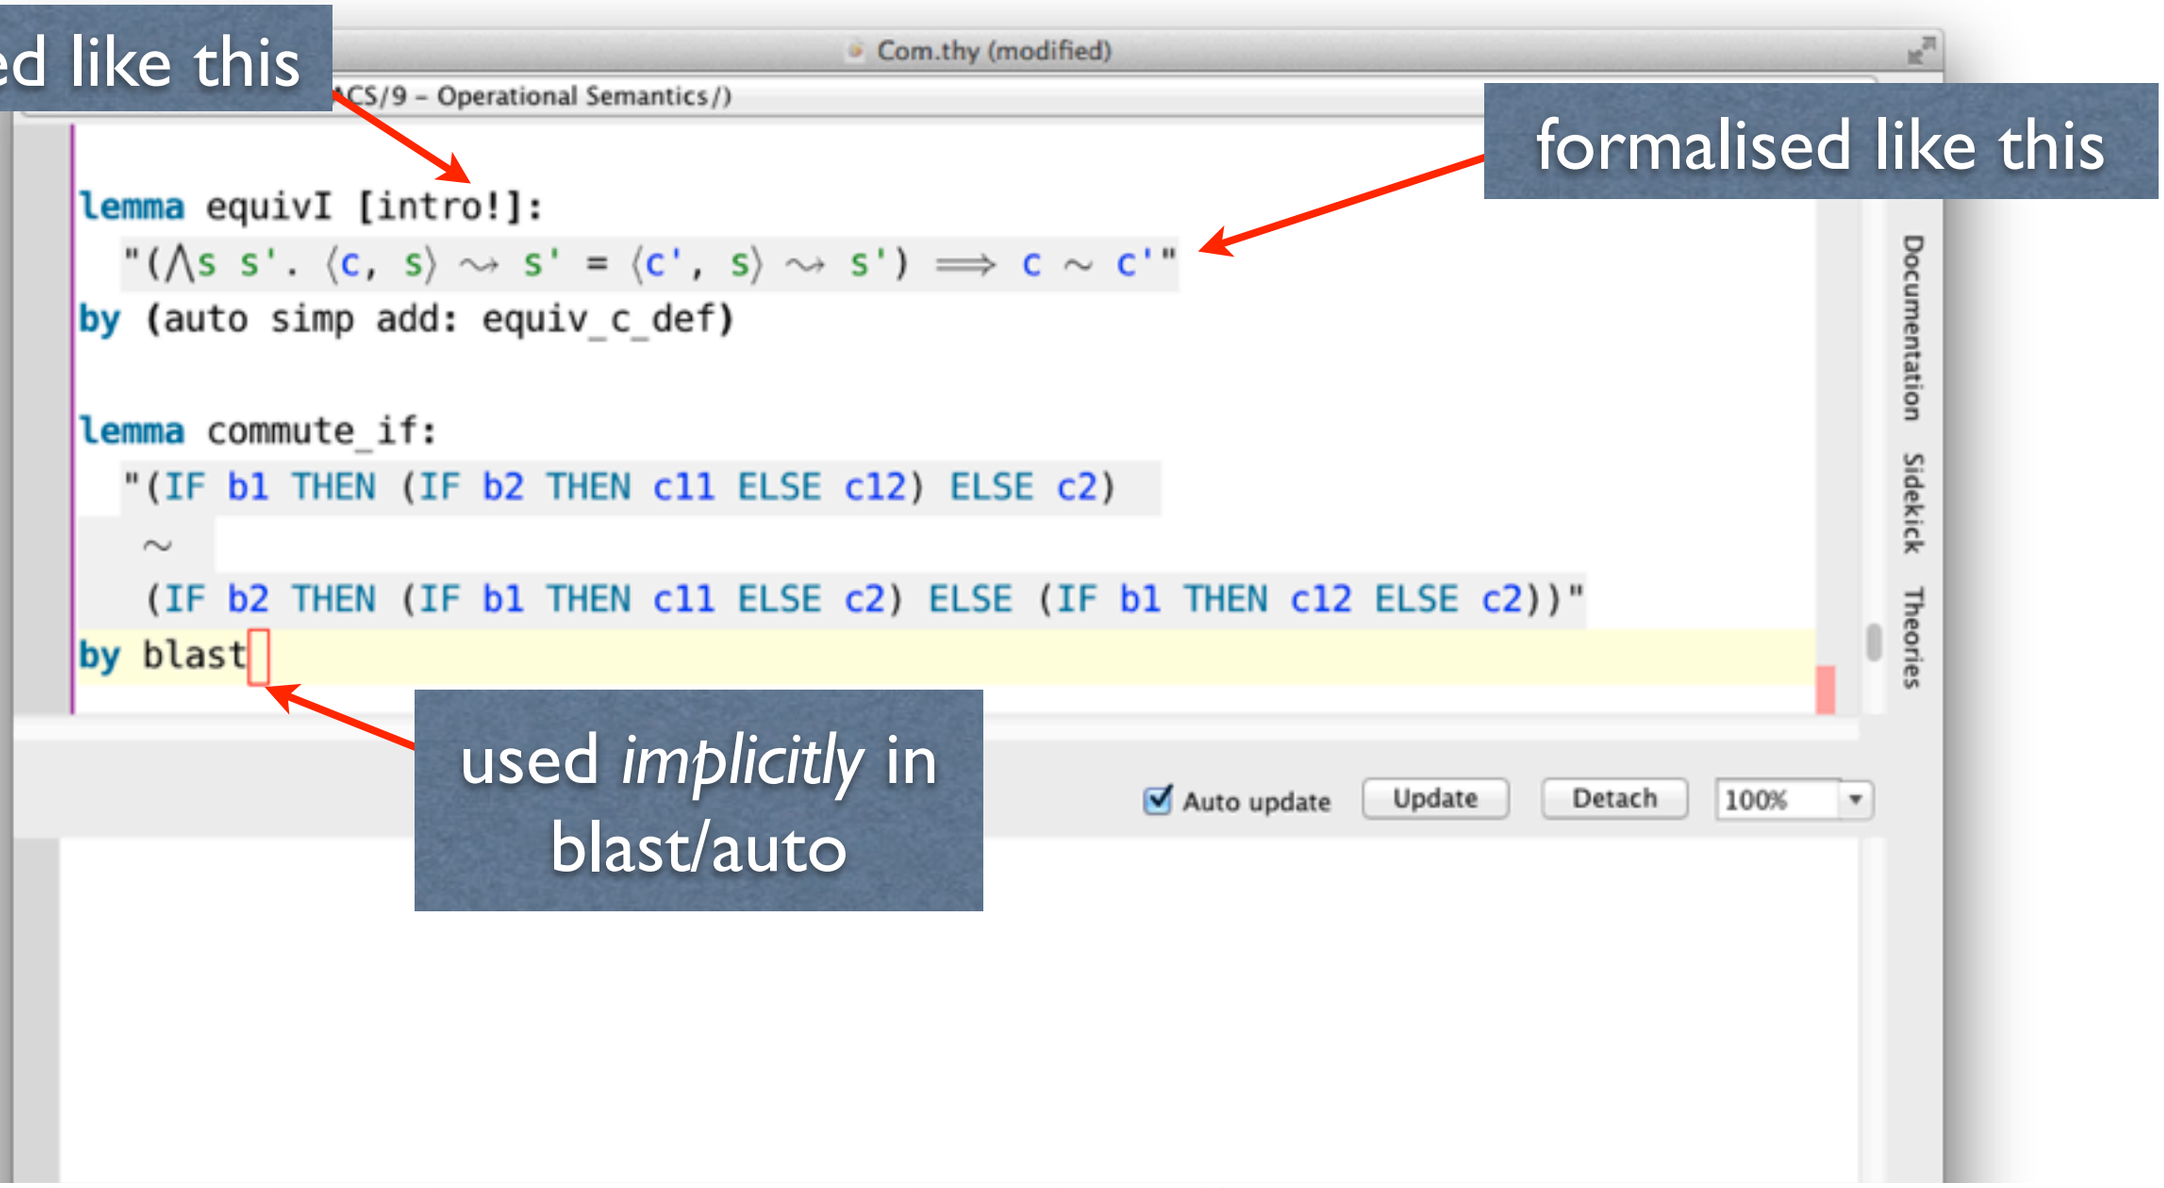

Giving the attribute `intro!` to a theorem informs Isabelle's automatic proof methods, including `auto`, `force` and `blast`, that this theorem should be used as an introduction rule. In other words, it should be used in backward-chaining mode: the conclusion of the rule is unified with the subgoal, continuing the search from that rule's premises. It is now unnecessary to mention this theorem when calling those proof methods. The theorem shown can now be proved using blast alone. We do not need to refer to `equivI` or to the definition of `equiv_c`. The approach used to prove other examples of semantic equivalence in this lecture do not terminate on this problem in a reasonable time. The proof shown only requires 12 ms.

The exclamation mark (`!`) tells Isabelle to apply the rule aggressively. It is appropriate when the premise of the rule is equivalent to the conclusion; equivalently, it is appropriate when applying the rule can never be a mistake. The weaker attribute `intro` should be used for a theorem that is one of many different ways of proving its conclusion.

Final Remarks on Semantics

- *Small-step semantics* can be treated similarly.
- *Variable binding* is crucial in larger examples, and should be formalised using the *nominal package*.
  - choosing a *fresh* variable
  - *renaming* bound variables consistently
- Serious proofs will be complex and difficult!

Documentation on the nominal package can be downloaded from `http://isabelle.in.tum.de/nominal/`

Many examples are distributed with Isabelle. See the directory `HOL/Nominal/Examples`.

Other relevant publications are available from Christian Urban's website: `http://www4.in.tum.de/~urbanc/publications.html`

Interactive Formal Verification
10: Structured Induction Proofs

Lawrence C Paulson
Computer Laboratory
University of Cambridge

Structured (Isar) Proofs

- As we've already seen:
  - Structured proofs are clearer than a series of commands, but verbose.
  - The Isar language is rich and complex, supporting a great many proof styles.
- *But there's more!*
- *Existential* reasoning: naming entities that "exist".
- Syntax for proof by *induction*.
  - No need to write out induction hypotheses.
  - Cases given by *name*; bound variables named.
- And the same syntax works for *case analysis*.

A Proof about Binary Trees

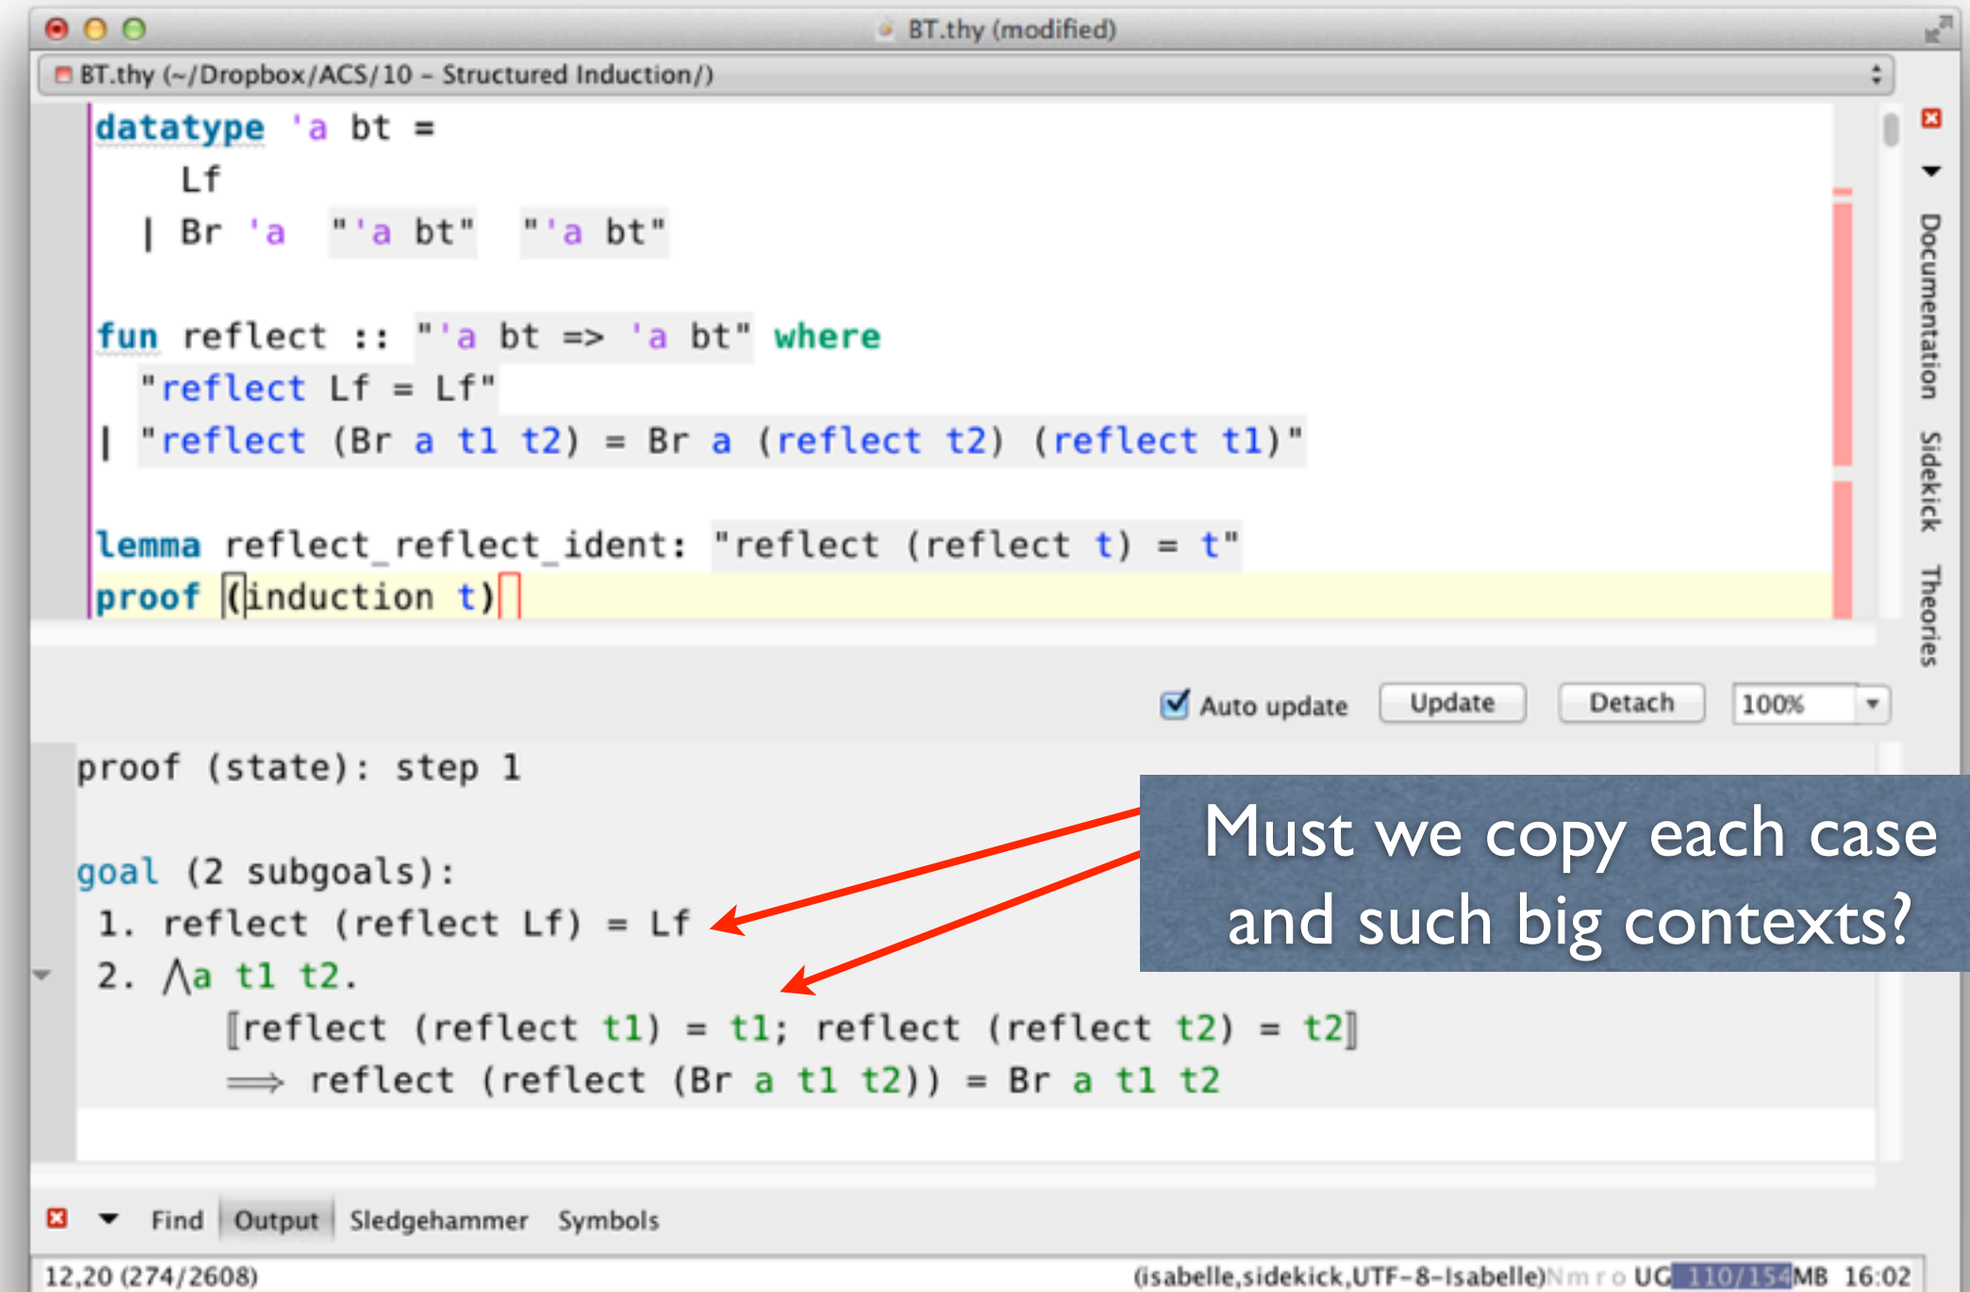

Inductive proofs frequently involve several subgoals, some of them with multiple assumptions and bound variables. Creating an Isar proof skeleton from scratch would be tiresome, and the resulting proof would be quite lengthy.

Finding Predefined Cases

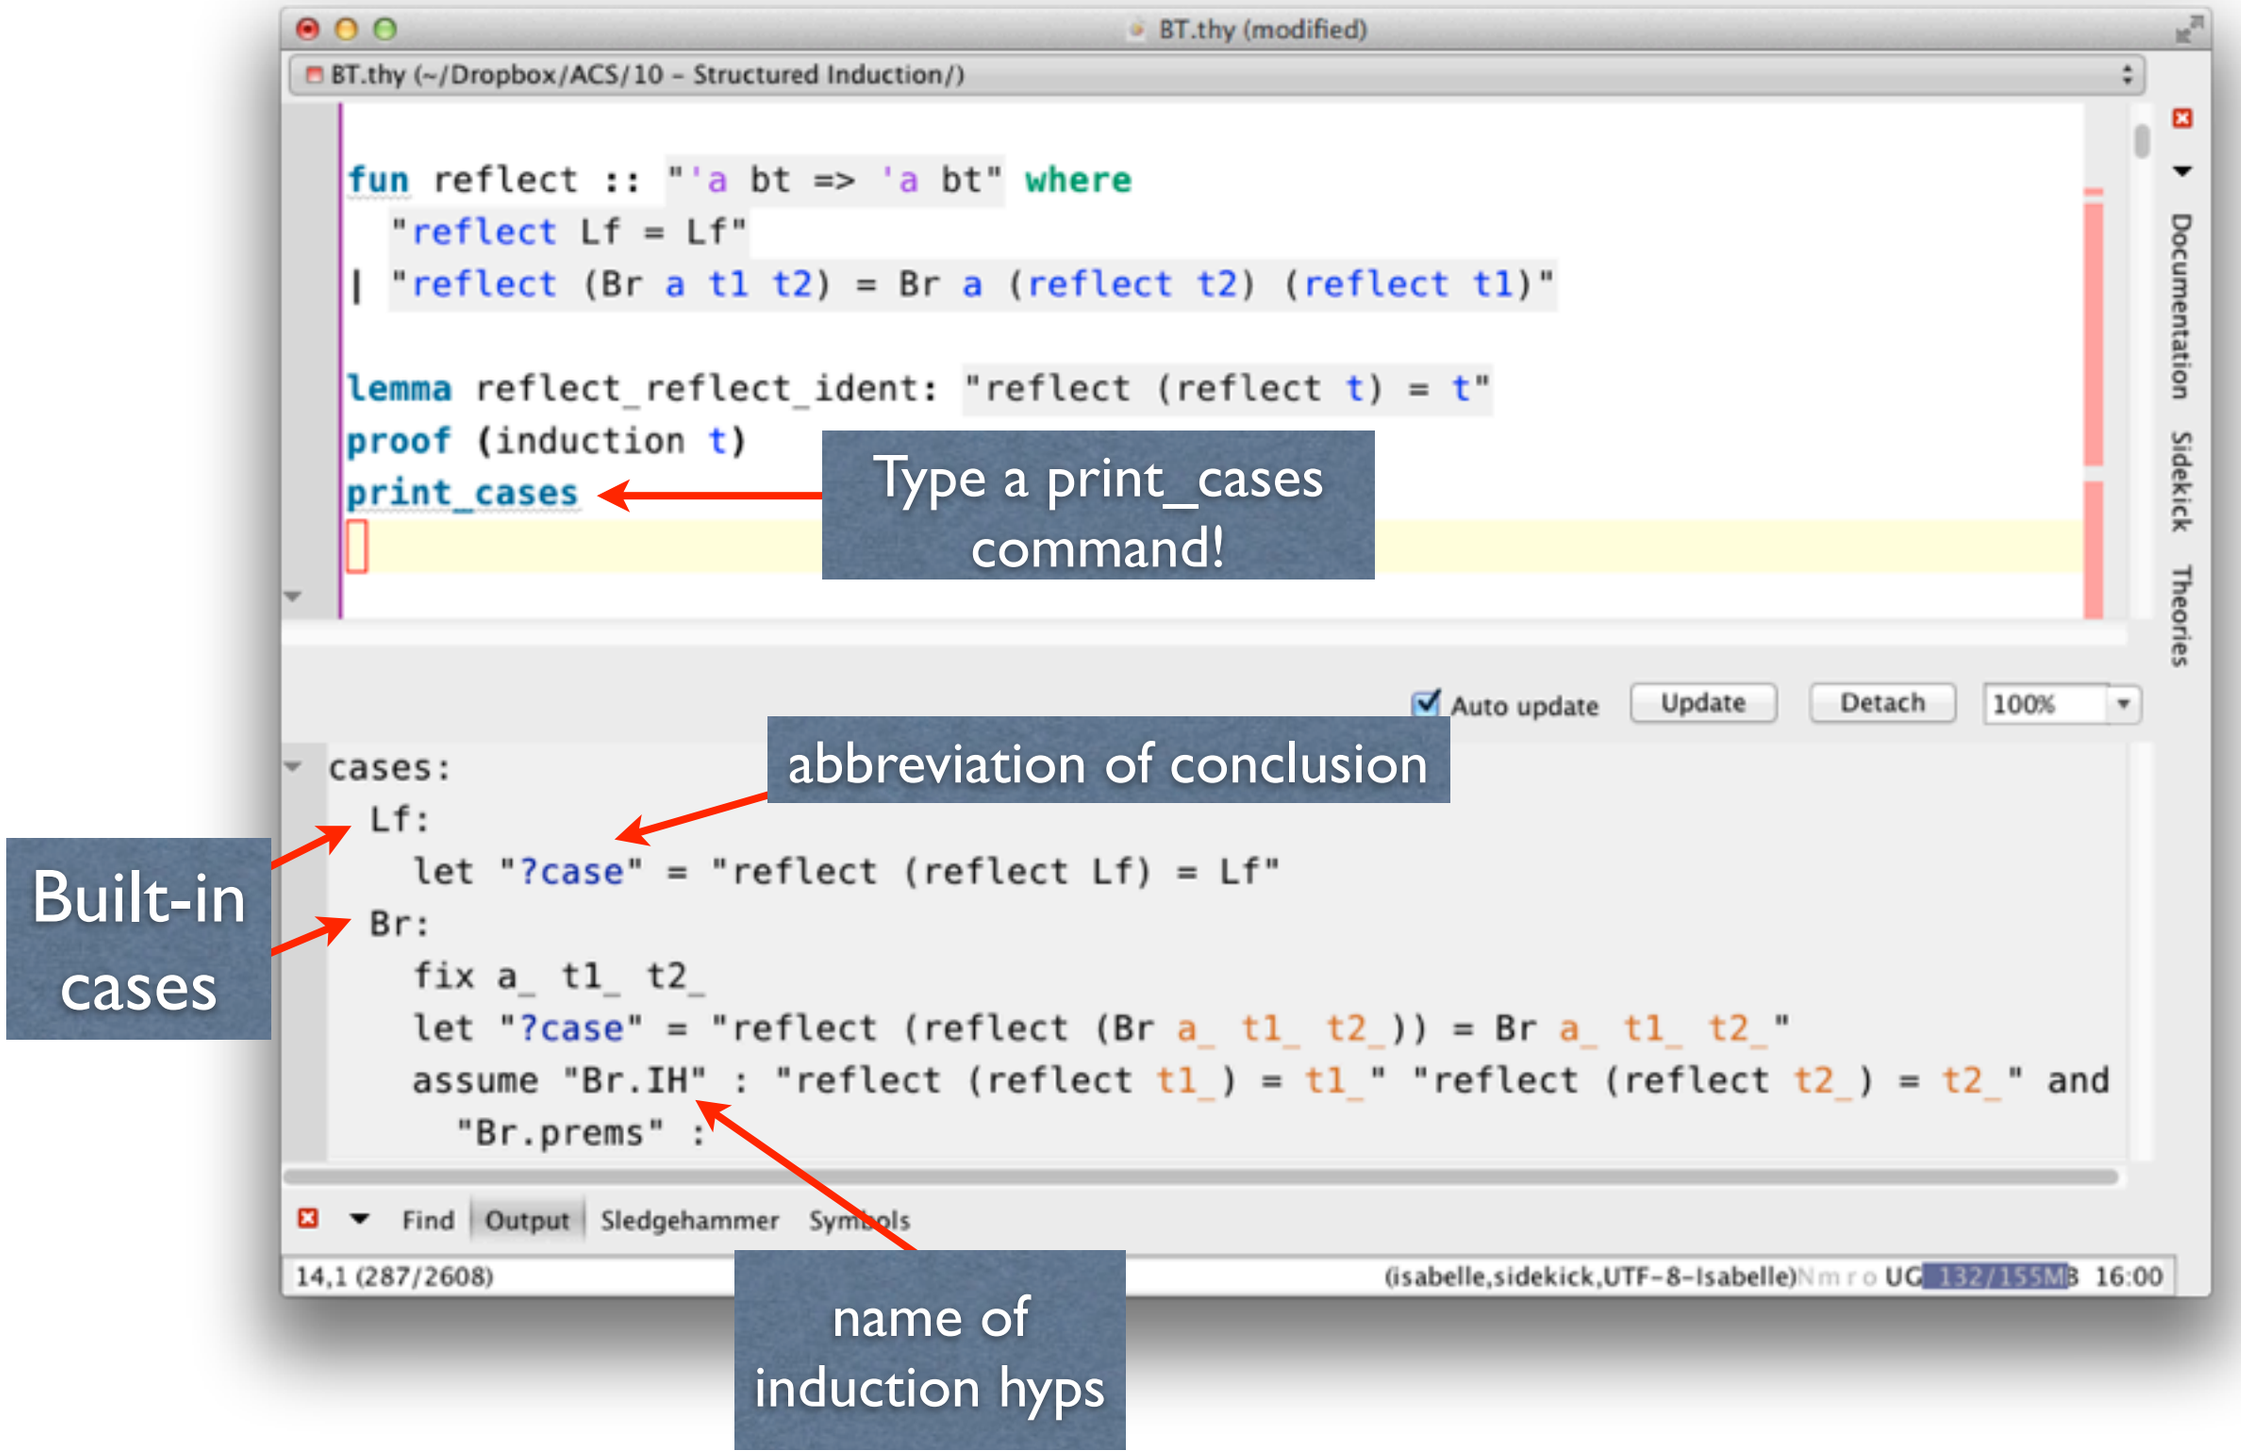

Many induction rules have attached cases designed for use with Isar. By referring to such a case, a proof script implicitly introduces the contexts shown above. There are placeholders for the bound variables (specific names must be given). Identifiers are introduced to denote induction hypotheses and other premises that accompany each case. Also, the identifier `?case` is introduced to abbreviate the required instance of the induction formula.

It is unfortunately necessary to type the command `print_cases` right in your document.

Proof Using Named Cases

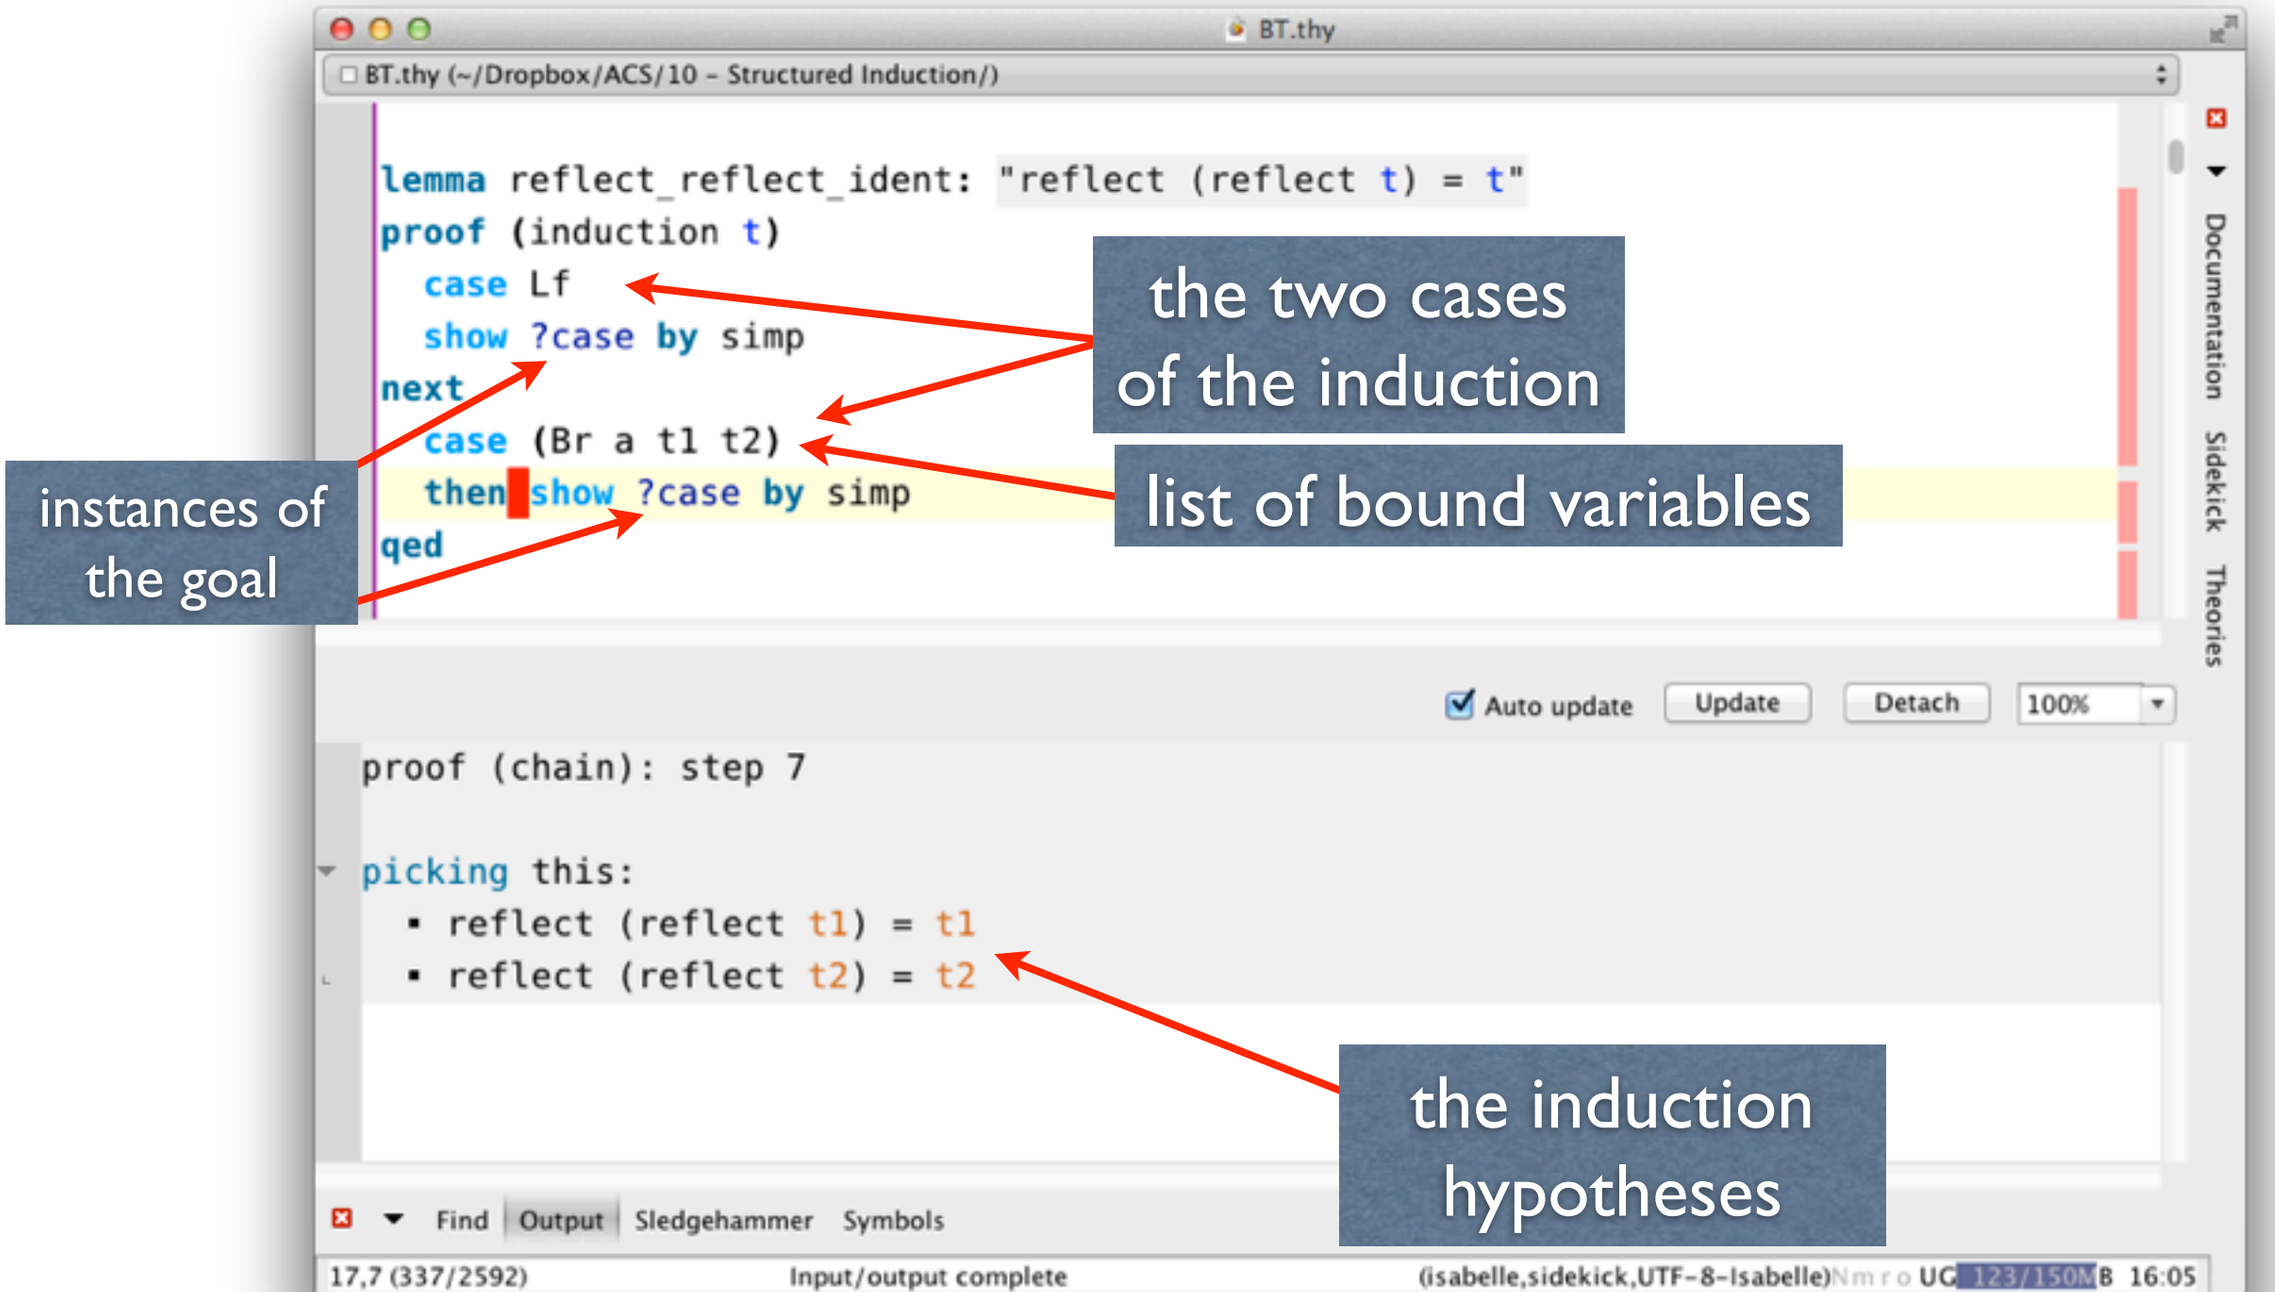

With all these abbreviations, the induction formula does not have to be repeated in its various instances. The instances that are to be proved are abbreviated as `?case`; they (and the induction hypotheses) are automatically generated from the supplied list of bound variables.

Observe the use of "`then`" with "`show`" in the inductive case, thereby providing the induction hypotheses to the method. In a more complicated proof, these hypotheses can be denoted by the identifier `Br.hyps`.

Induction with a Context

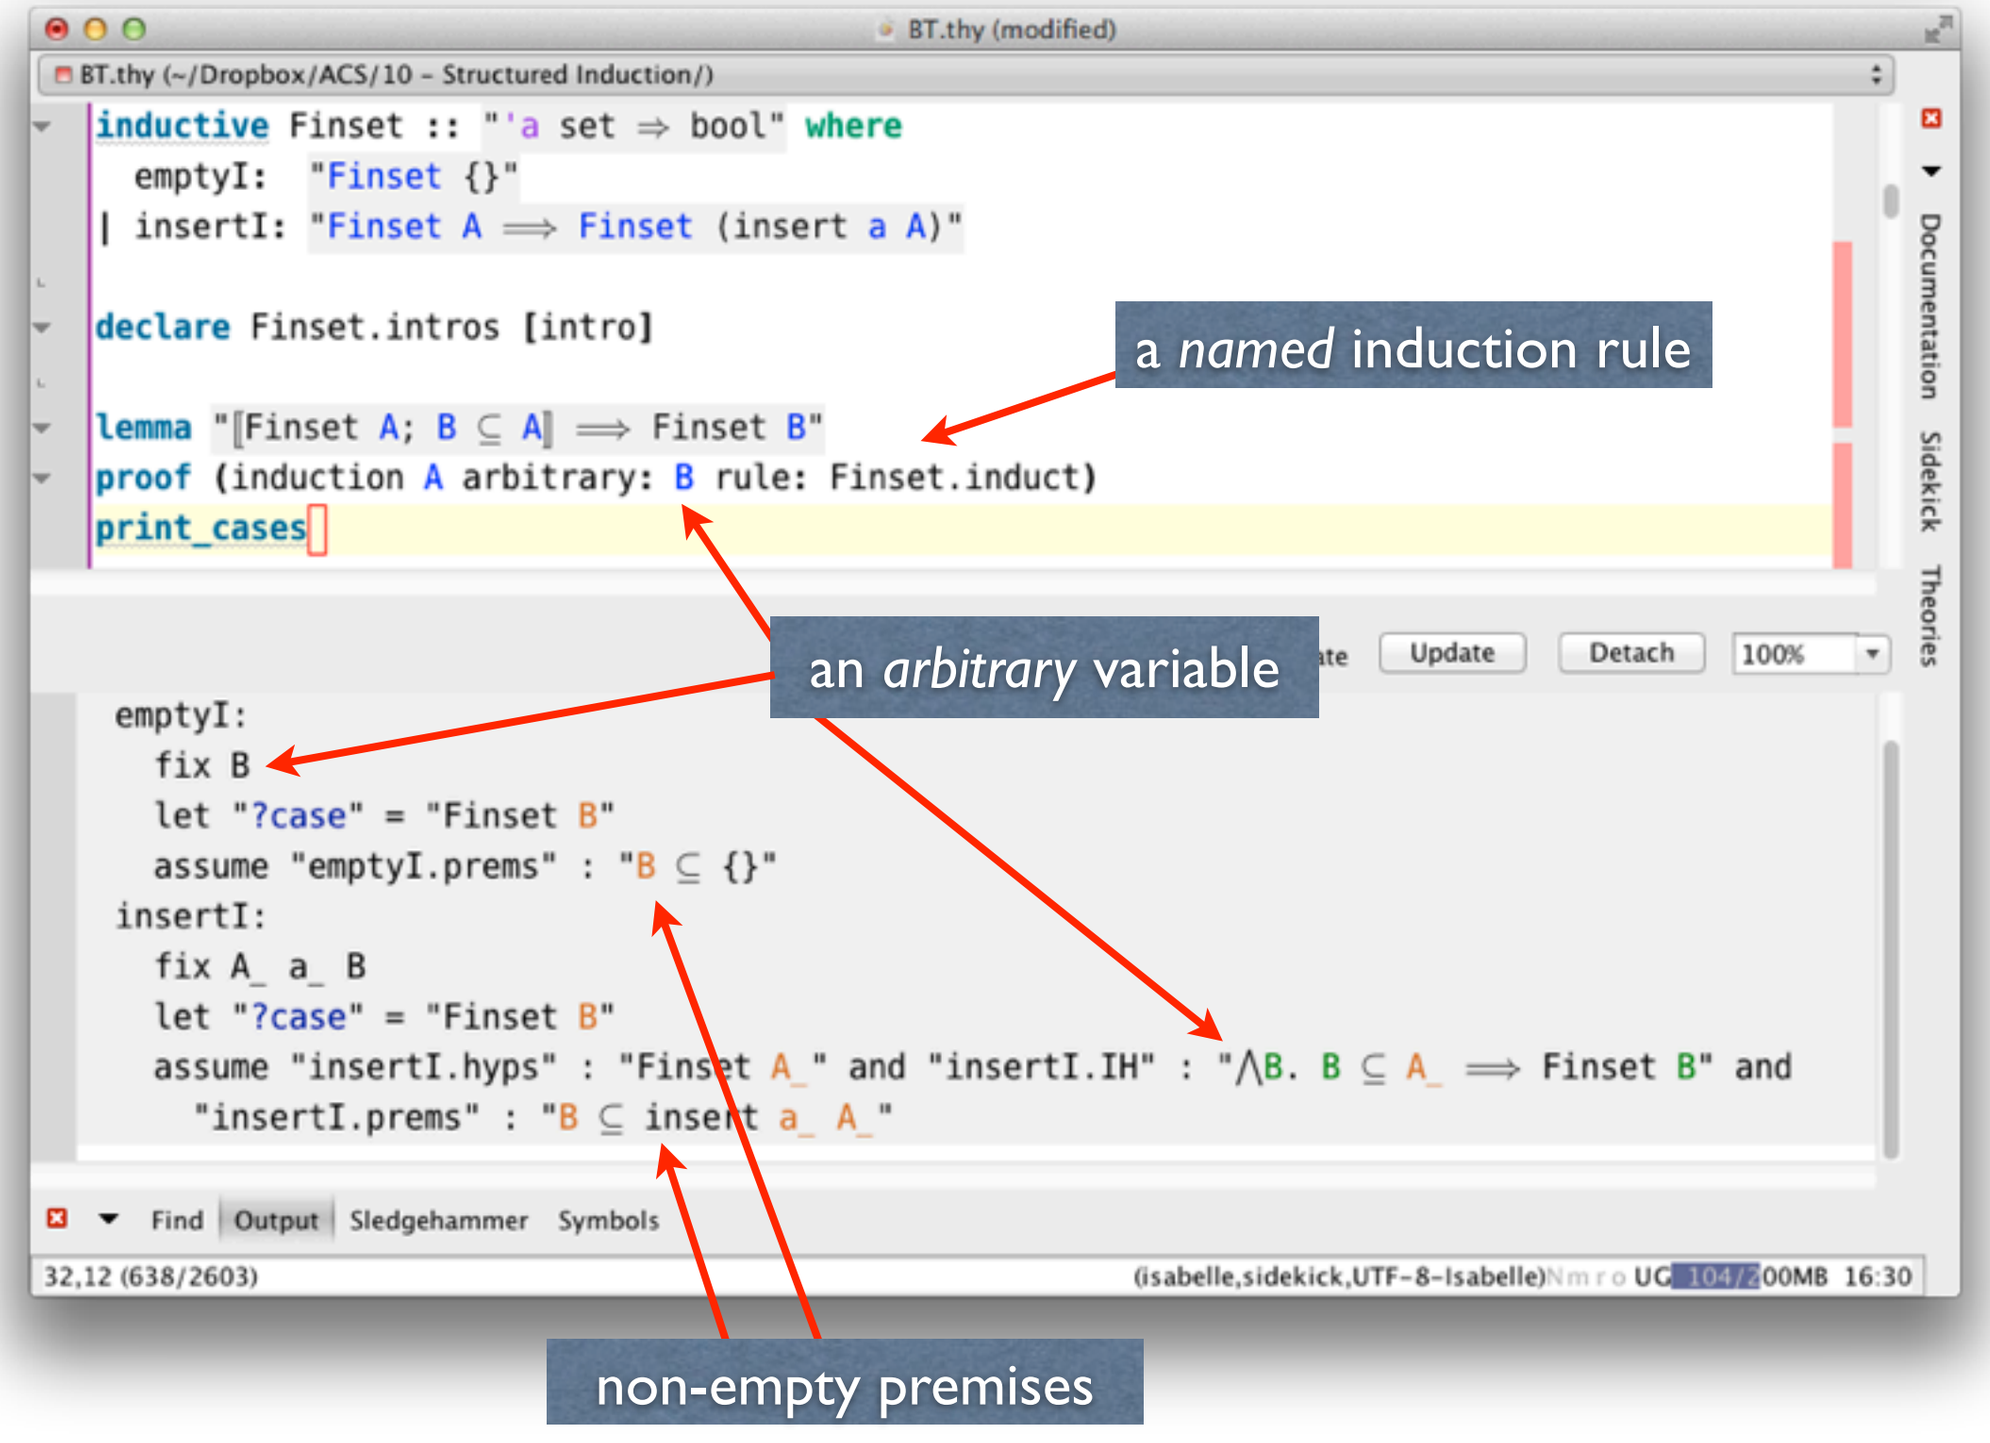

An inductive definition generates an induction rule with one case (correspondingly named) for each introduction rule. This particular proof requires the variable B to be taken as arbitrary, which means, universally quantified: it becomes an additional bound variable in each case. This proof also carries along a further premise, B⊆A, instances of which are attached to both subgoals.

Proving the Base Case

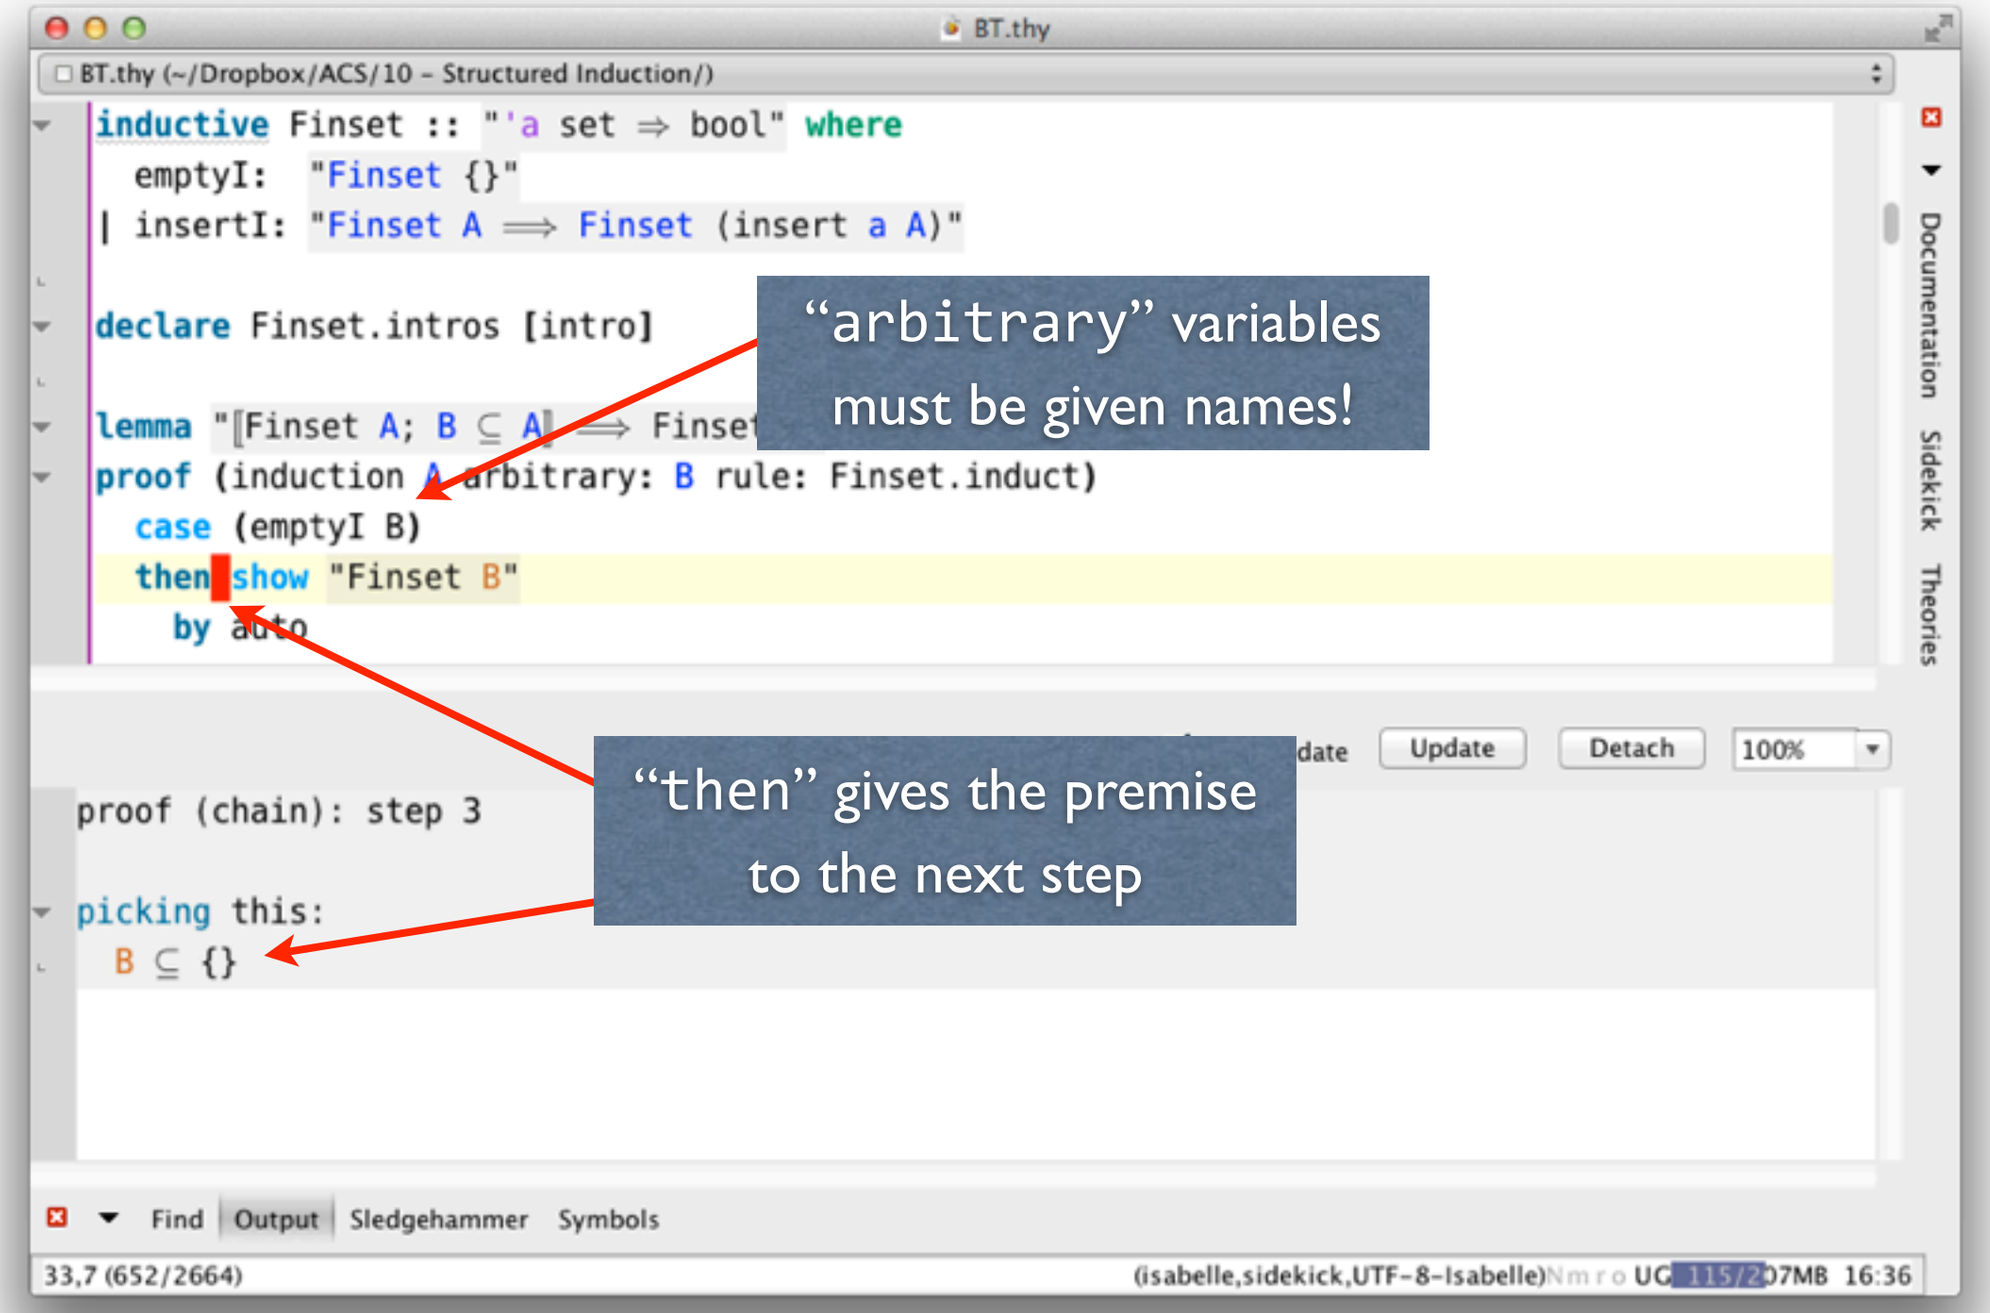

The base case would normally be just `emptyI`. But here, there is an additional bound variable. Note that we could have written, for example, `(emptyI C)` and Isabelle would have adjusted everything to use C instead of B.

A Nested Case Analysis

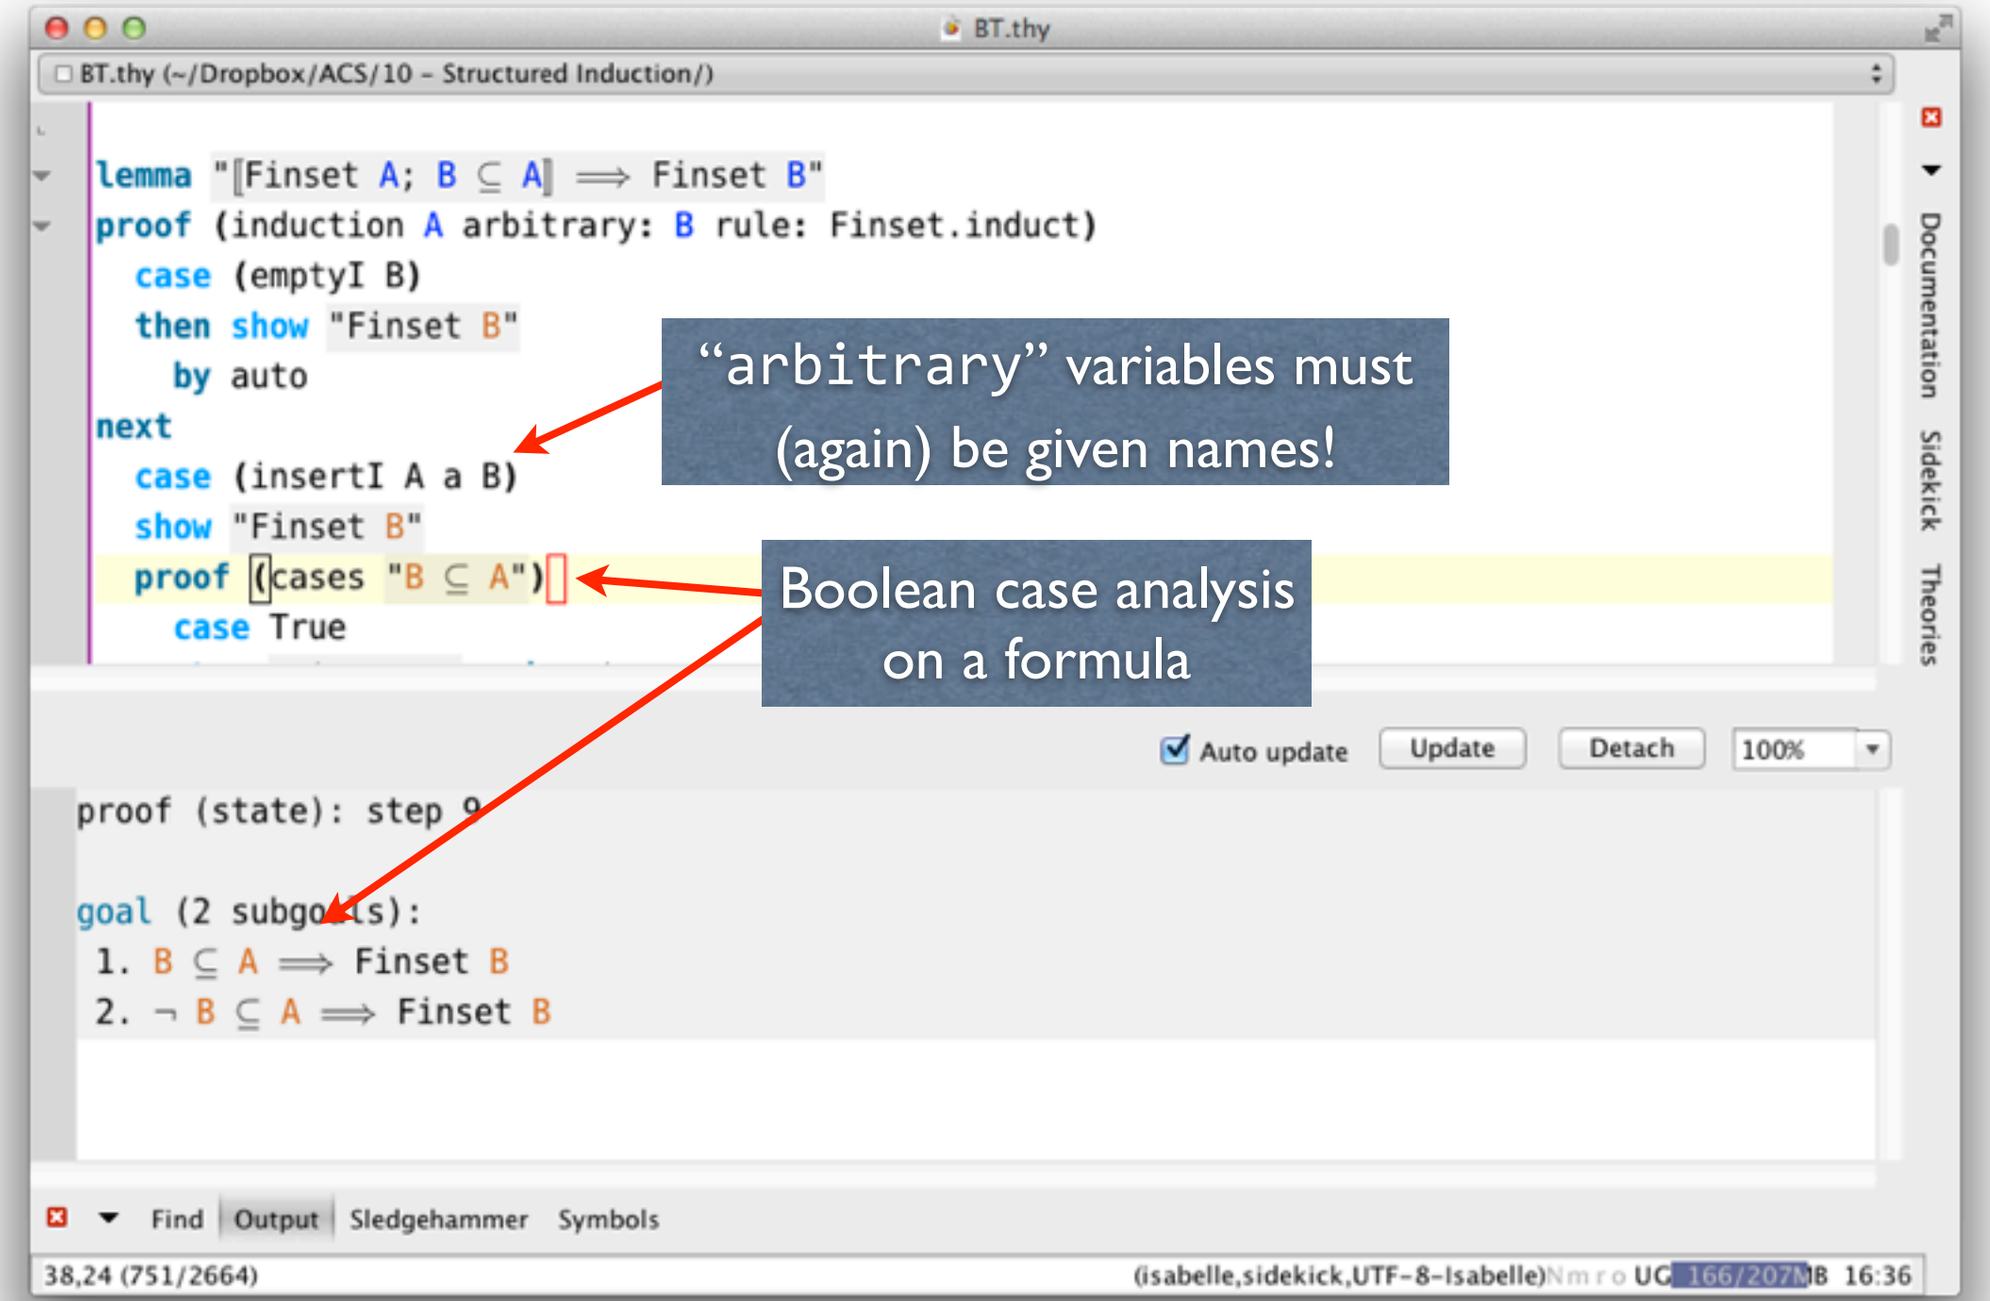

Here we know B ⊆ insert a A, as it is the inherited premise of this case. But do we in fact know B⊆A?

The Complete Proof

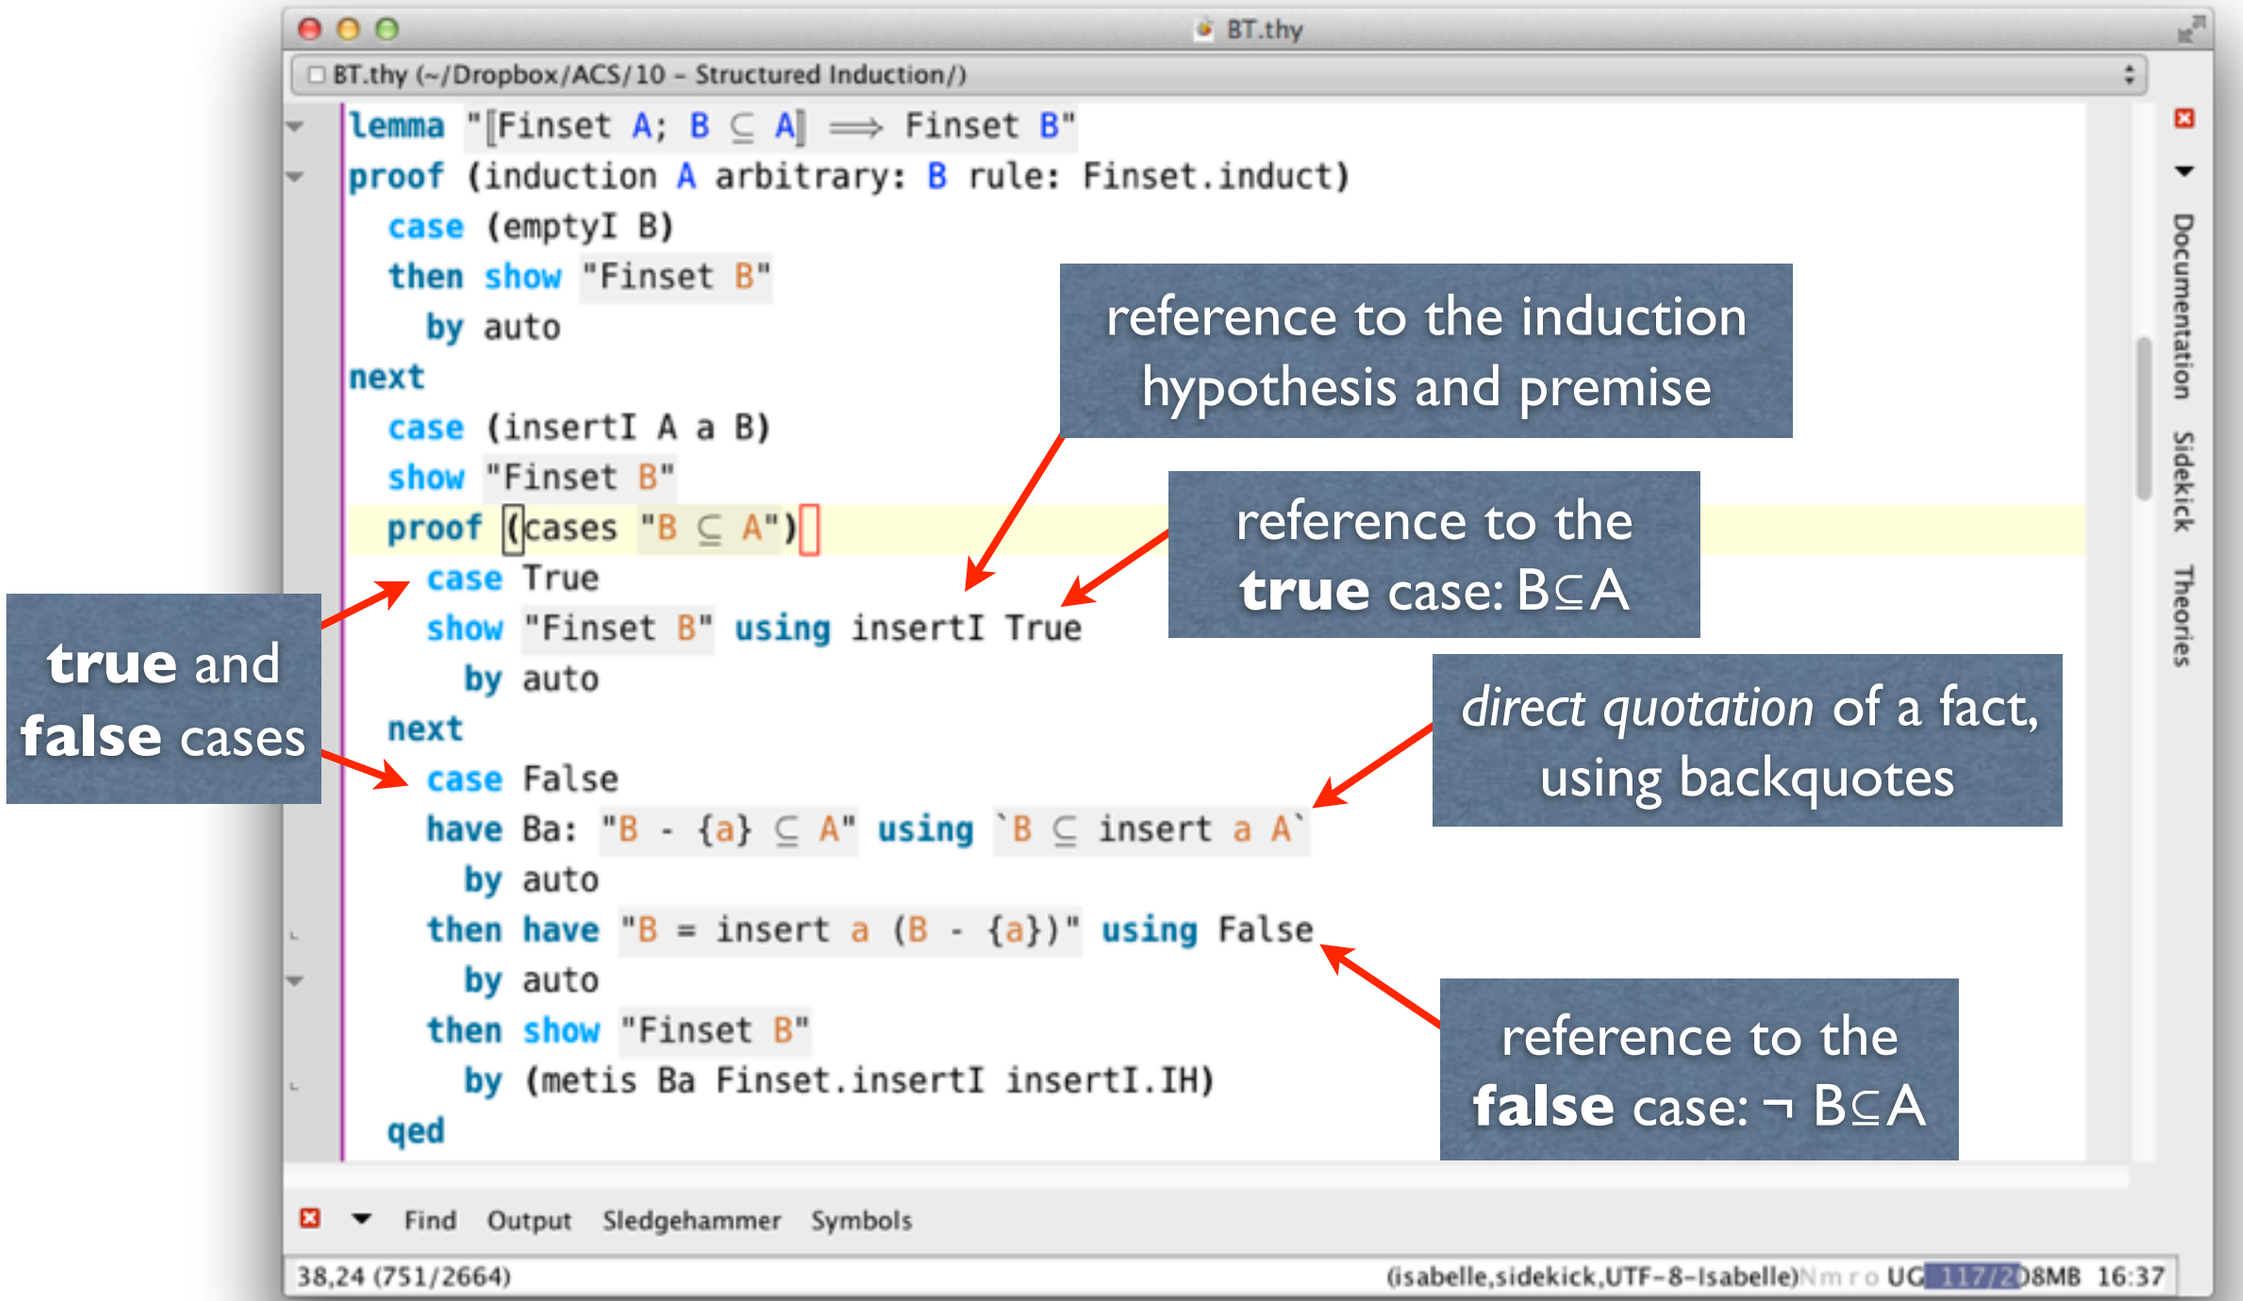

Here is an outline of the proof. If B⊆A, then it is trivial, as we can immediately use the induction hypothesis. If not, then we apply the induction hypothesis to the set B-{a}. We deduce that B-{a} ∈ Fin, and therefore B = insert a (B-{a}) ∈ Fin.

This proof script contains many references to facts. The facts attached to the case of an inductive proof or case analysis are denoted by the name of that case, for example, insertI, True or False. We can also refer to a theorem by enclosing the actual theorem statement in backward quotation marks. We see this above in the proof of B-{a} ⊆ A.

Additional Proof Structures

```
case (insertI A a B)
show "Finset B"
proof (cases "B ⊆ A")
  case True
  show "Finset B" using insertI True
    by auto
next
  case False
  have Ba: "B - {a} ⊆ A" using `B ⊆ insert a A`
    by auto
  then have "B = insert a (B - {a})" using False
    by auto
  then show "Finset B"
    by (metis Ba Finset.insertI insertI.IH)
qed
```

→

```
case (insertI A a B)
show "Finset B"
proof (cases "B ⊆ A")
  case True
  with insertI show "Finset B"
    by auto
next
  case False
  from `B ⊆ insert a A` have Ba: "B - {a} ⊆ A"
    by auto
  with False have "B = insert a (B - {a})"
    by auto
  with Ba insertI.IH show "Finset B"
    by (metis Finset.insertI)
qed
```

from ⟨*facts*⟩ … = … using ⟨*facts*⟩

with ⟨*facts*⟩ … = then from ⟨*facts*⟩ …

Full details, probably much more than you want at this stage, can be found in *The Isabelle/Isar Reference Manual* by Makarius Wenzel.

Viewing Available Facts

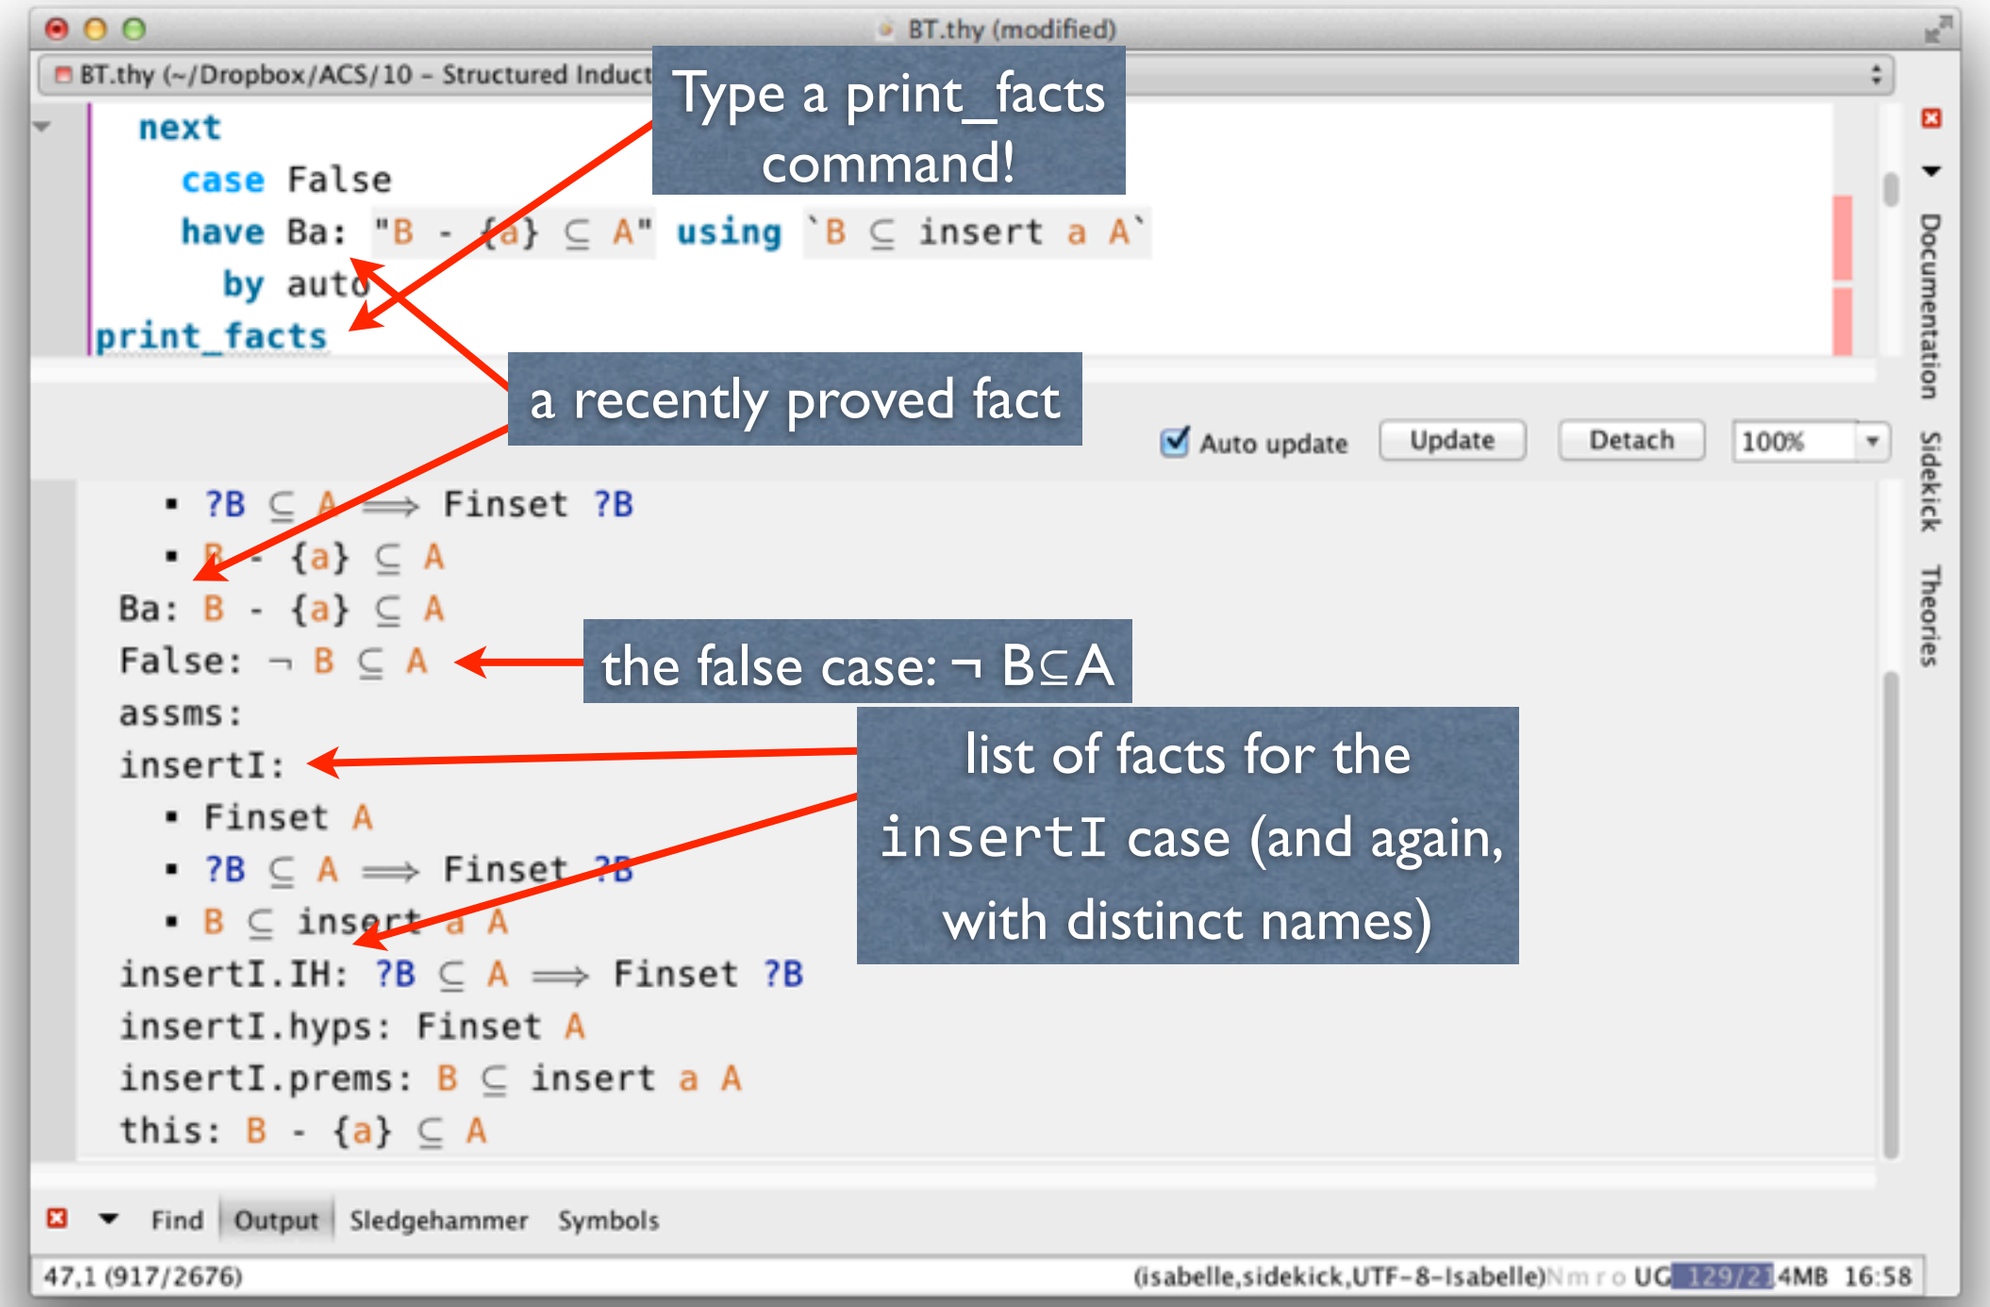

It is unfortunately necessary to type the command `print_facts` right in your document.

Popups

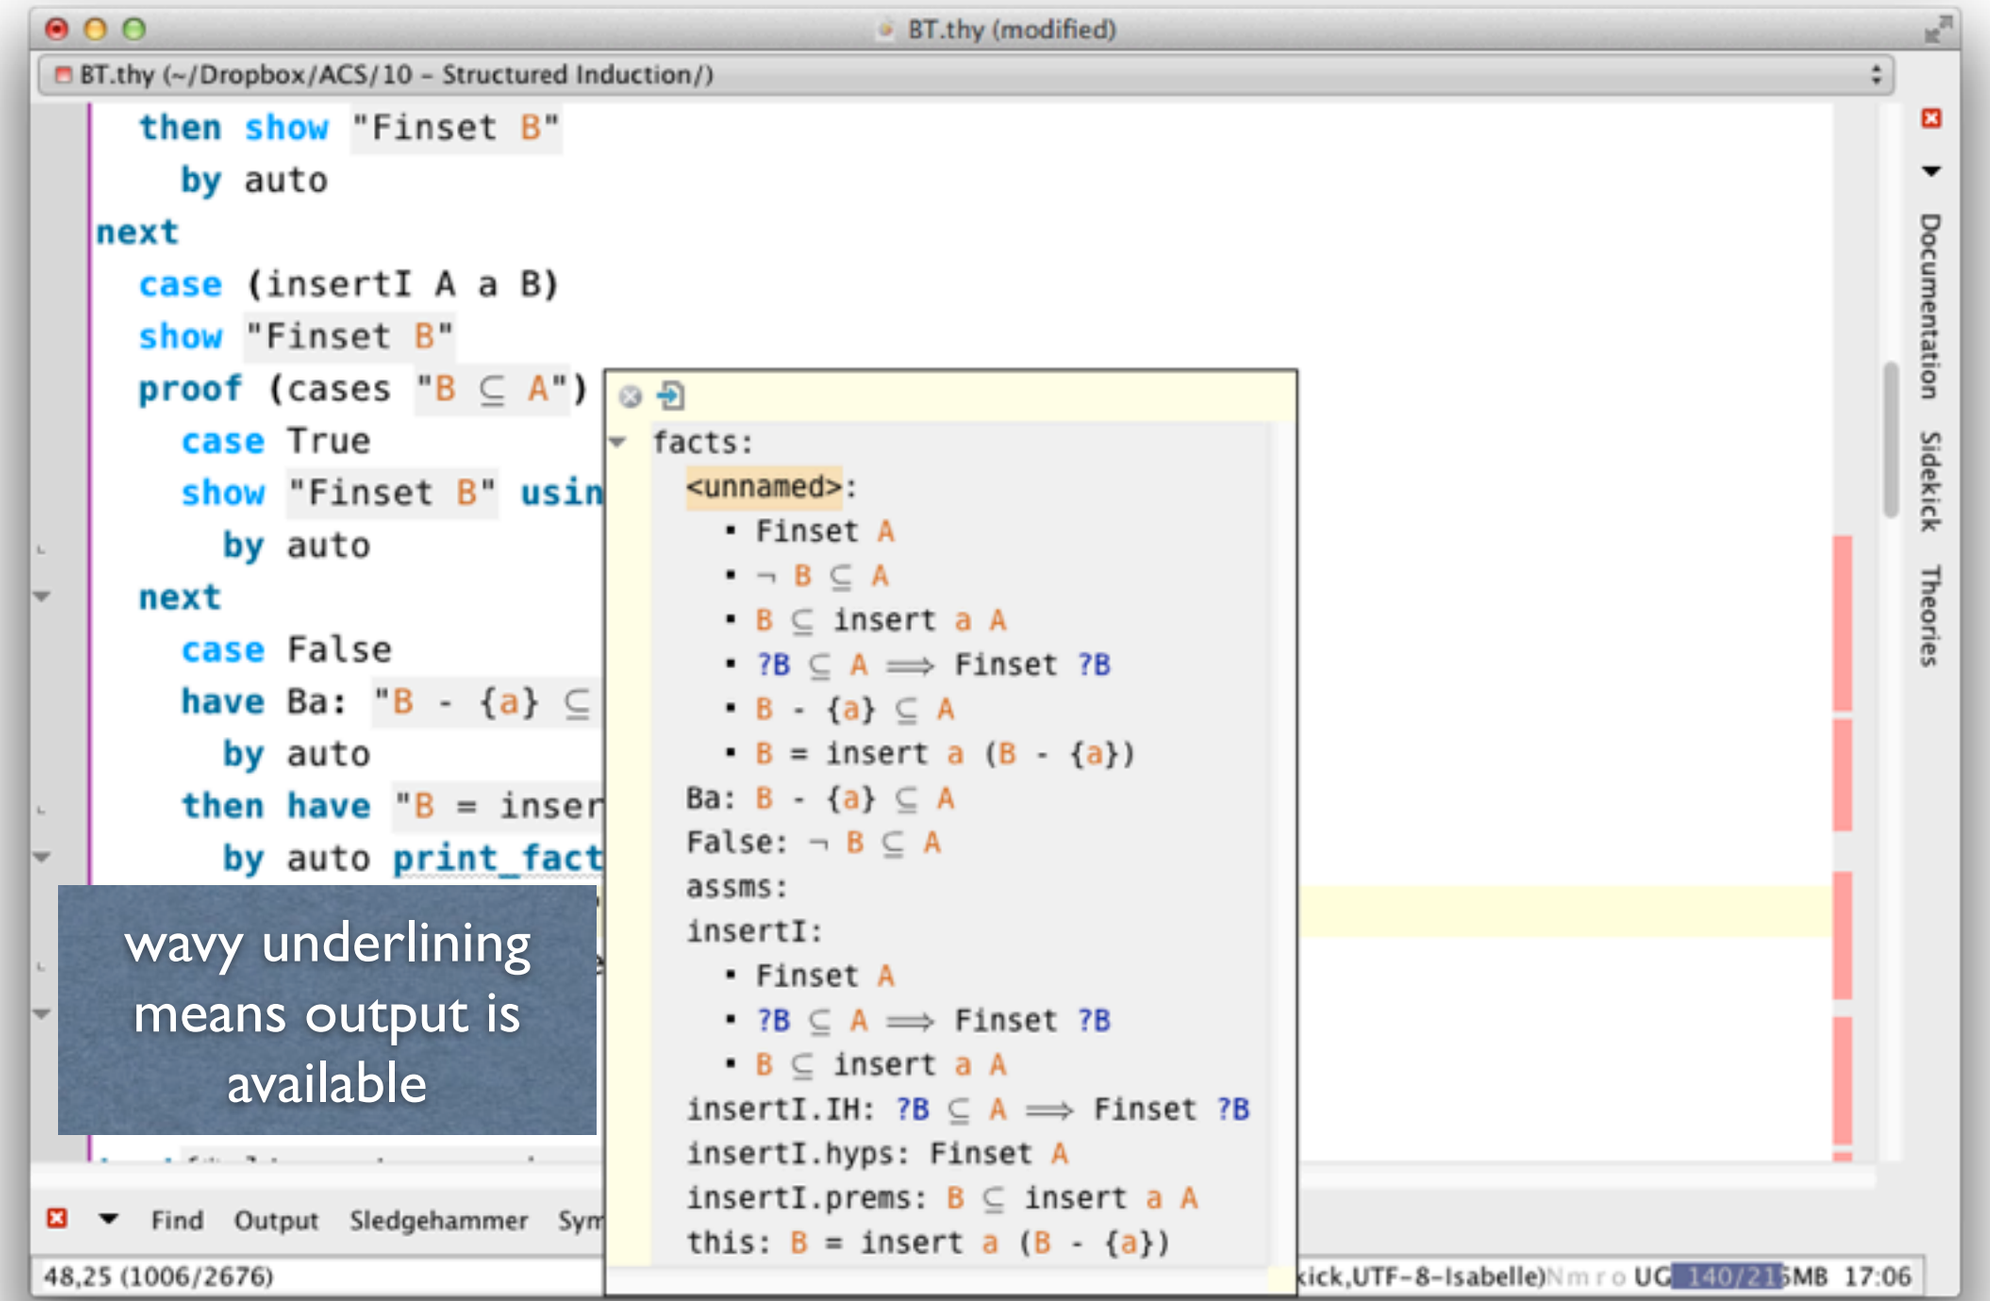

Simply hover with the mouse over any text where you see wavy underlining.

moreover / ultimately

```
notepad
begin

  have l1: "fact1" sorry
  have l2: "fact2" sorry
  have l3: "fact3" sorry
  from l1 l2 l3

oops
```

```
proof (chain): step 7

picking this:
  ▪ fact1
  ▪ fact2
  ▪ fact3
```

```
notepad
begin

  have "fact1" sorry
  moreover
  have "fact2" sorry
  moreover
  have "fact3" sorry
  ultimately

oops
```

```
calculation:
  ▪ fact1
  ▪ fact2
  ▪ fact3
proof (chain): step 9
```

These two keywords are useful when the conclusion is derived from a series of facts. The need for labels is eliminated (assuming that there are no other references to those facts) and the overall structure becomes much clearer. Here we also see the `notepad` construct, which is handy for typing in experimental proofs.

Existential Claims: "obtain"

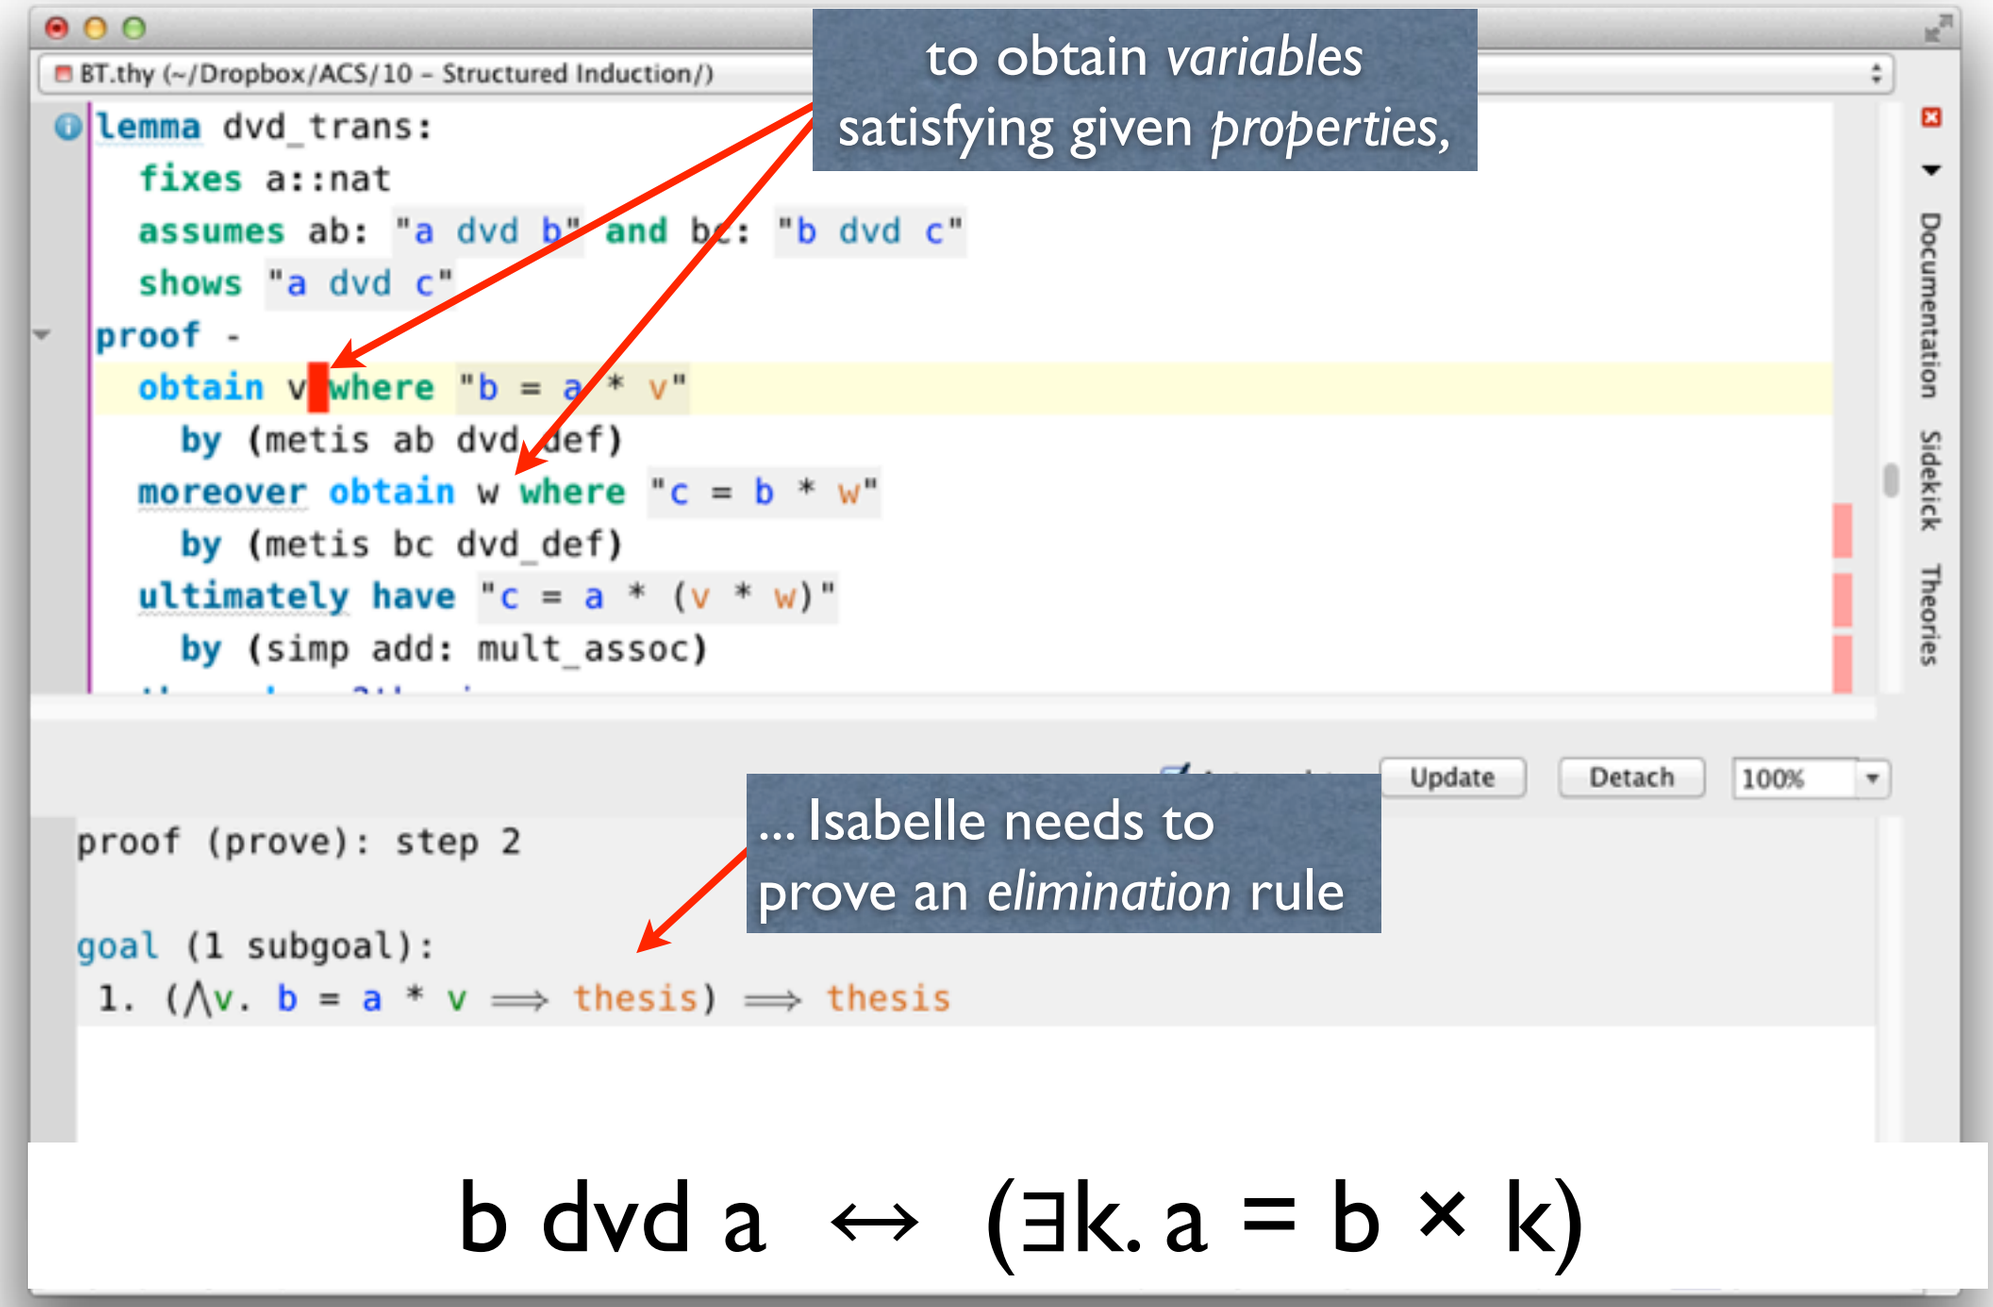

Frequently, our reasoning involves quantities (such as j above) that are known to satisfy certain properties. Here, the "divides" premise implies the existence of a divisor, j. Proof attempts involving "obtain" can be difficult to understand, especially when they fail. Isabelle proves a theorem having the general form of an elimination rule, which in the premise introduces one or more bound variables: the variables that we "obtain".

Chaining Facts: "moreover"

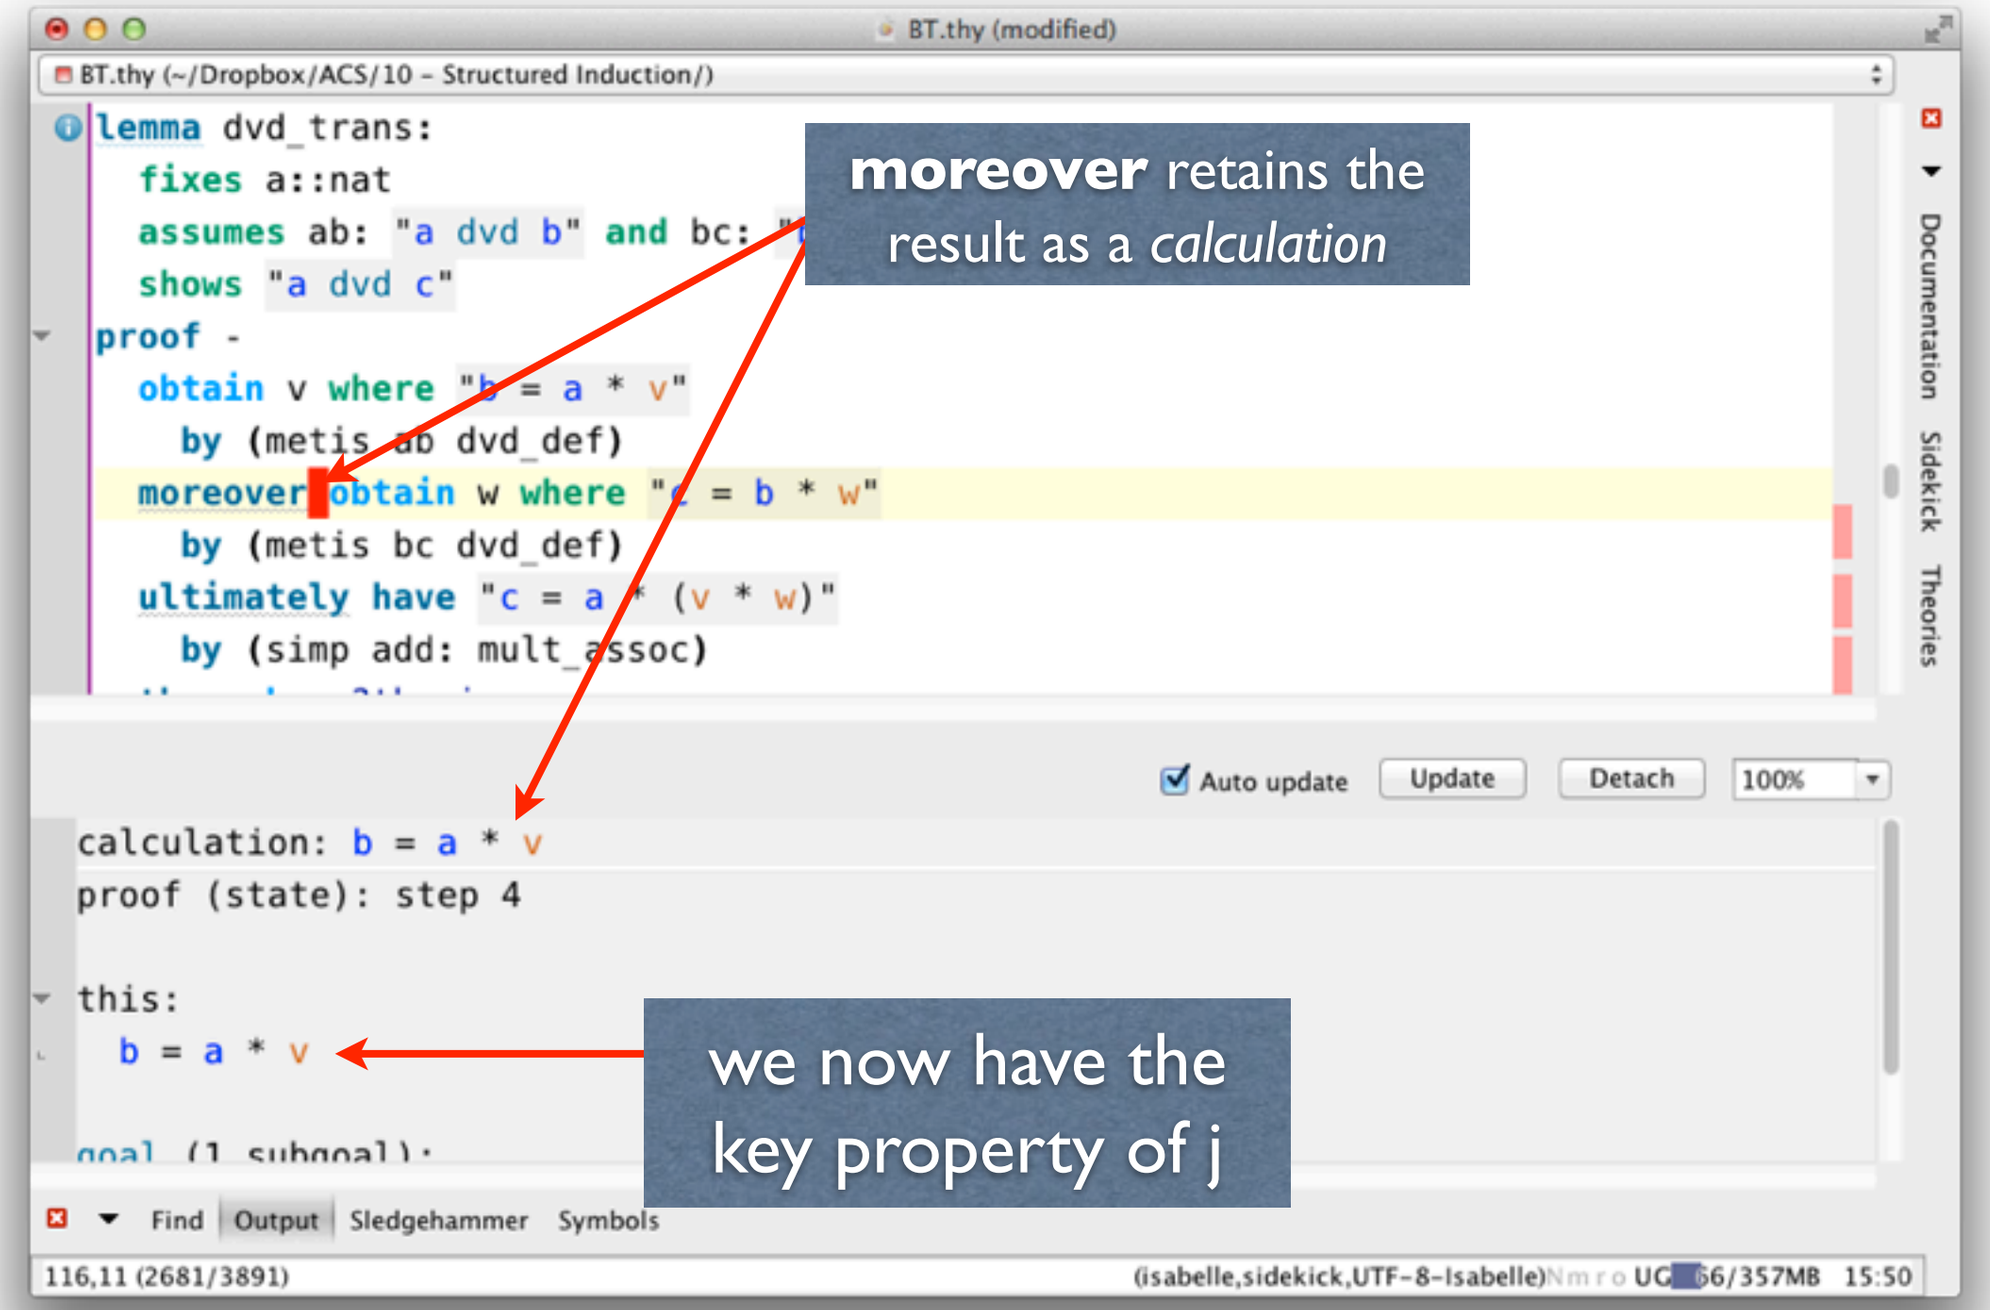

Delivering Facts: "ultimately"

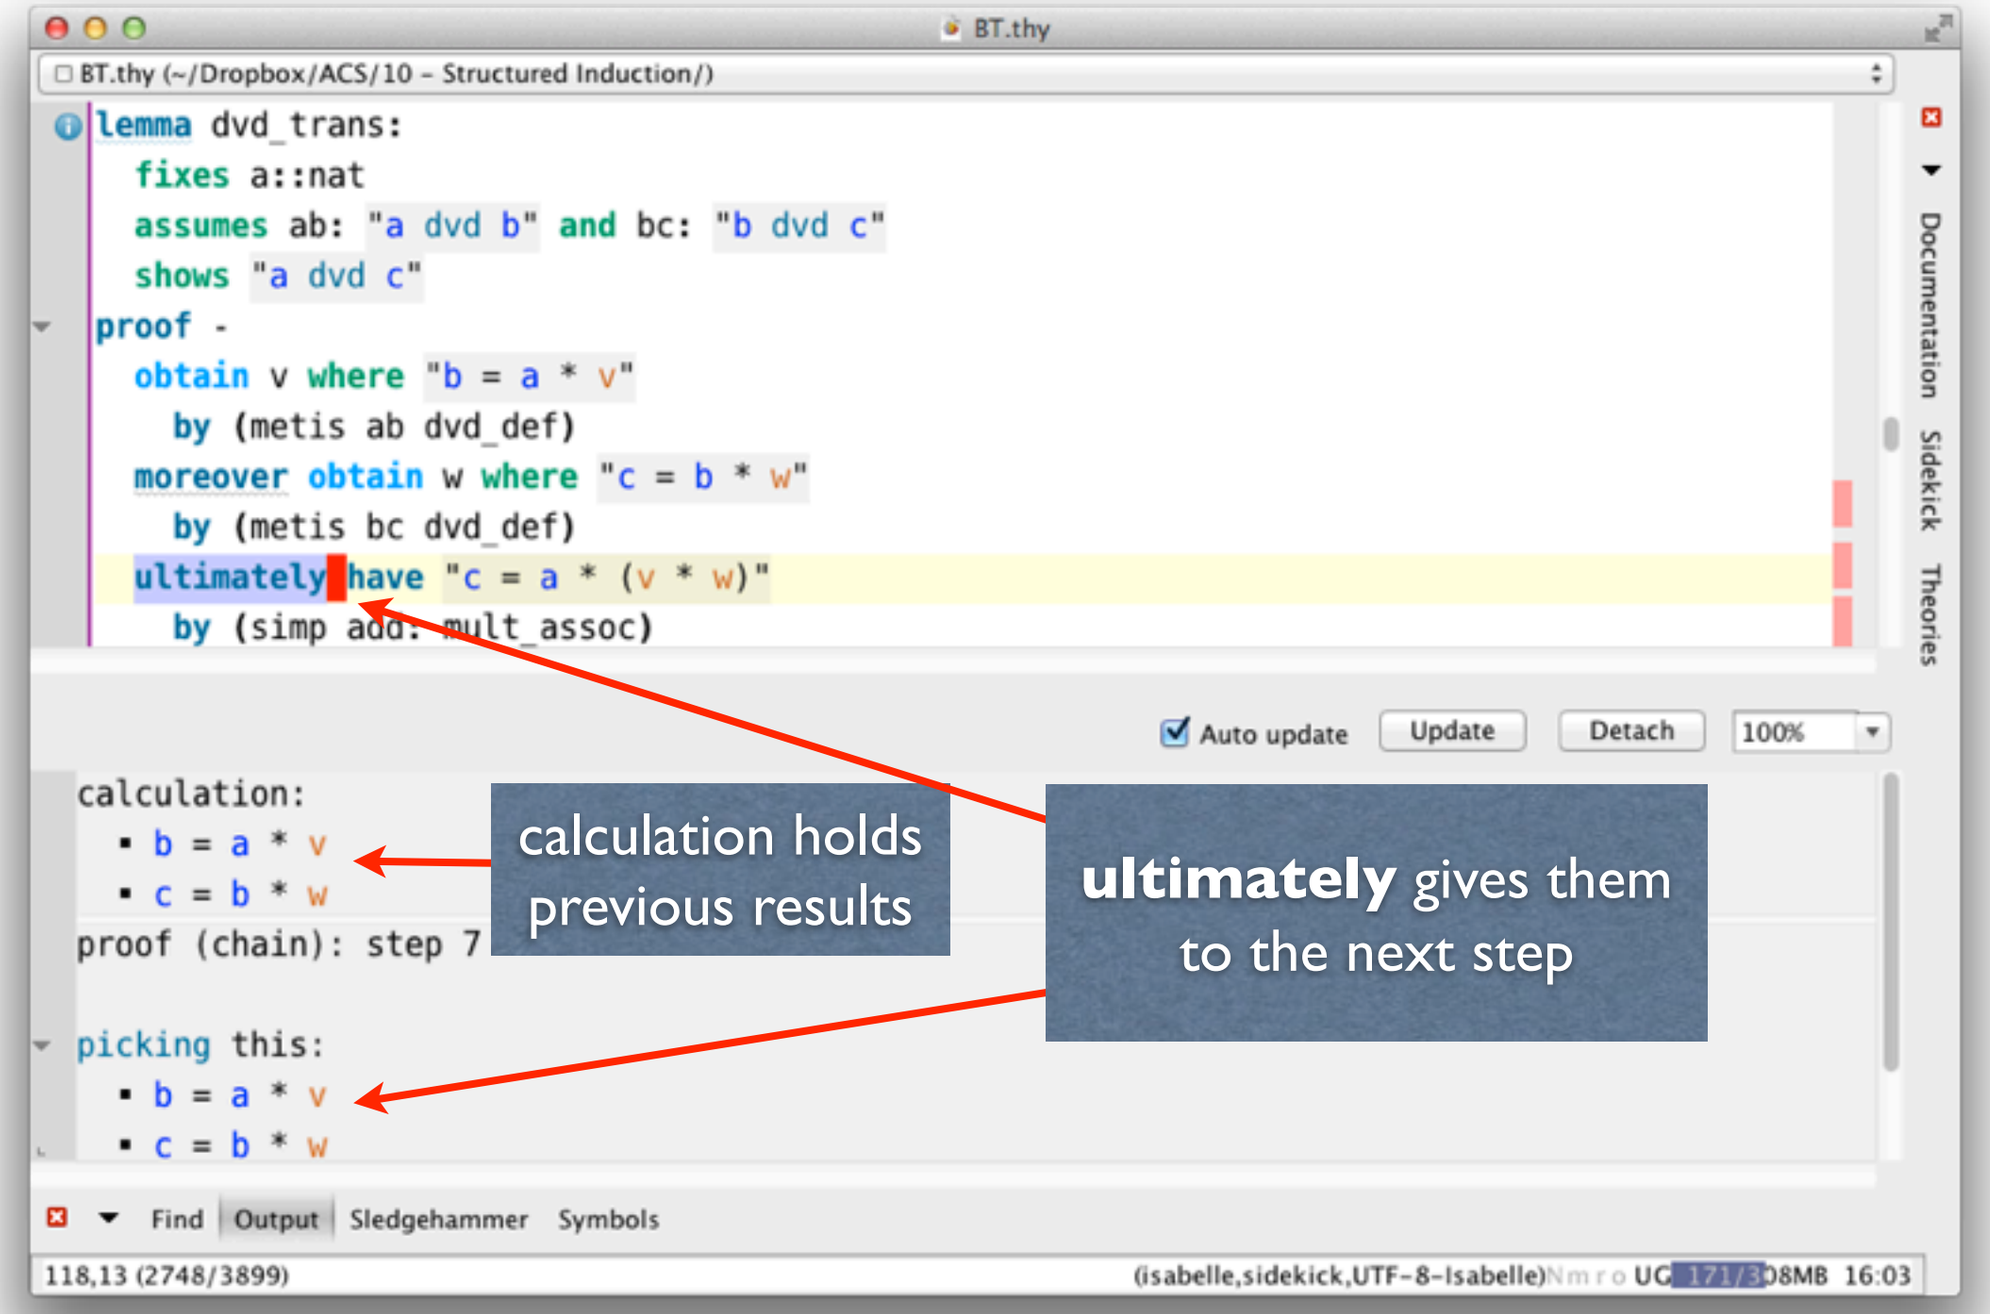

The Finished Proof

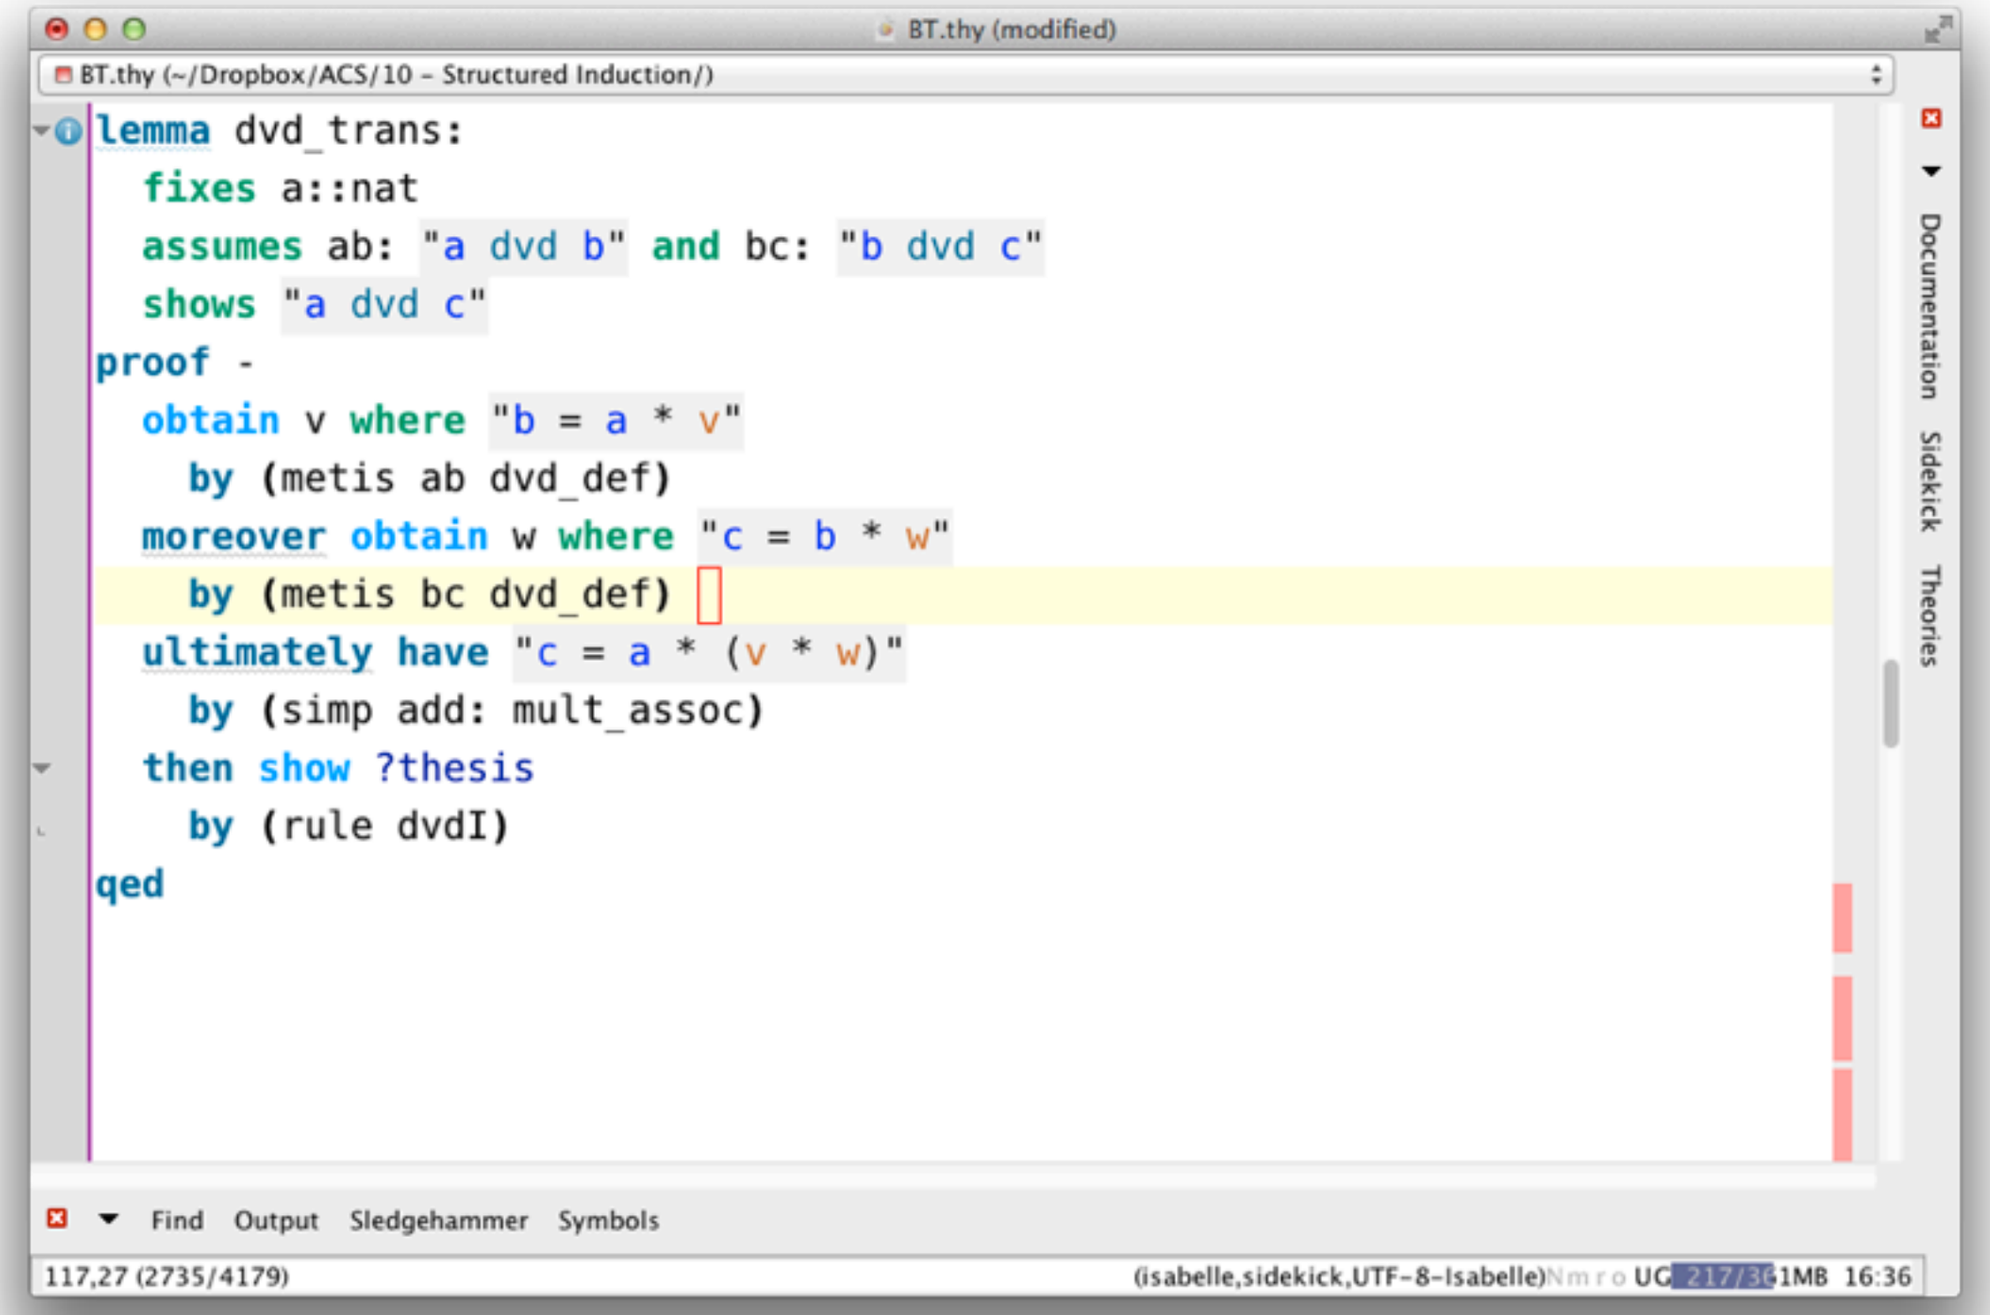

A Simpler Proof

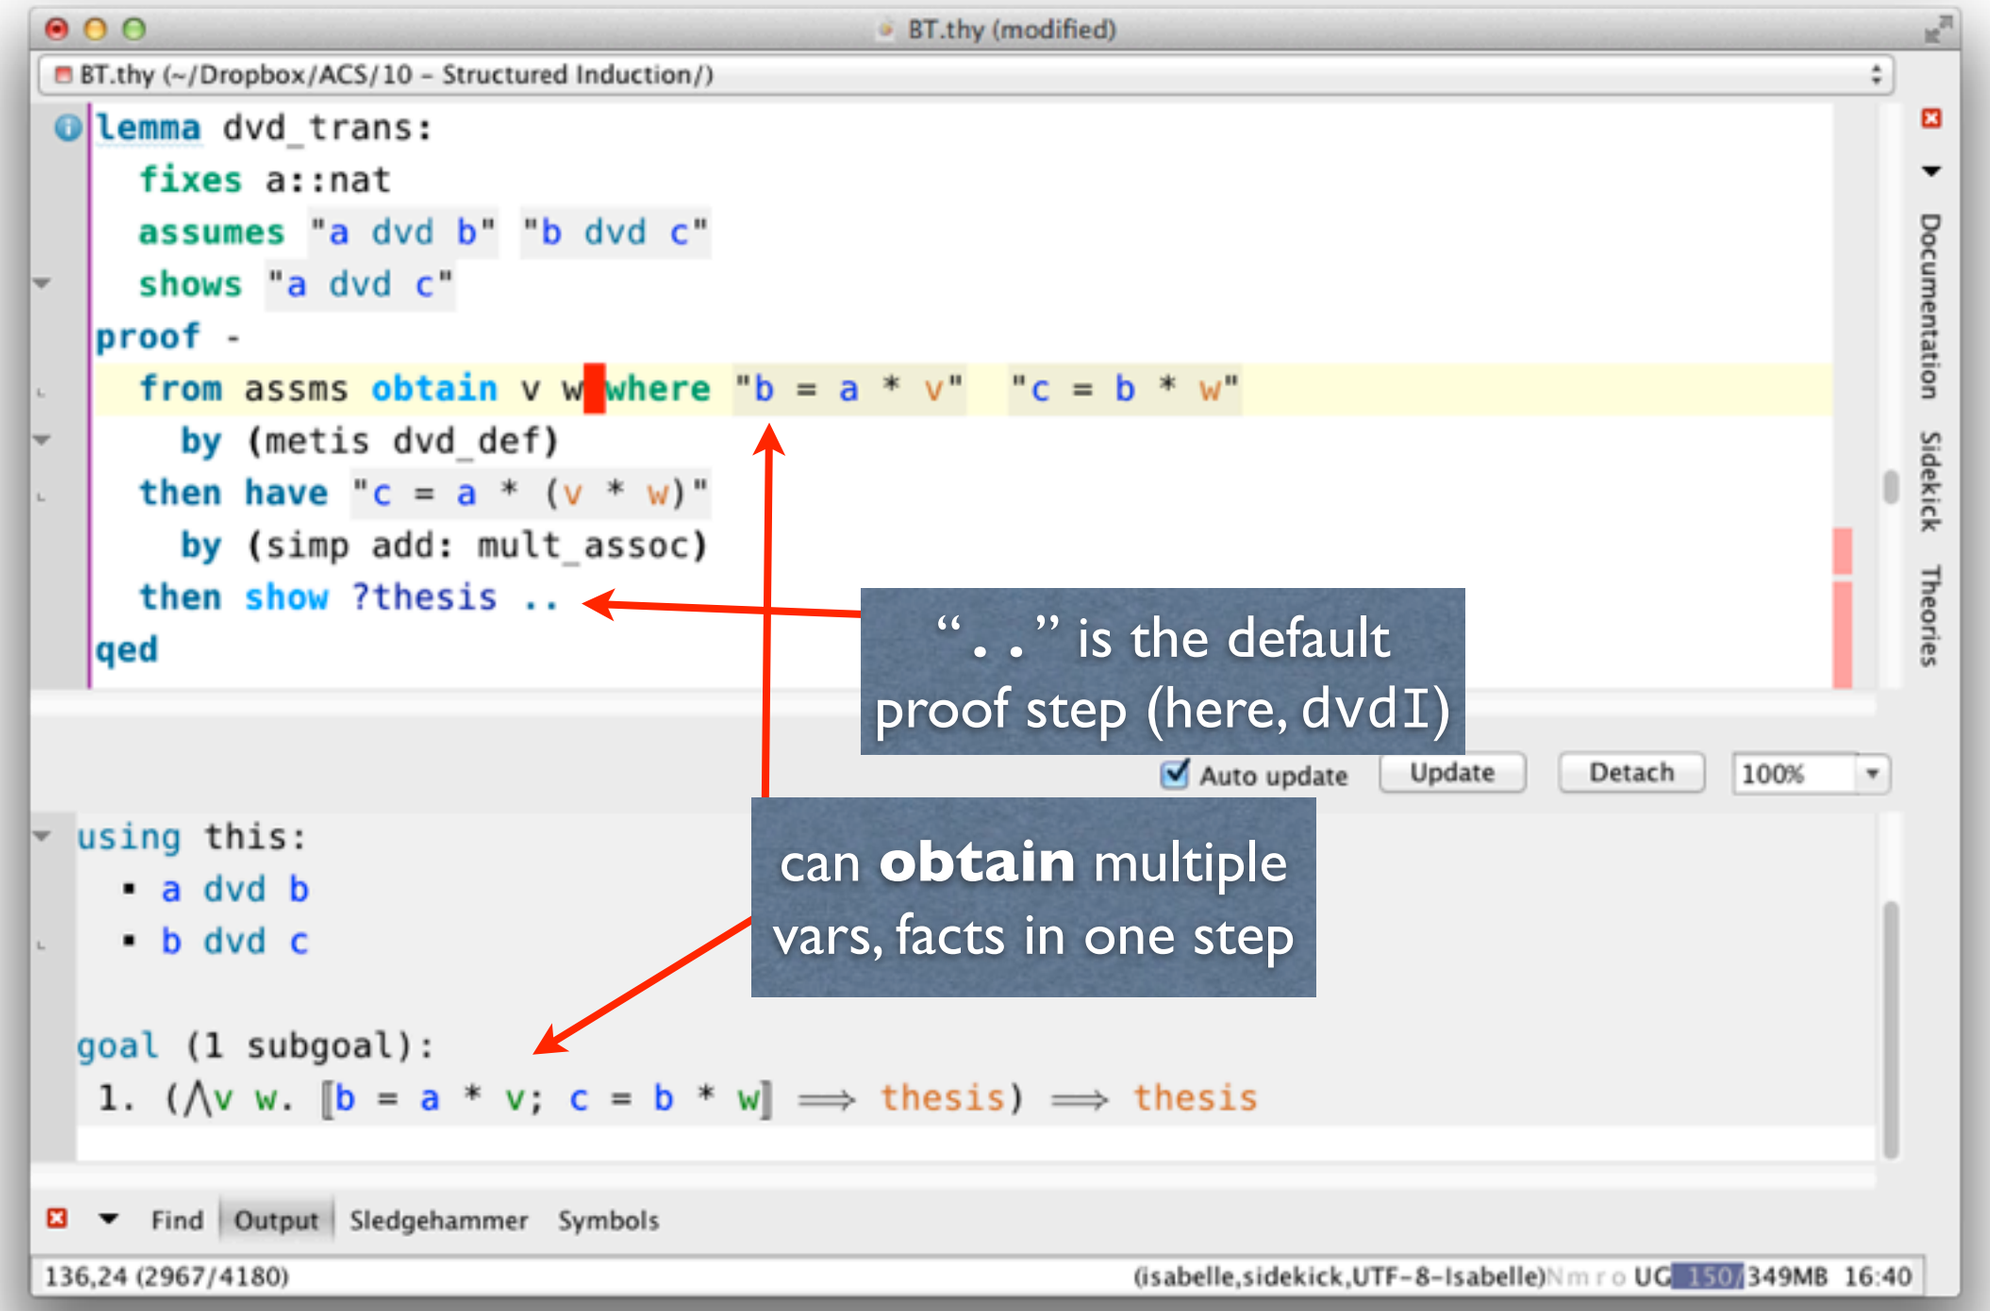

Any proof can be written in a variety of different ways. The concluding step is surprising. The mysterious `..` symbol denotes the default proof step, which in this case happens to be a rule called `dvdI`. This rule exactly matches the given premise and conclusion. In practice, however, default proof steps are seldom used.

Advanced Proof Structures

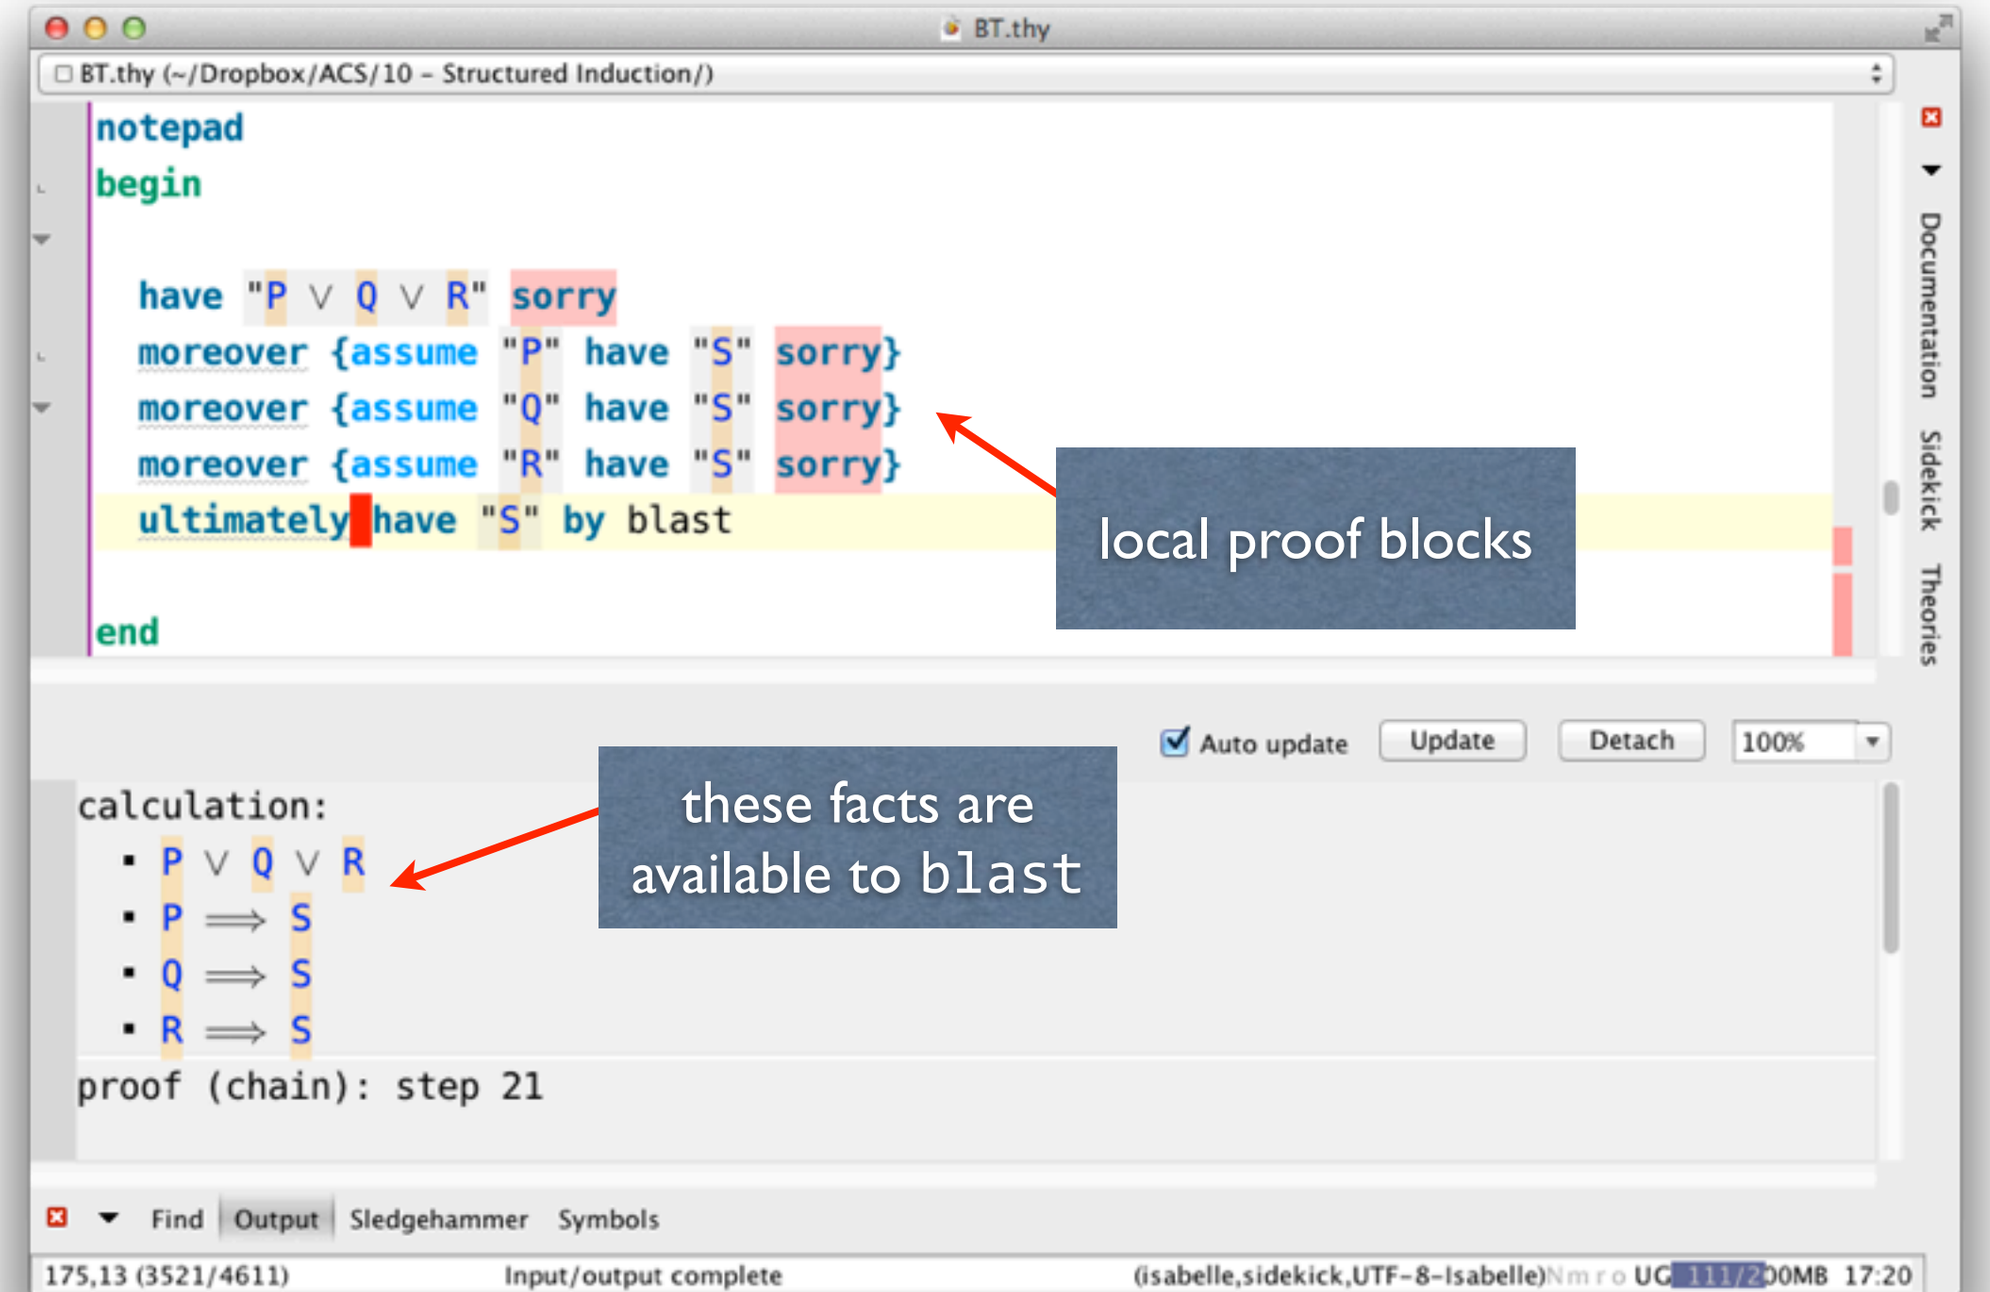

Here we see a three-way case distinction. Local blocks have many other uses.

Interactive Formal Verification
I I: Modelling Hardware

Lawrence C Paulson
Computer Laboratory
University of Cambridge

Outline

- General modelling techniques
- Hardware verification in higher-order logic
- Additional elements of the Isar language, for instantiating theorems

Basic Principles of Modelling

- Define *mathematical abstractions* of the objects of interest (systems, hardware, protocols,...).
- Whenever possible, use *definitions* — not axioms!
- Ensure that the abstractions capture enough detail.
- Unrealistic models have unrealistic properties.
- Inconsistent models will satisfy *all* properties.

All models involving the real world are *approximate*!

Constructing models using definitions exclusively is called the definitional approach. A purely definitional theory is guaranteed to be consistent. Axioms are occasionally necessary in abstract models, where the behaviour is too complex to be captured by definitions. However, a system of axioms can easily be inconsistent, which means that they imply every theorem. The most famous example of an inconsistent theory is Frege's, which was refuted by Russell's paradox. A surprising number of Frege's constructions survived this catastrophe. Nevertheless, an inconsistent theory is almost worthless.

Useful models are abstract, eliminating unnecessary details in order to focus on the crucial points. The frictionless surfaces and pulleys found in school physics problems are a well-known example of abstraction. Needless to say, the real world is not frictionless and this particular model is useless for understanding everyday physics such as walking. But even models that introduce friction use abstractions, such as the assumption that the force of friction is linear, which cannot account for such phenomena as slipping on ice. Abstraction is always necessary in models of the real world, with its unimaginable complexity; it is often necessary even in a purely mathematical context if the subject material is complicated.

Hardware Verification

- Pioneered by Prof. M. J. C. Gordon and his students, using successive versions of the HOL system.

- Used to model substantial hardware designs:

  - VIPER chip verification, by Avra Cohn (1988)

  - The ARM6 processor, by Anthony Fox (2003)

- Works *hierarchically* from arithmetic units and memories right down to flip-flops and transistors.

- Crucially uses *higher-order logic*, modelling signals as boolean-valued functions over time.

The material in this lecture is based on Prof Gordon's lecture notes for *Specification and Verification II*, which are available on the Internet at `http://www.cl.cam.ac.uk/~mjcg/Teaching/SpecVer2/`

Devices as Relations

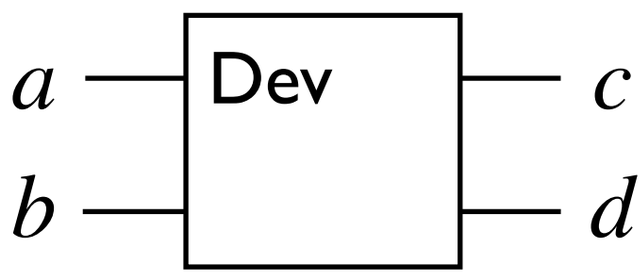

A *relation* in a, b, c, d

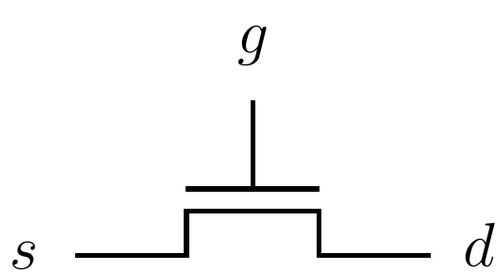

$$g \rightarrow s = d$$

The relation describes the possible combinations of values on the ports.

Values could be bits, words, signals (functions from time to bits), etc

The second device on the slide above is an N-type field effect transistor, which can be conceived as a switch: when the gate goes high, the source and drain are connected. The logical implication shown next to the transistor formalises this behaviour. Note that the connection between the source and drain is *bidirectional*, with no suggestion that information flows from one port to the other.

Relational Composition

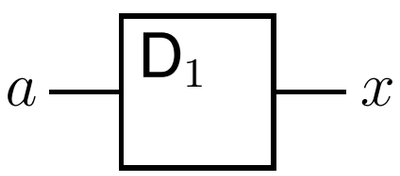

$\mathbf{S}_1[a, x]$

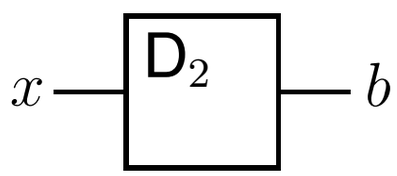

$\mathbf{S}_2[x, b]$

two devices modelled by two formulas

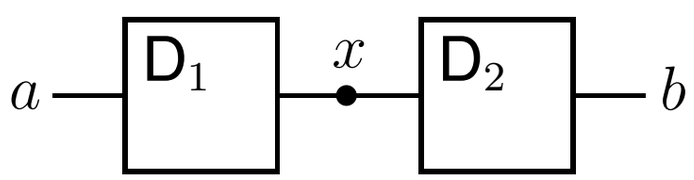

$\mathbf{S}_1[a, x] \wedge \mathbf{S}_2[x, b]$

the connected ports have the *same* value

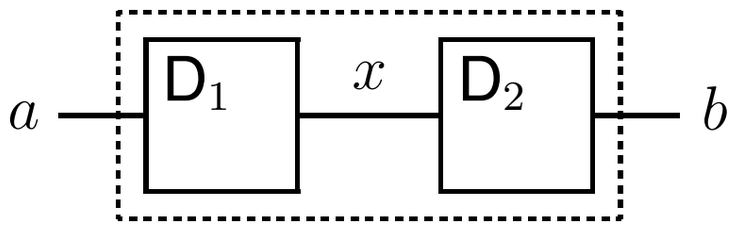

$\exists x.\, \mathbf{S}_1[a, x] \wedge \mathbf{S}_2[x, b]$

the connected ports have *some* value

The diagrams are taken from Prof Gordon's lecture notes.

Because we model devices by relations, connecting devices together must be modelled by relational composition. Syntactically, we specify circuits by logical terms that denote relations and we express relational composition using the existential quantifier. The quantifier creates a local scope, thereby hiding the internal wire.

Specifications and Correctness

- The *implementation* of a device in terms of other devices can be expressed by composition.
- The *specification* of the device's intended behaviour can be given by an abstract formula.
- Sometimes the implementation and specification can be proved *equivalent*: *Imp*⇔*Spec*.
- The property *Imp*⇒*Spec* ensures that every possible behaviour of the *Imp* is permitted by *Spec*.

Impossible implementations satisfy *all* specifications!

The implementation describes a circuit, while the specification should be based on mathematical definitions that were established prior to the implementation. A limitation of this approach is that impossible implementations can be expressed: in the most extreme case, implementations that identify the values true and false. In hardware, this represents a short circuit connecting power to ground, possibly a short circuit that only occurs when a particular combination of values appears on other wires, activating an unfortunate series of transistors. In the real world, short circuits have catastrophic effects, while in logic, identifying true with false allows anything to be proved. Therefore, absence of short circuits needs to be established somehow if this relational approach is to be used safely.

For combinational circuits (those without time), both the implementation and the specification express truth tables with no concept of a "don't care" entry, so logical equivalence should be provable. Sequential circuits involve time, and frequently the specification samples the clock only a specific intervals, ignoring the situation otherwise. Specifications can involve many other forms of abstraction. In general, we cannot expect to prove logical equivalence.

Proving the logical equivalence of the implementation with the specification does not prove the absence of short circuits, but it does prove that the short circuits coincide with inconsistencies in the specification itself. Needless to say, a correct specification should be free of inconsistencies, but there is no way in general to guarantee this. How then do we benefit from using logic? Specifications tend to be much simpler than implementations and they are less likely to contain errors. Moreover, the attempt to prove properties relating specifications and implementations frequently identifies errors, even if we cannot promise all embracing guarantees.

The *Switch Model* of CMOS

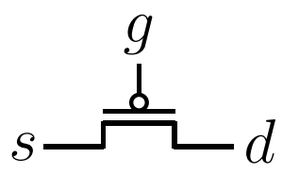

Ptran$(g, s, d) = (\neg g \Rightarrow (d = s))$

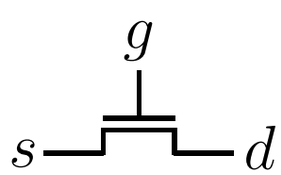

Ntran$(g, s, d) = (g \Rightarrow (d = s))$

Gnd $g = (g =$ F$)$

Pwr $p = (p =$ T$)$

```
subsection{* Specification of CMOS primitives *}

text{* P and N transistors *}
definition "Ptran = (λ(g,a,b). (~g ⟶ a = b))"
definition "Ntran = (λ(g,a,b). (g ⟶ a = b))"

text{* Power and Ground*}
definition "Pwr p = (p = True)"
definition "Gnd p = (p = False)"
```

CMOS (complementary metal oxide semiconductor) technology combines P- and N-type transistors on a chip to make gates and other devices. The slide shows primitive concepts: the two types of transistors, ground (modelled by the value False) and power (model by the value True). The corresponding Isabelle definitions are easily expressed. Lambda-notation is a convenient way to express a function is argument is a triple.

Full Adder: Specification

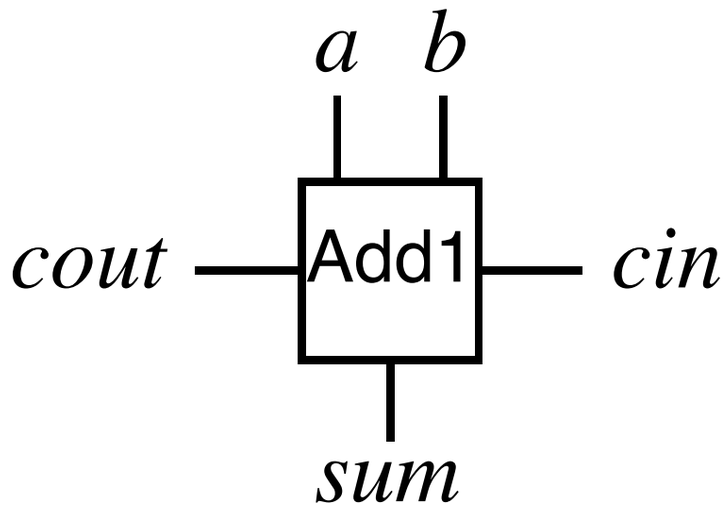

$$2 \times cout + sum = a + b + cin$$

```
text{* 1-bit full adder specification *}

text{* Convert boolean to number (0 or 1) *}
definition bit_val :: "bool ⇒ nat"  where
  "bit_val p = (if p then 1 else 0)"

definition "Add1Spec = (λ(a,b,cin,sum,cout).
                2*(bit_val cout) + bit_val sum =
                bit_val a + bit_val b + bit_val cin)"
```

A full adder forms the sum of three one-bit inputs, yielding a two-bit result. The higher-order output bit is called “carry out”, and it will typically be connected to the “carry in” of the next stage. Because we typically use True and False to designate hardware bit values, the obvious conversion to 1 and 0 is necessary in order to express arithmetic properties. Even with this small step, expressing the specification in higher-order logic is trivial. The identifier denotes the abstract relation satisfied by a full adder, namely the legal combinations of values on the various ports.

Full Adder: Implementation

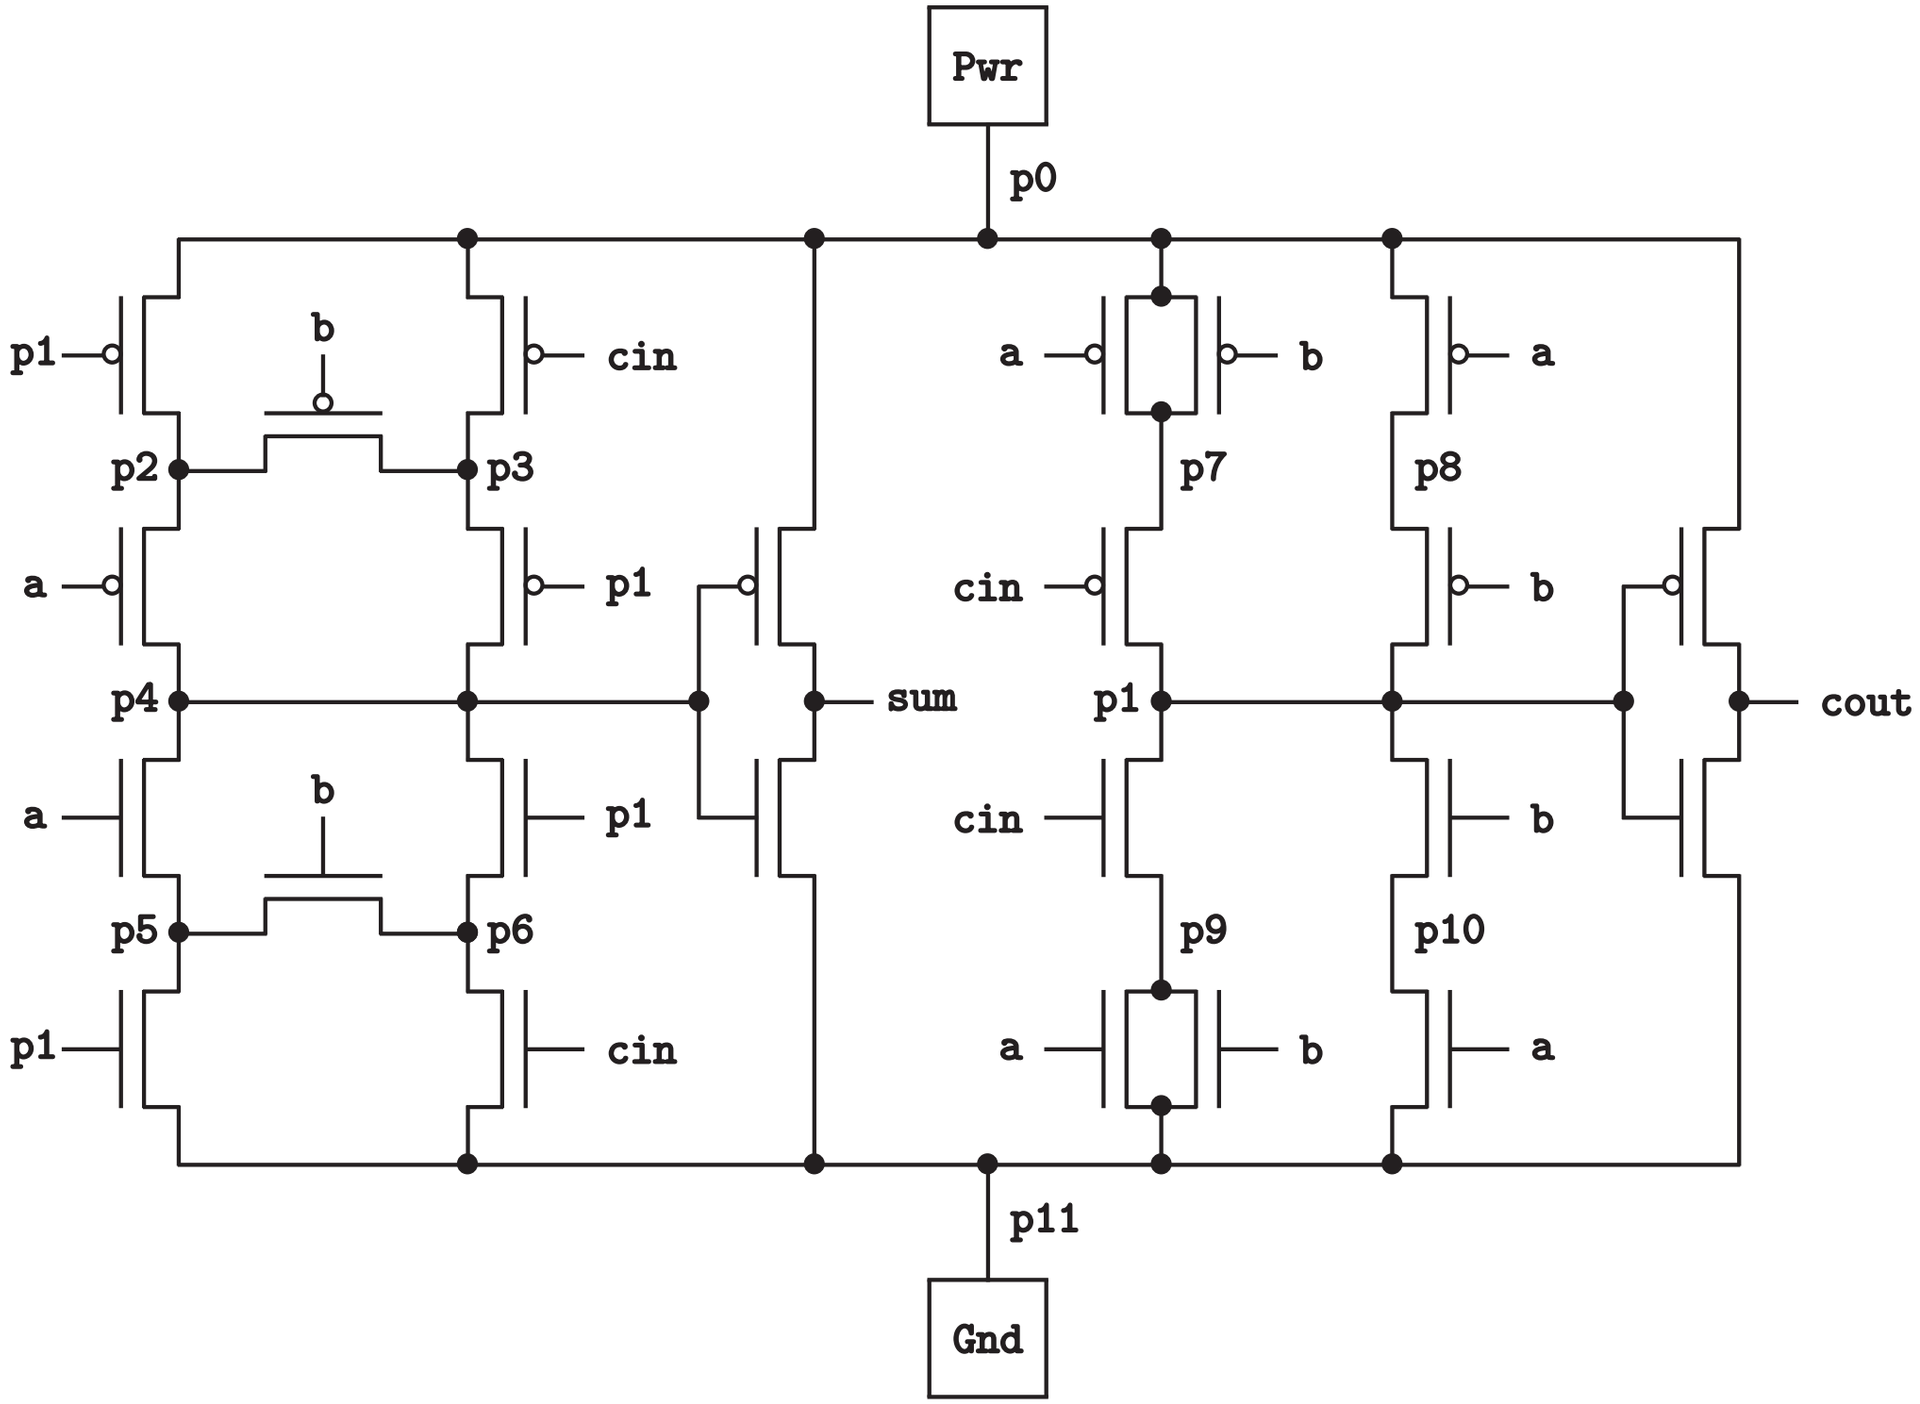

A full adder is easily expressed at the gate level in terms of exclusive-OR (to compute the `sum`) and other simple gating to compute the carry. The diagram above, again from Prof Gordon's notes, expresses a full adder as would be implemented directly in terms of transistors.

Full Adder in Isabelle

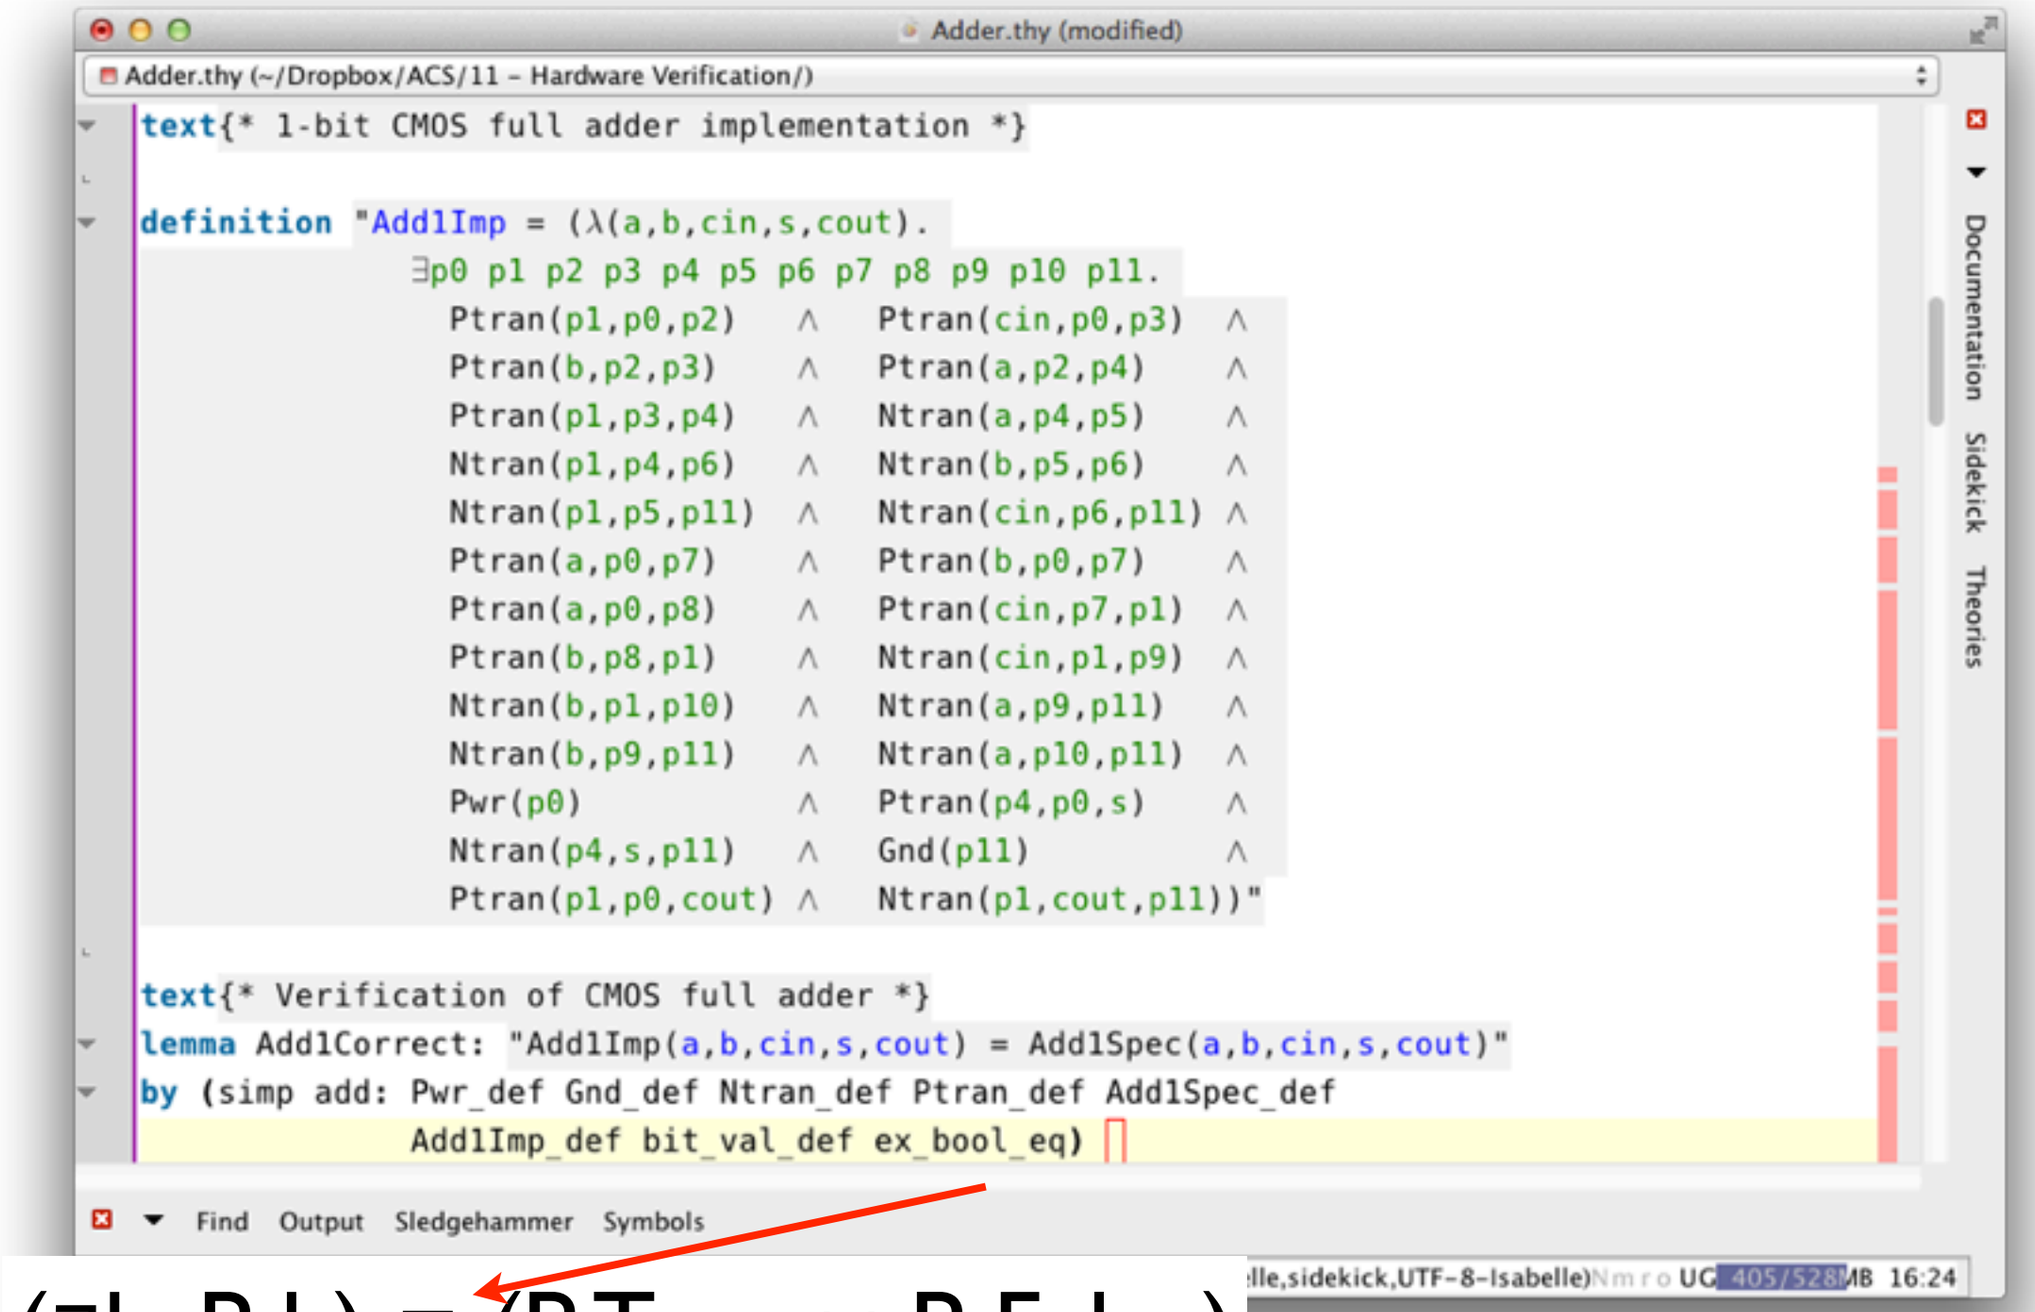

The logical formula above is a direct translation of the diagram on the previous slide. Needless to say, the translation from diagram to formula should ideally be automatic, and better still, driven by the same tools that fabricate the actual chip.

The theorem expresses the logical equivalence between the implementation (in terms of transistors) and the specification (in terms of arithmetic). This type of proof is trivial for reasoning tools based on BDDs or SAT solvers. Isabelle is not ideal for such proofs, and this one requires over four seconds of CPU time. In the simplifier call, the last theorem named is crucial, because it forces a case split on every existentially quantified wire.

An *n*-bit Ripple-Carry Adder

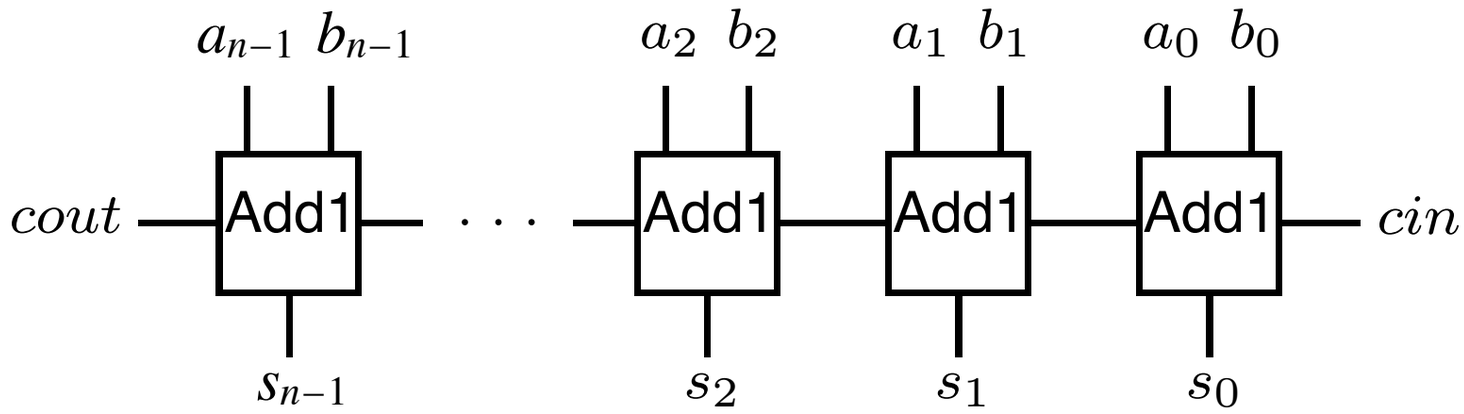

$$(2^n \times cout) + s = a + b + cin$$

- Cascading several full adders yields an *n*-bit adder.
- The implementation is expressed recursively.
- The specification is obvious mathematics.

Adder Specification

$$(2^n \times cout) + s = a + b + cin$$

values of n-bit words

```
text{* Unsigned number denoted by bitstring f(n-1)...f(0) *}

fun bits_val where
  "bits_val f 0       = 0"
| "bits_val f (Suc n) = 2^n * bit_val(f n) + bits_val f n"

text{* Specification of an n-bit adder *}

definition
  "AdderSpec n = (λ(a, b, cin, sum, cout).
    2^n * bit_val cout + bits_val sum n =
    bits_val a n + bits_val b n + bit_val cin)"
```

The function `bits_val` converts a binary numeral (supplied in the form of a boolean valued function, `f`) to a non-negative integer.The specification of the adder then follows the obvious arithmetic specification closely. When n=0, the specification merely requires cin=cout.

Adder Implementation

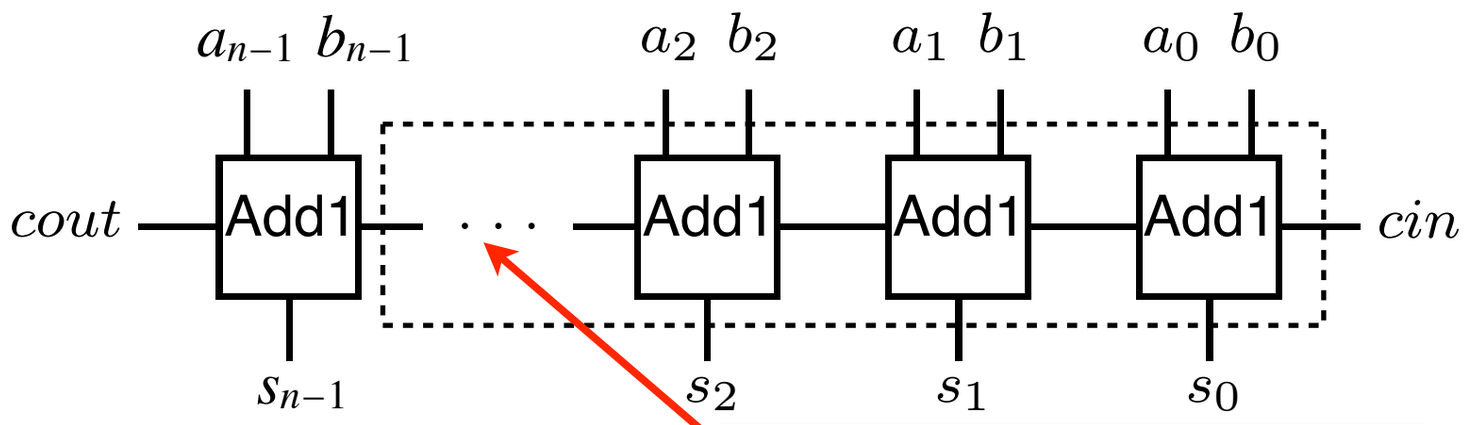

internal wire, to be hidden

```
text{* Implementation of an n-bit ripple-carry adder*}

fun AdderImp where
  "AdderImp 0 (a, b, cin, sum, cout) = (cout = cin)"
| "AdderImp (Suc n) (a, b, cin, sum, cout) =
    (∃c. AdderImp n (a, b, cin, sum, c) ∧
         Add1Imp (a n, b n, c, sum n, cout))"
```

a zero-bit adder simply connects the carry lines!

An (n+1)-bit adder consists of a full adder connected to an n-bit adder. Note that `AdderImp n` specifies an n-bit adder, and in particular, a 0-bit adder is nothing but a wire connecting carry in to carry out.

Partial Correctness Proof

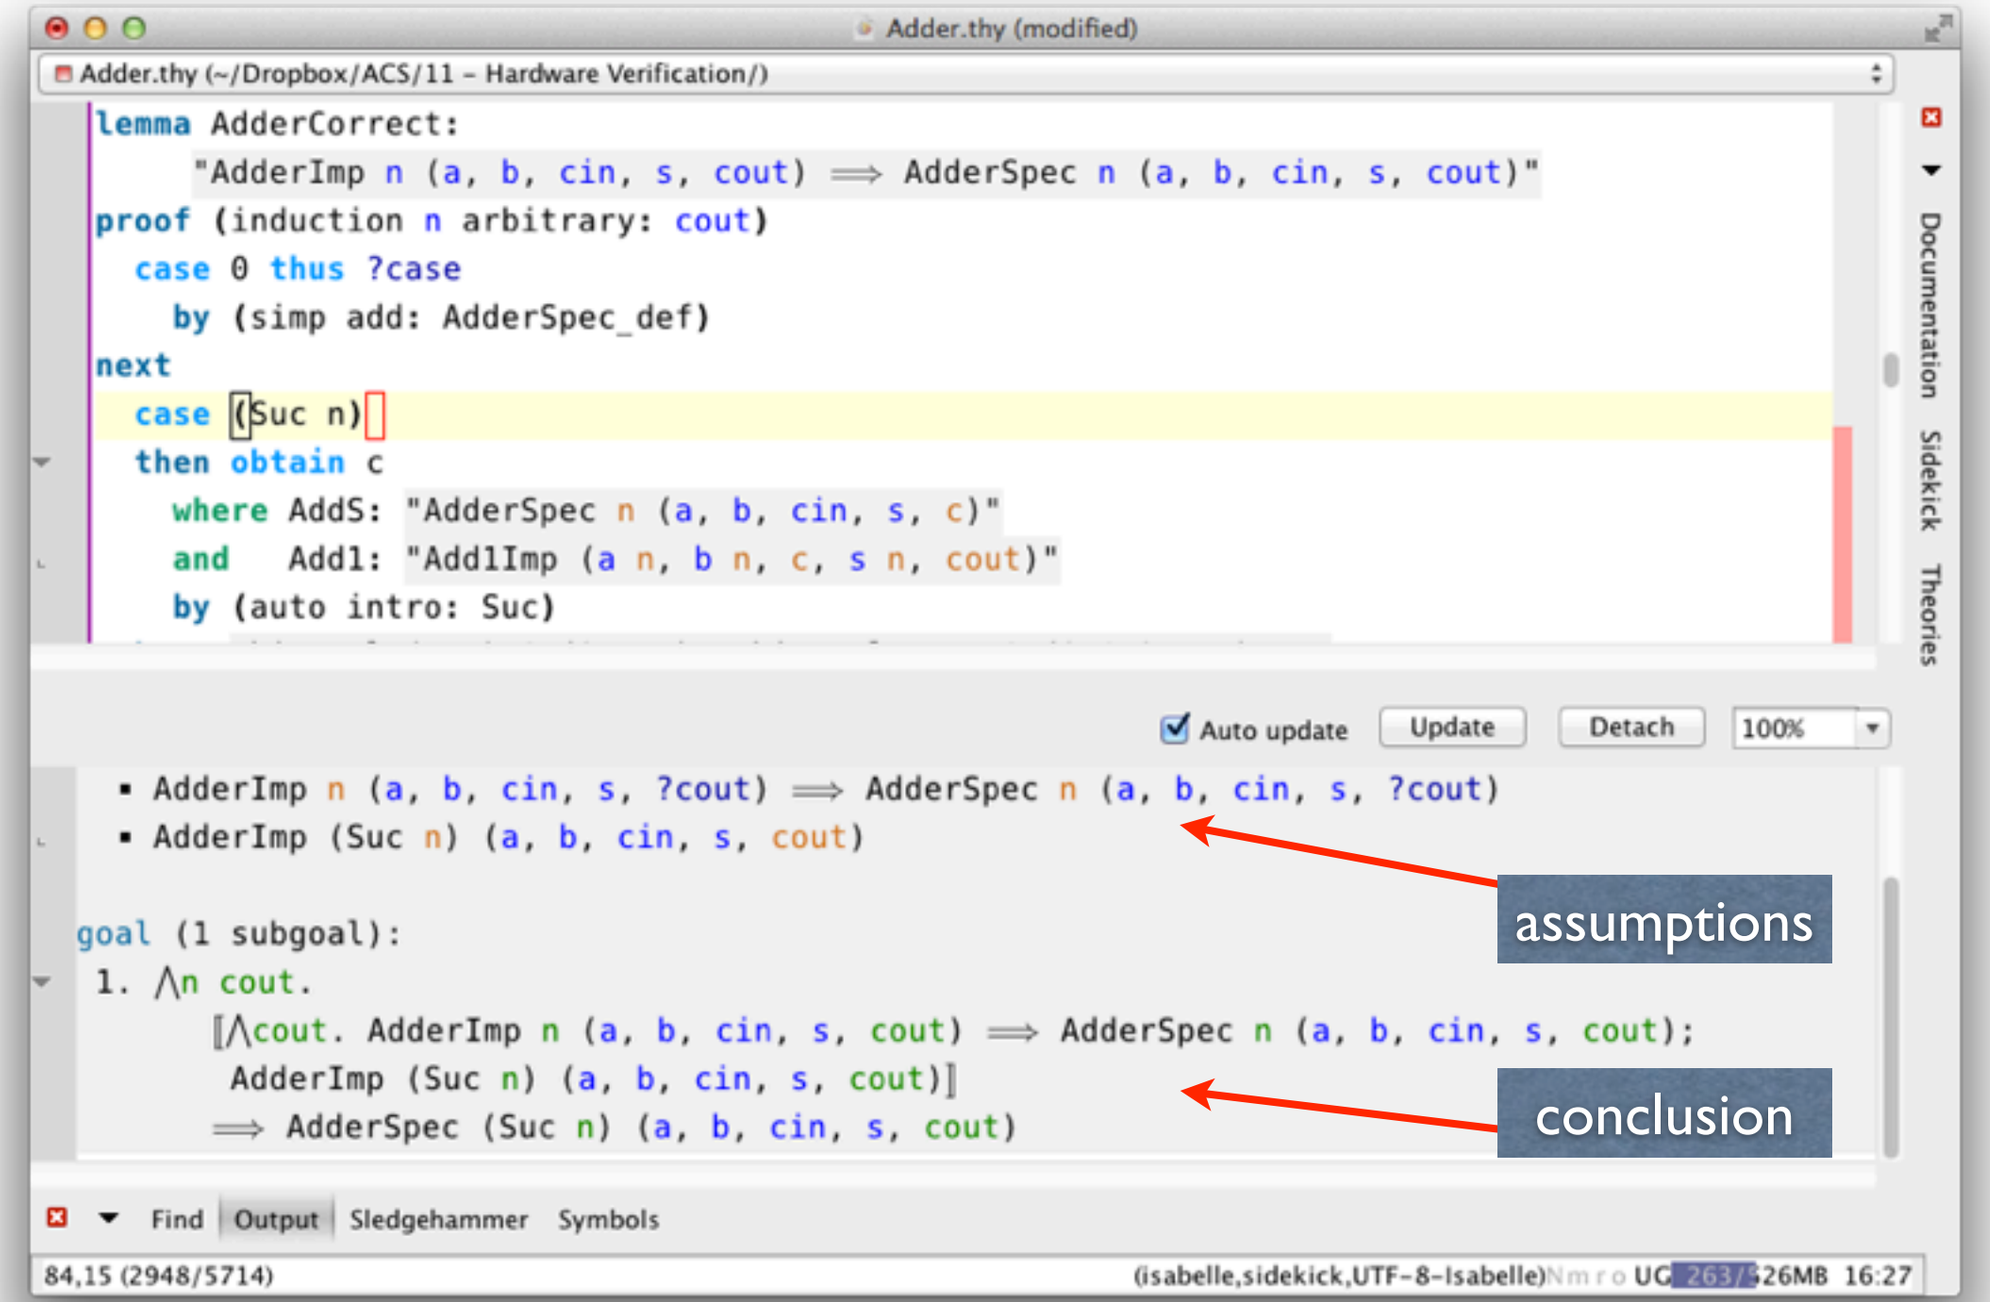

We are proving *partial correctness* only: that the implementation implies the specification. The term “partial correctness” here refers to a limitation of the approach, namely that an inconsistent implementation (one with short circuits) can imply any specification. Termination, obviously, plays no role in this circuit.

The base case is trivial. Our task in the induction step Is shown on the slide. It is expressed in terms of predicates for the implementation and specification. The induction hypothesis asserts that the implementation implies the specification for n. We now assume the implementation for n+1 and must prove the corresponding specification.

Using the Induction Hypothesis

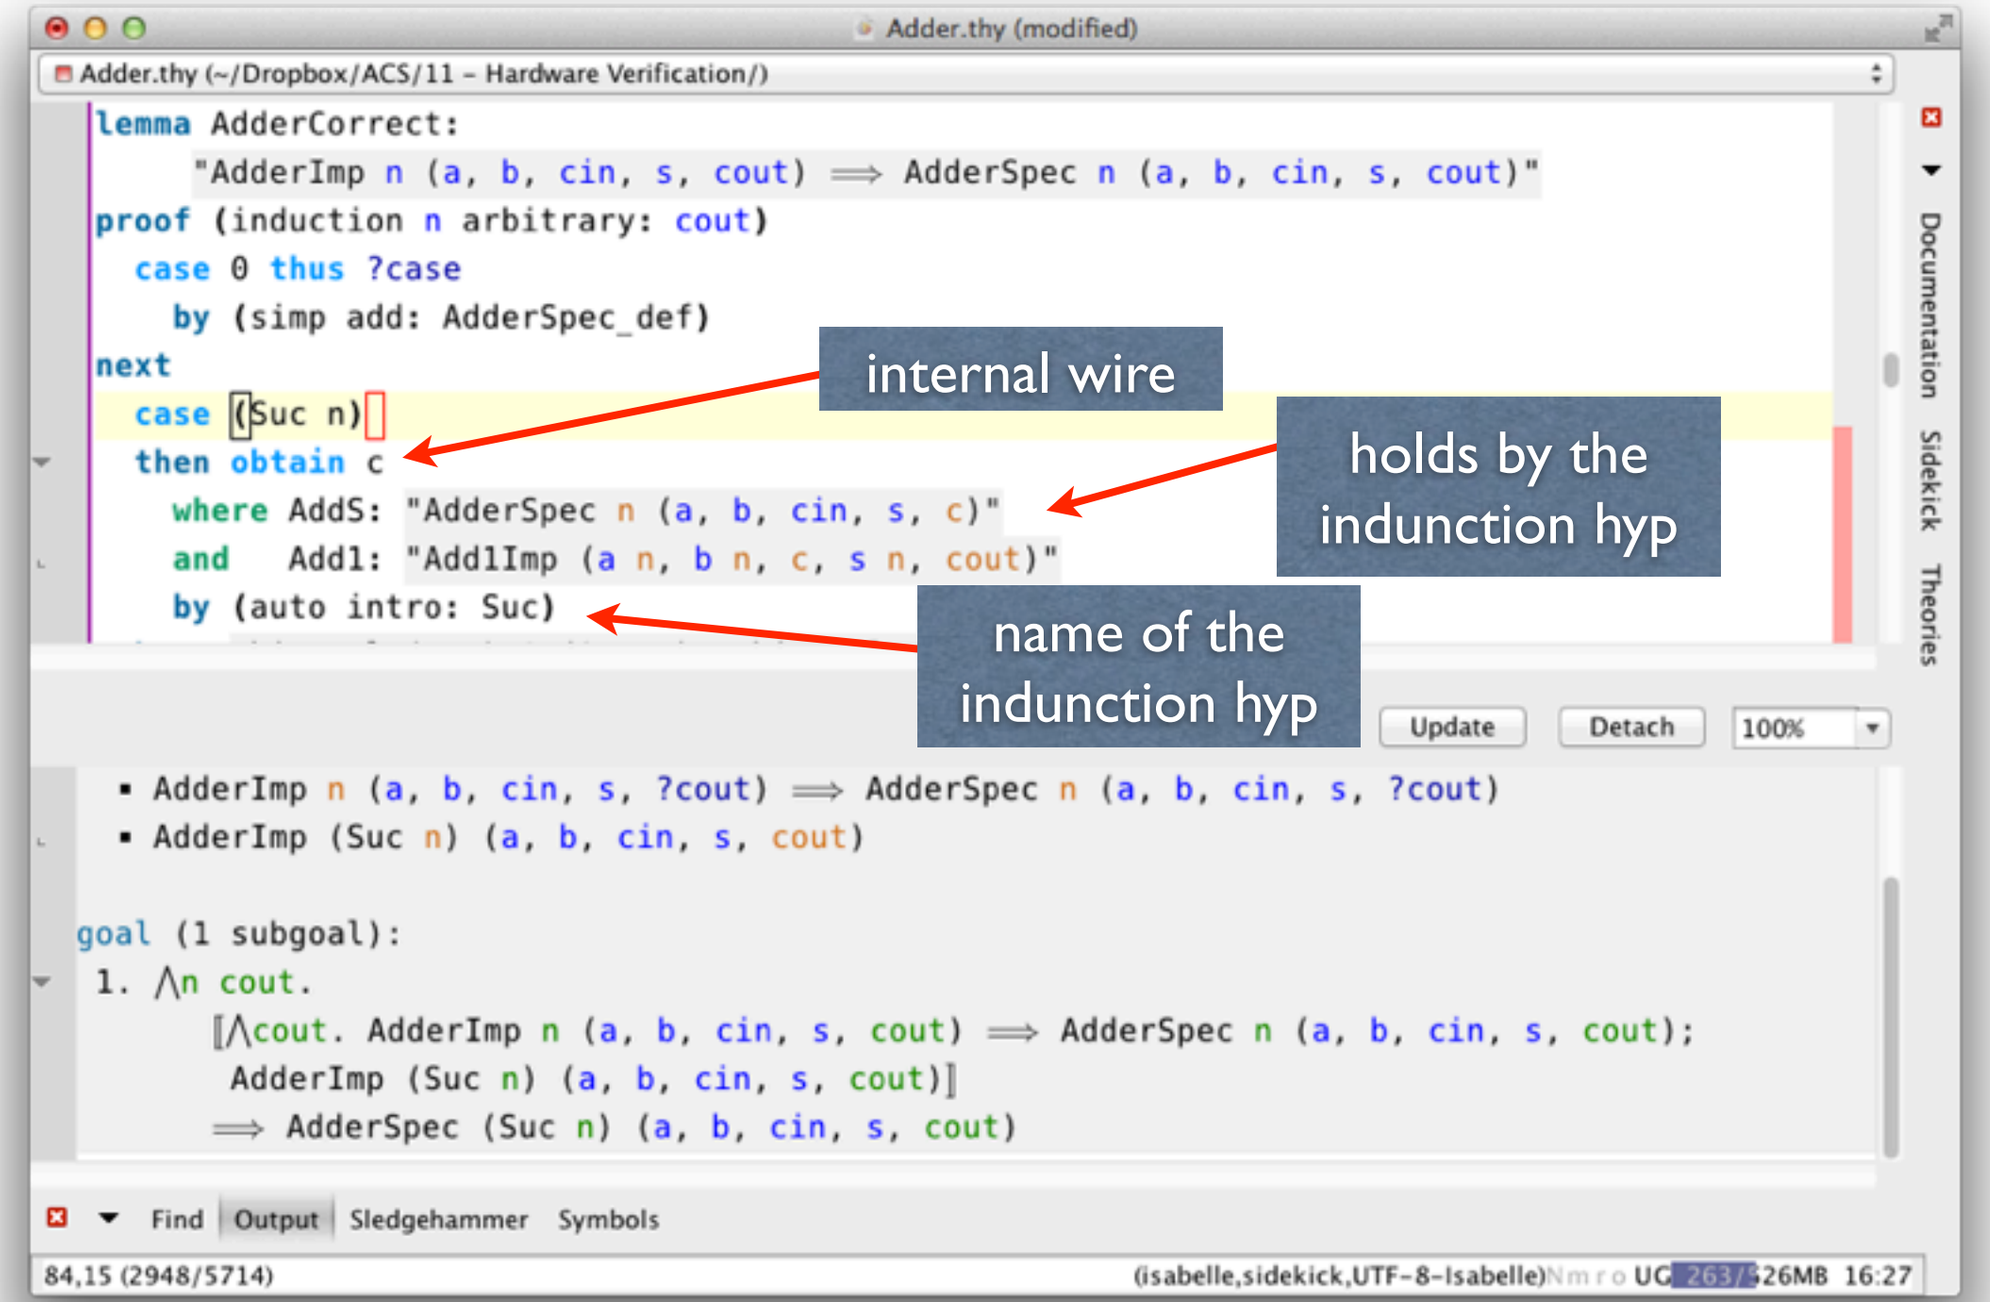

By assumption, we have `AdderImp(Suc n)` and therefore both `AdderImp n` and `Add1Imp`. The simplest use of "`obtain`" would derive those assumptions, but we can skip a step and go directly to `AdderSpec n` by referring to the induction hypothesis.

A Tiresome Calculation

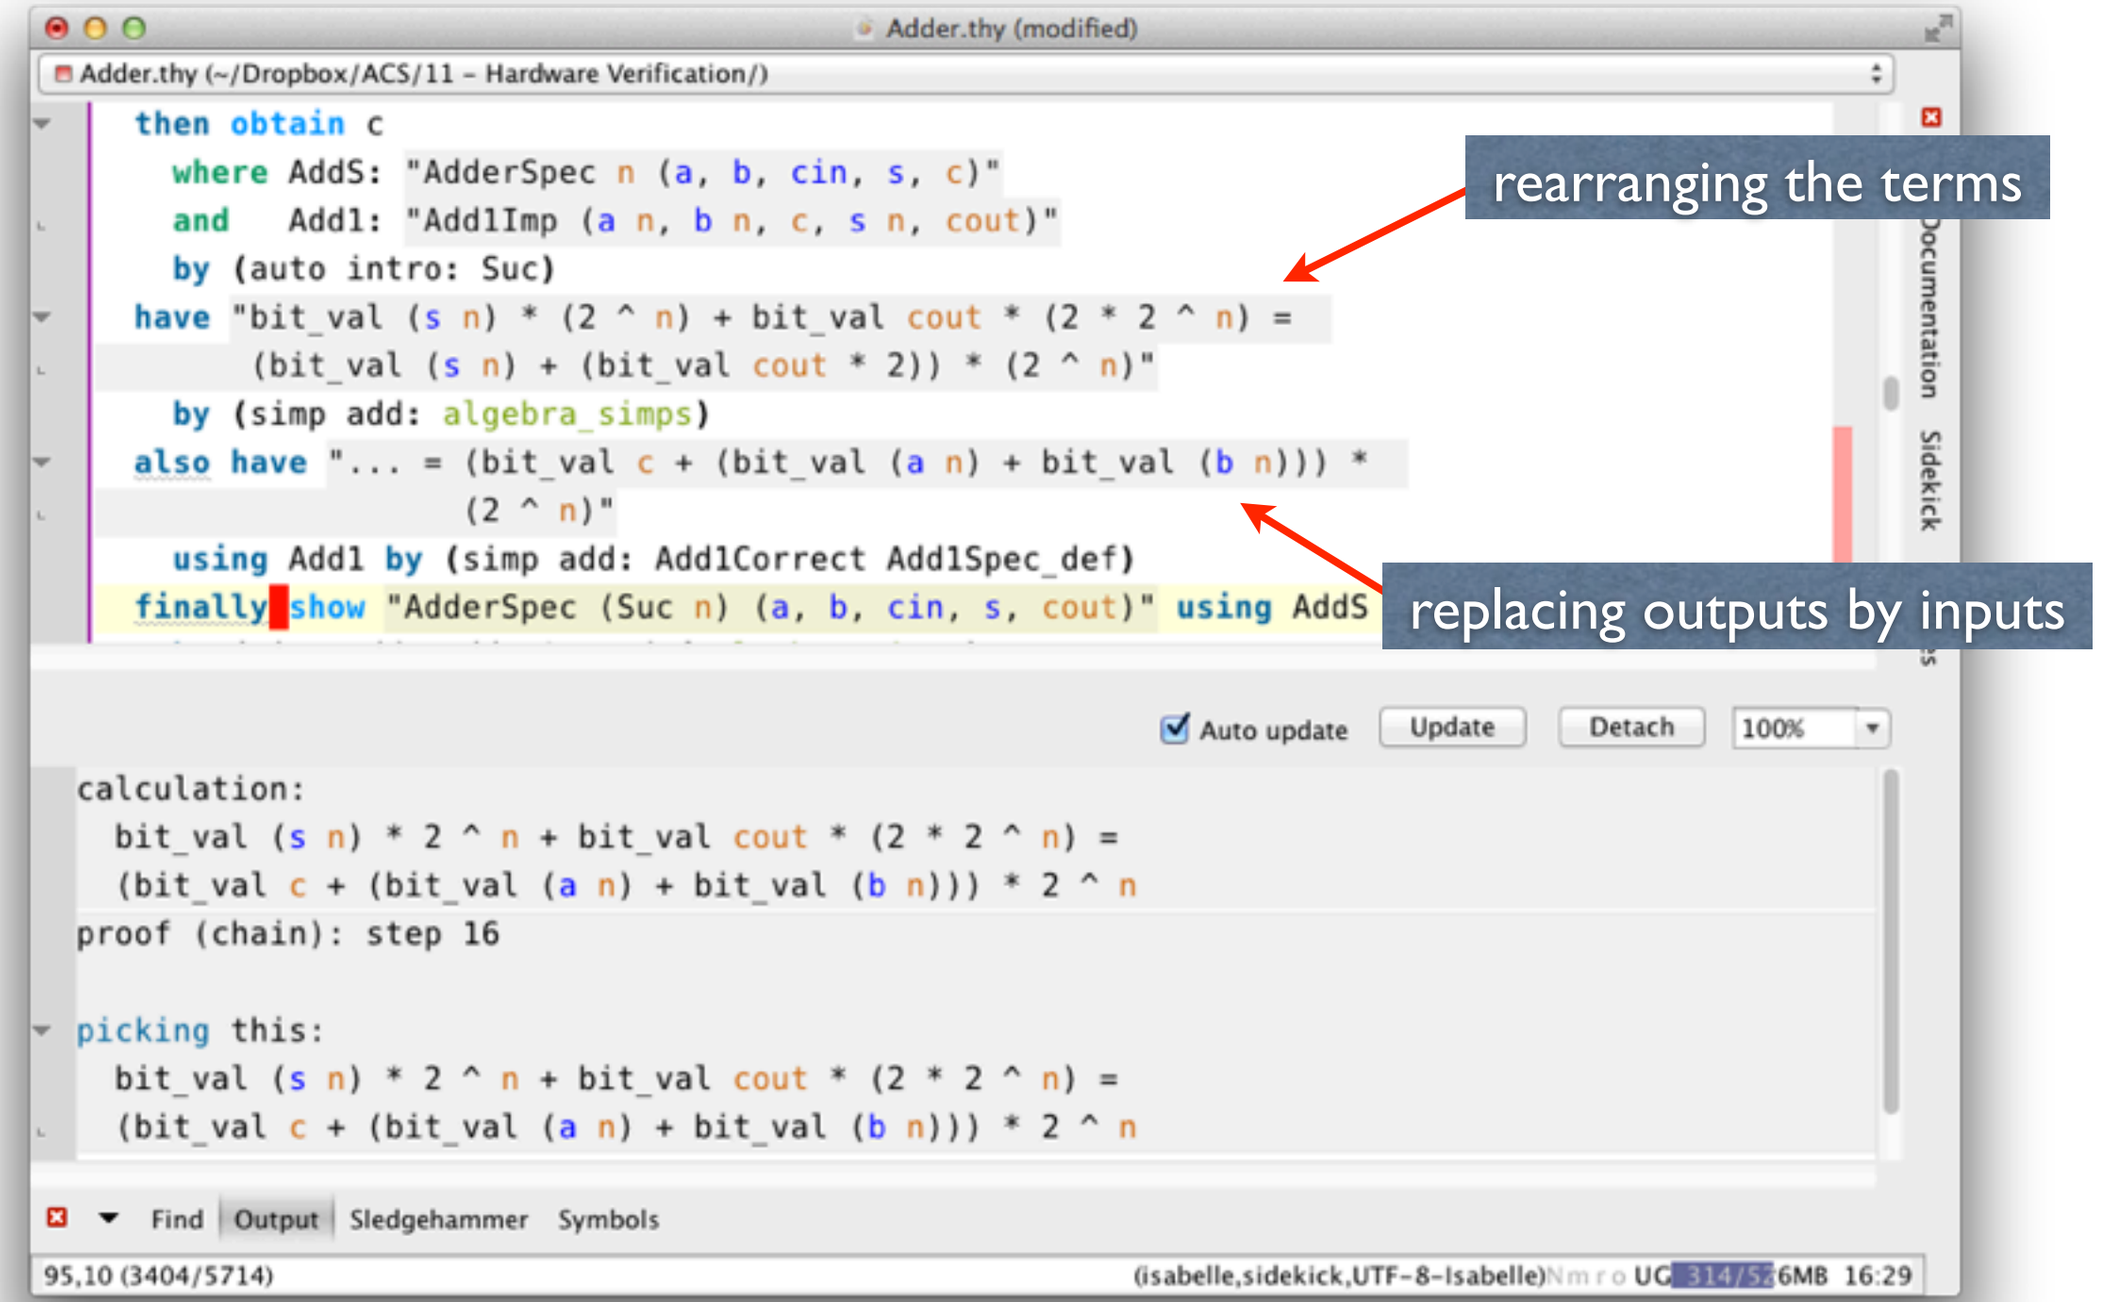

This equation is suggested by earlier attempts to prove the induction step directly. The proof involves using the correctness of a full adder to replace `Add1Imp` by `Add1Spec`, then unfolding the latter to get the sum c + a n + b n. The precise form of the left-hand side has been chosen to match a term that will appear in the main proof. This kind of reasoning is tedious even with the help of Isar. Better support for arithmetic could make this proof almost automatic.

Partial Correctness is Proved!

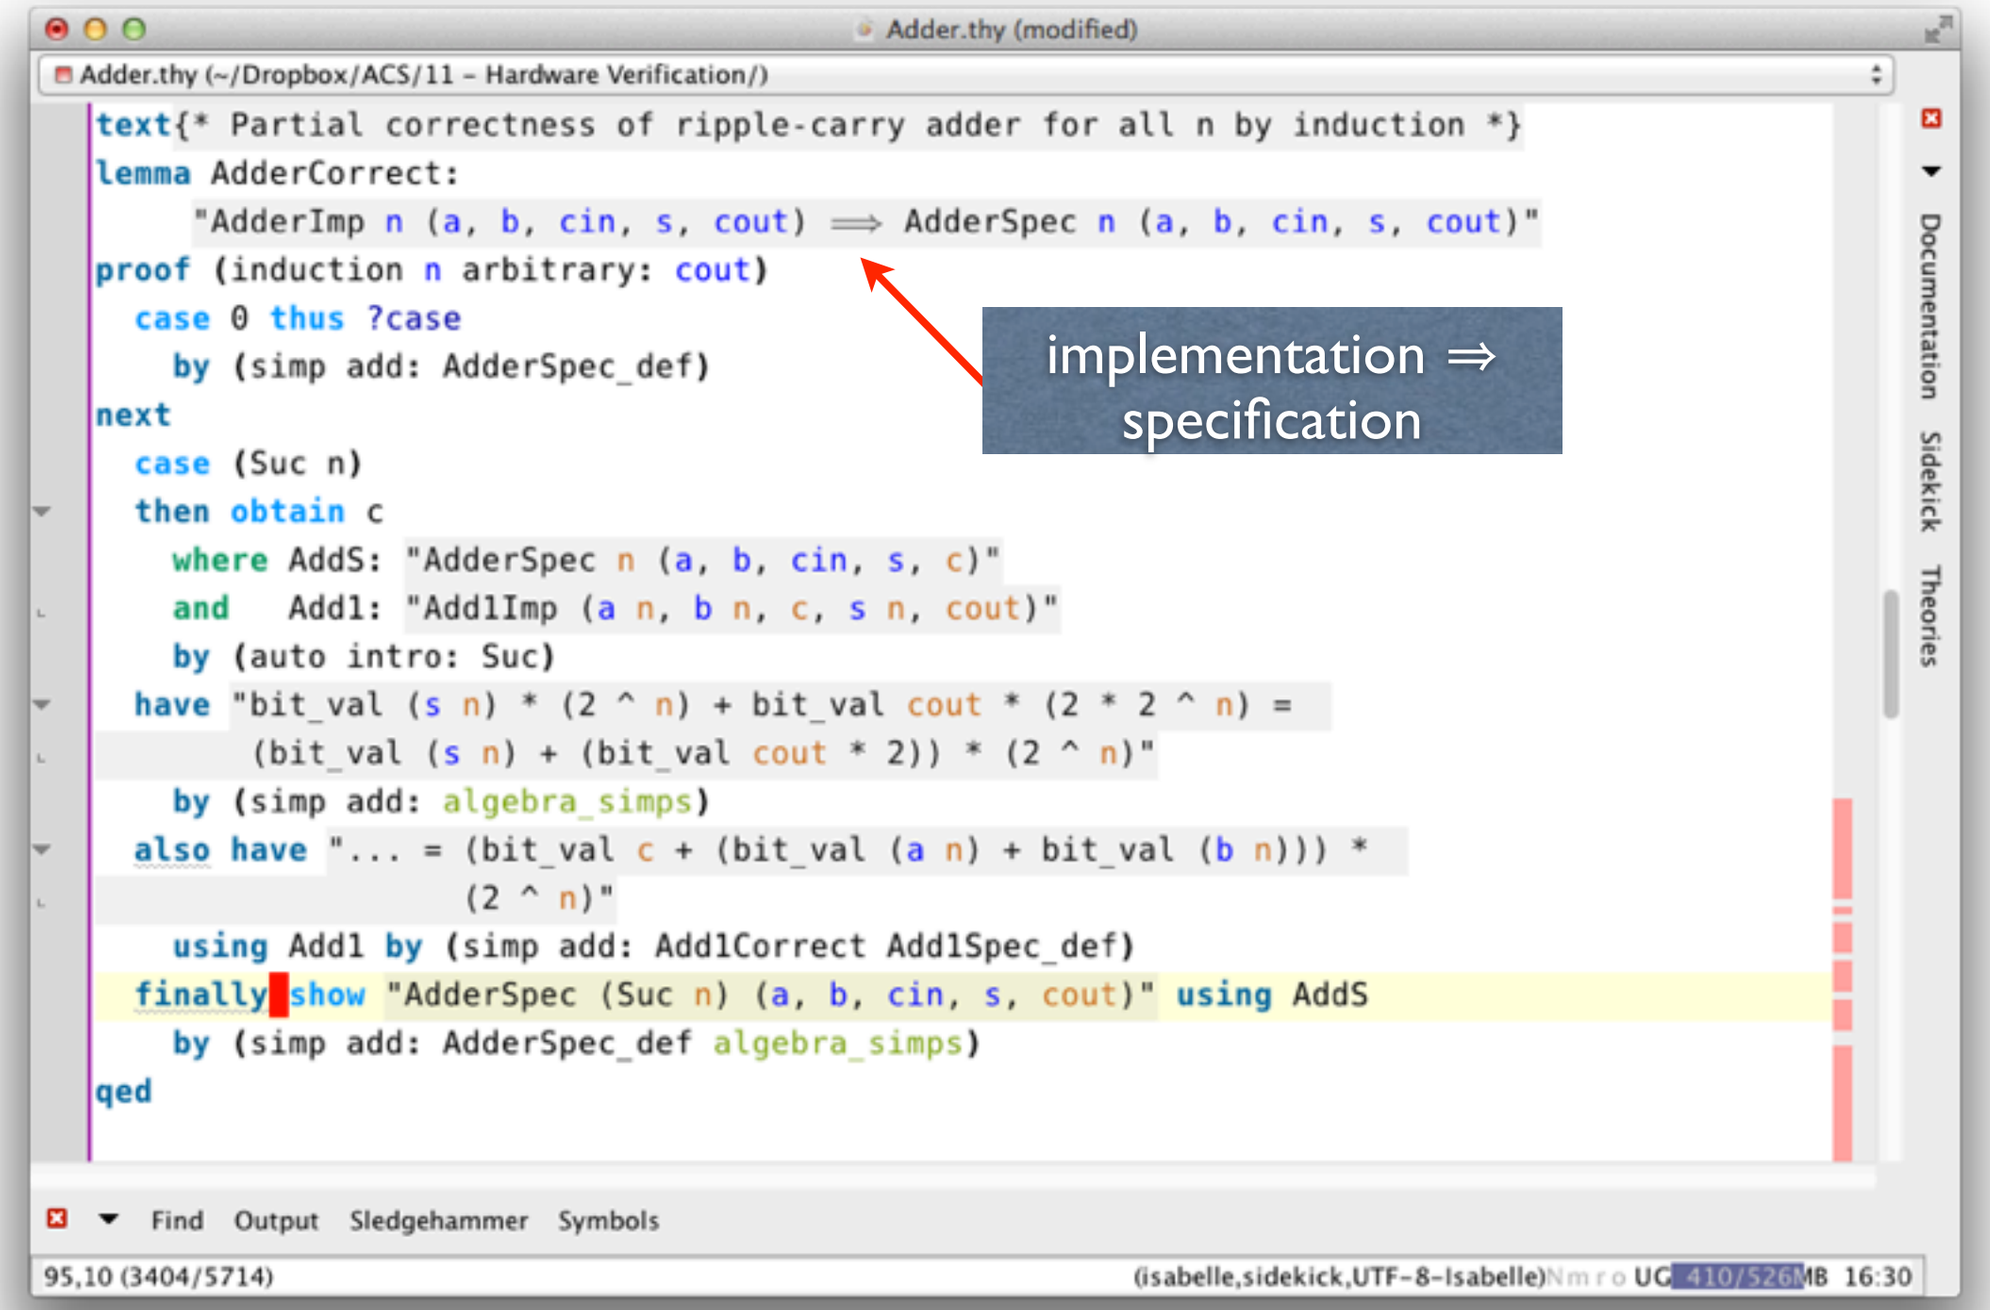

We end up with a fairly simple structure. Note that we could have used it `Add1Correct` earlier in the proof, obtaining Add1: "`Add1Spec` ..." directly.

To repeat: we have proved that every possible configuration involving the connectors to our circuit satisfies the specification of an n-bit adder. Tools based on BDDs or SAT solvers can prove instances of this result for fixed values of n, but not in the general case.

Proving Equivalence

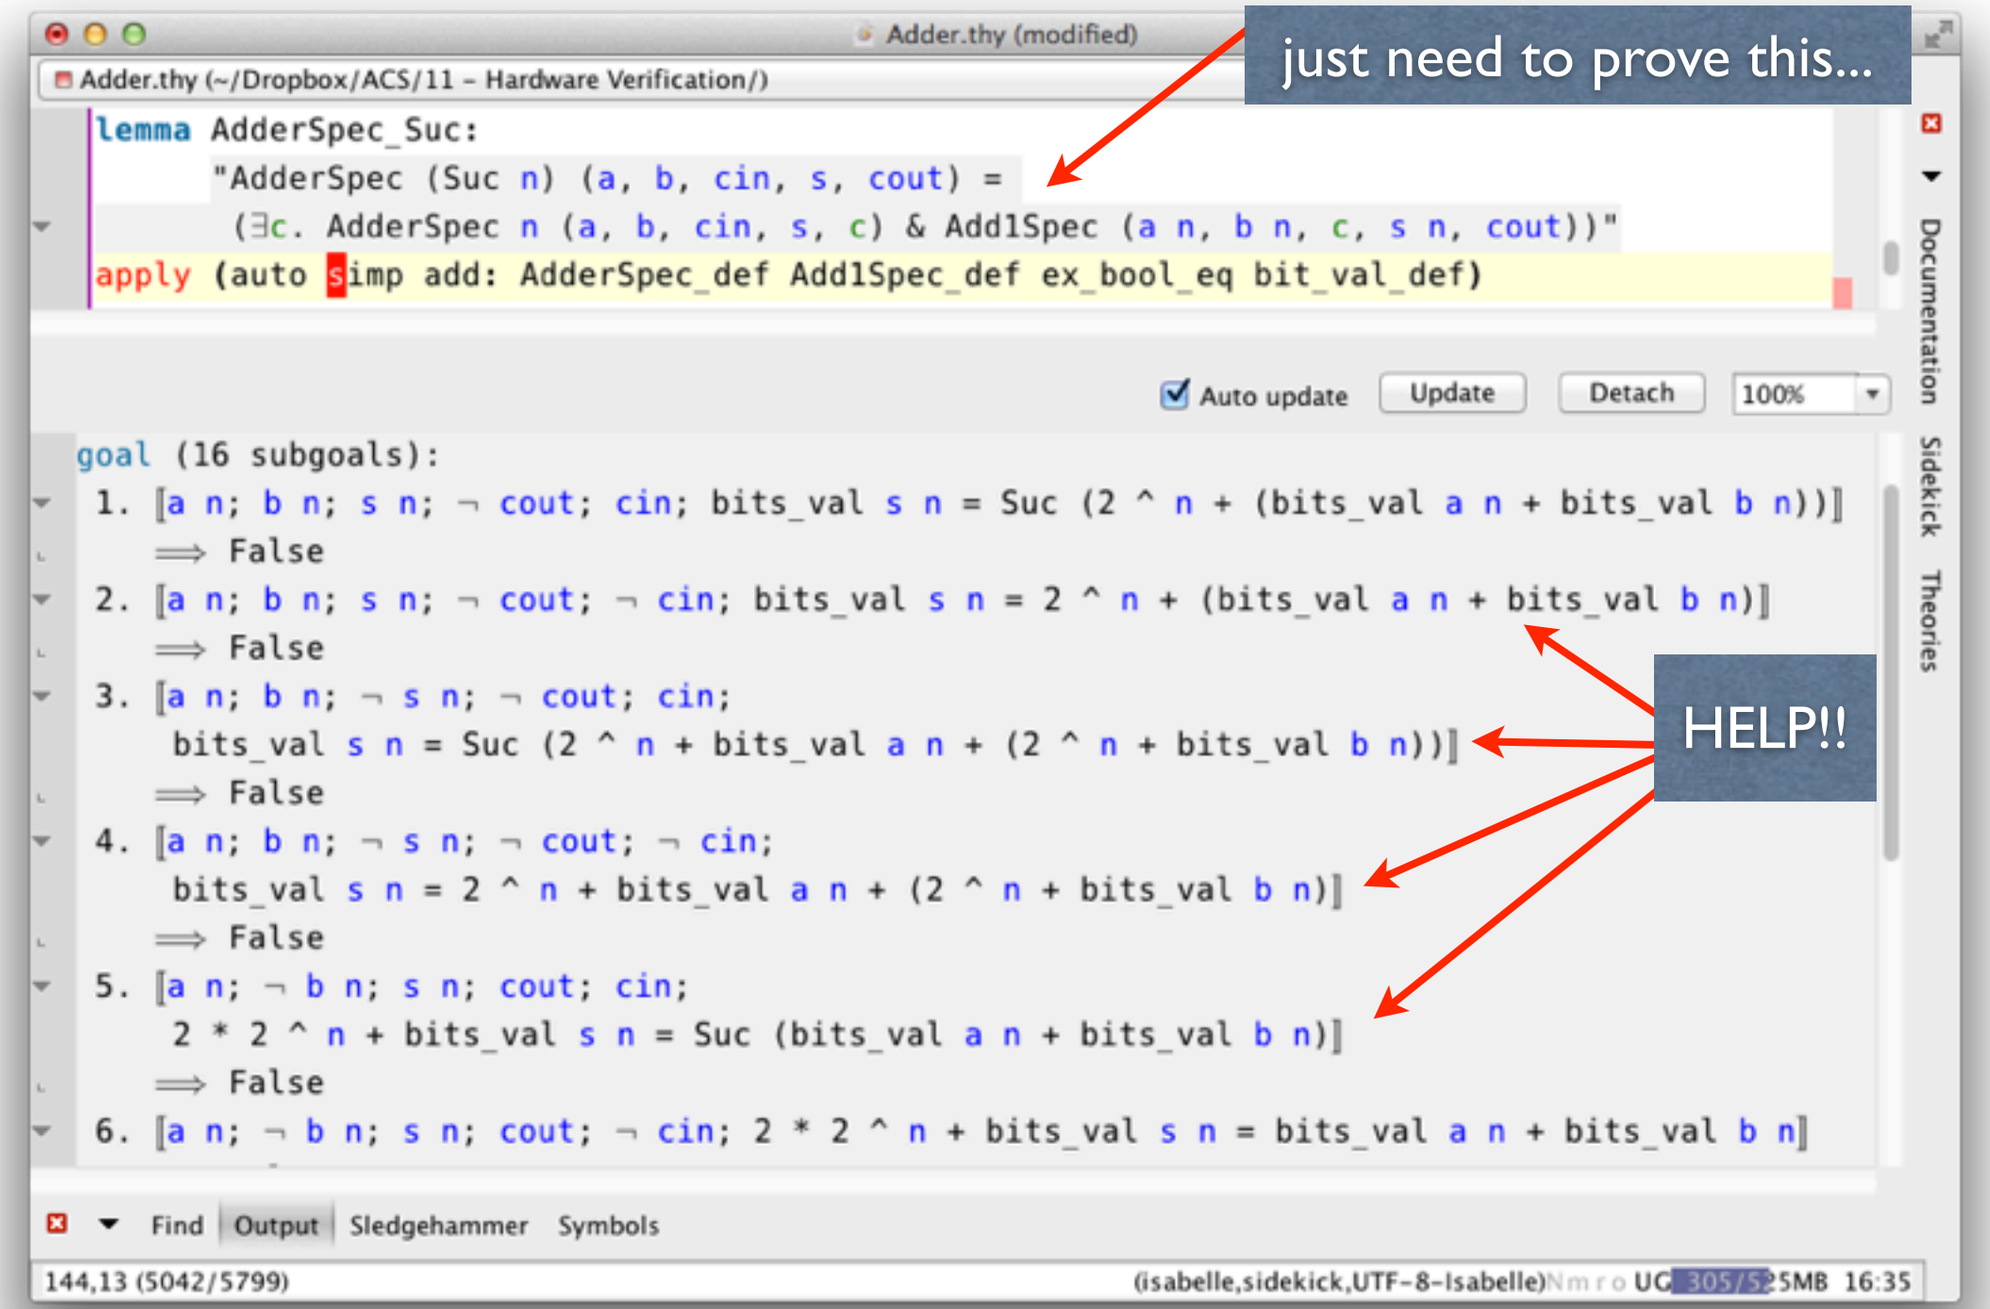

To prove that the specification implies the implementation would yield their exact equivalence. It would also guarantee the lack of short circuits in the implementation, as the specification is obviously correct.

The verification requires the lemma shown above, which resembles the recursive case of `AdderImp`. We might expect its proof to be straightforward. Unfortunately, the obvious proof attempt leaves us with 16 subgoals. A bit of thought informs us that these cases represent impossible combinations of bits. These arithmetic equations cannot hold. But how can we prove this theorem with reasonable effort?

A Crucial Lemma

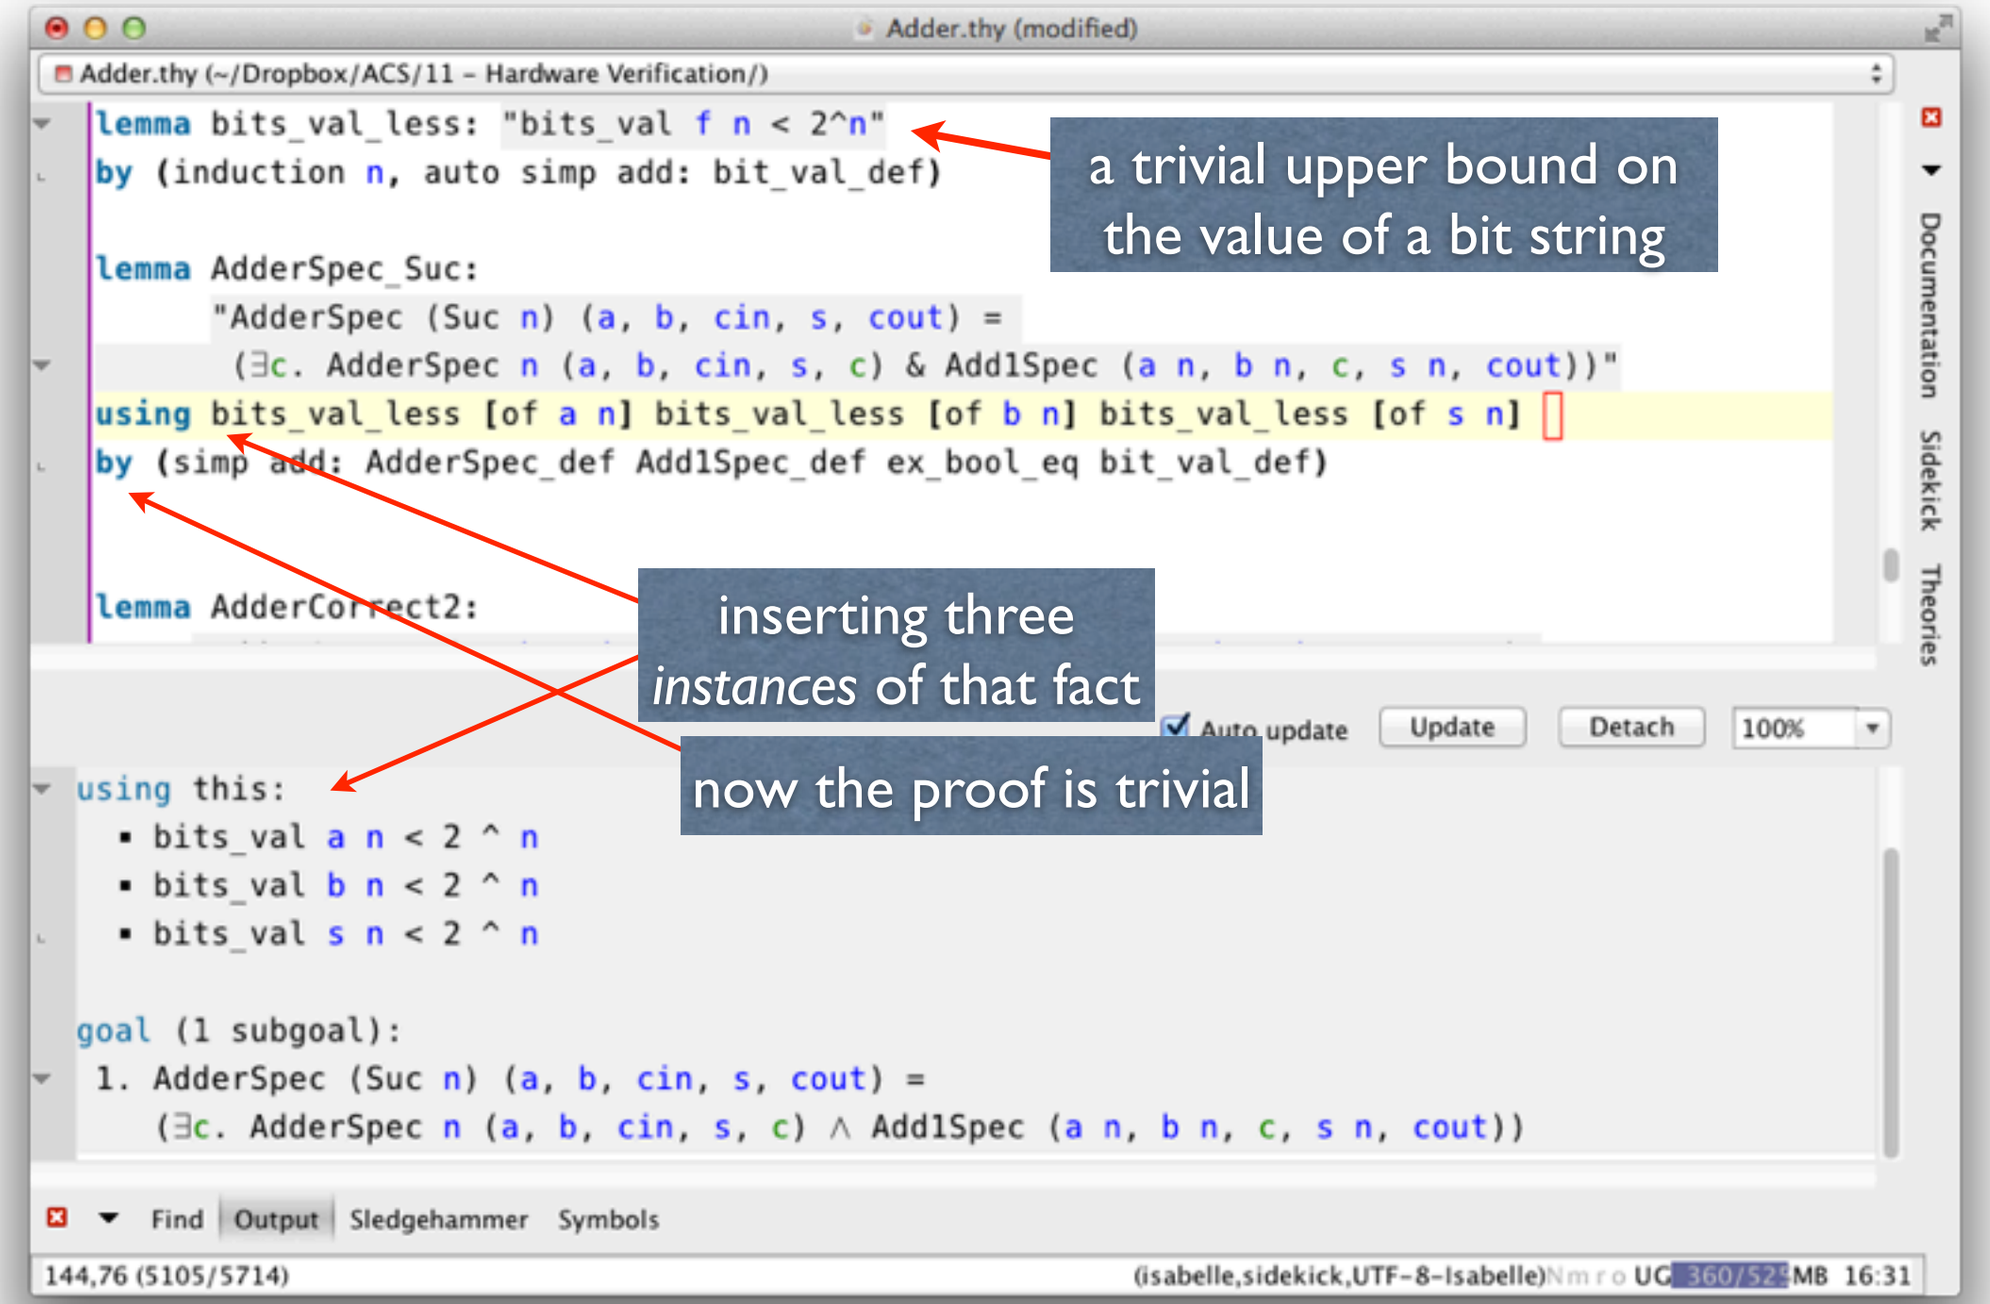

The crucial insight is that all of the impossible cases involve bit strings that have impossibly high values. It is trivial to prove the obvious upper bound on an n-bit string. Less obvious is that Isabelle's arithmetic decision procedures can dispose of the impossible cases with the help of that upper bound. We use a couple of tricks. One is that "`using`" can be inserted before the "`apply`" command, where it makes the given theorems available. The other trick is the keyword "`of`", which is described below.

The Opposite Implication

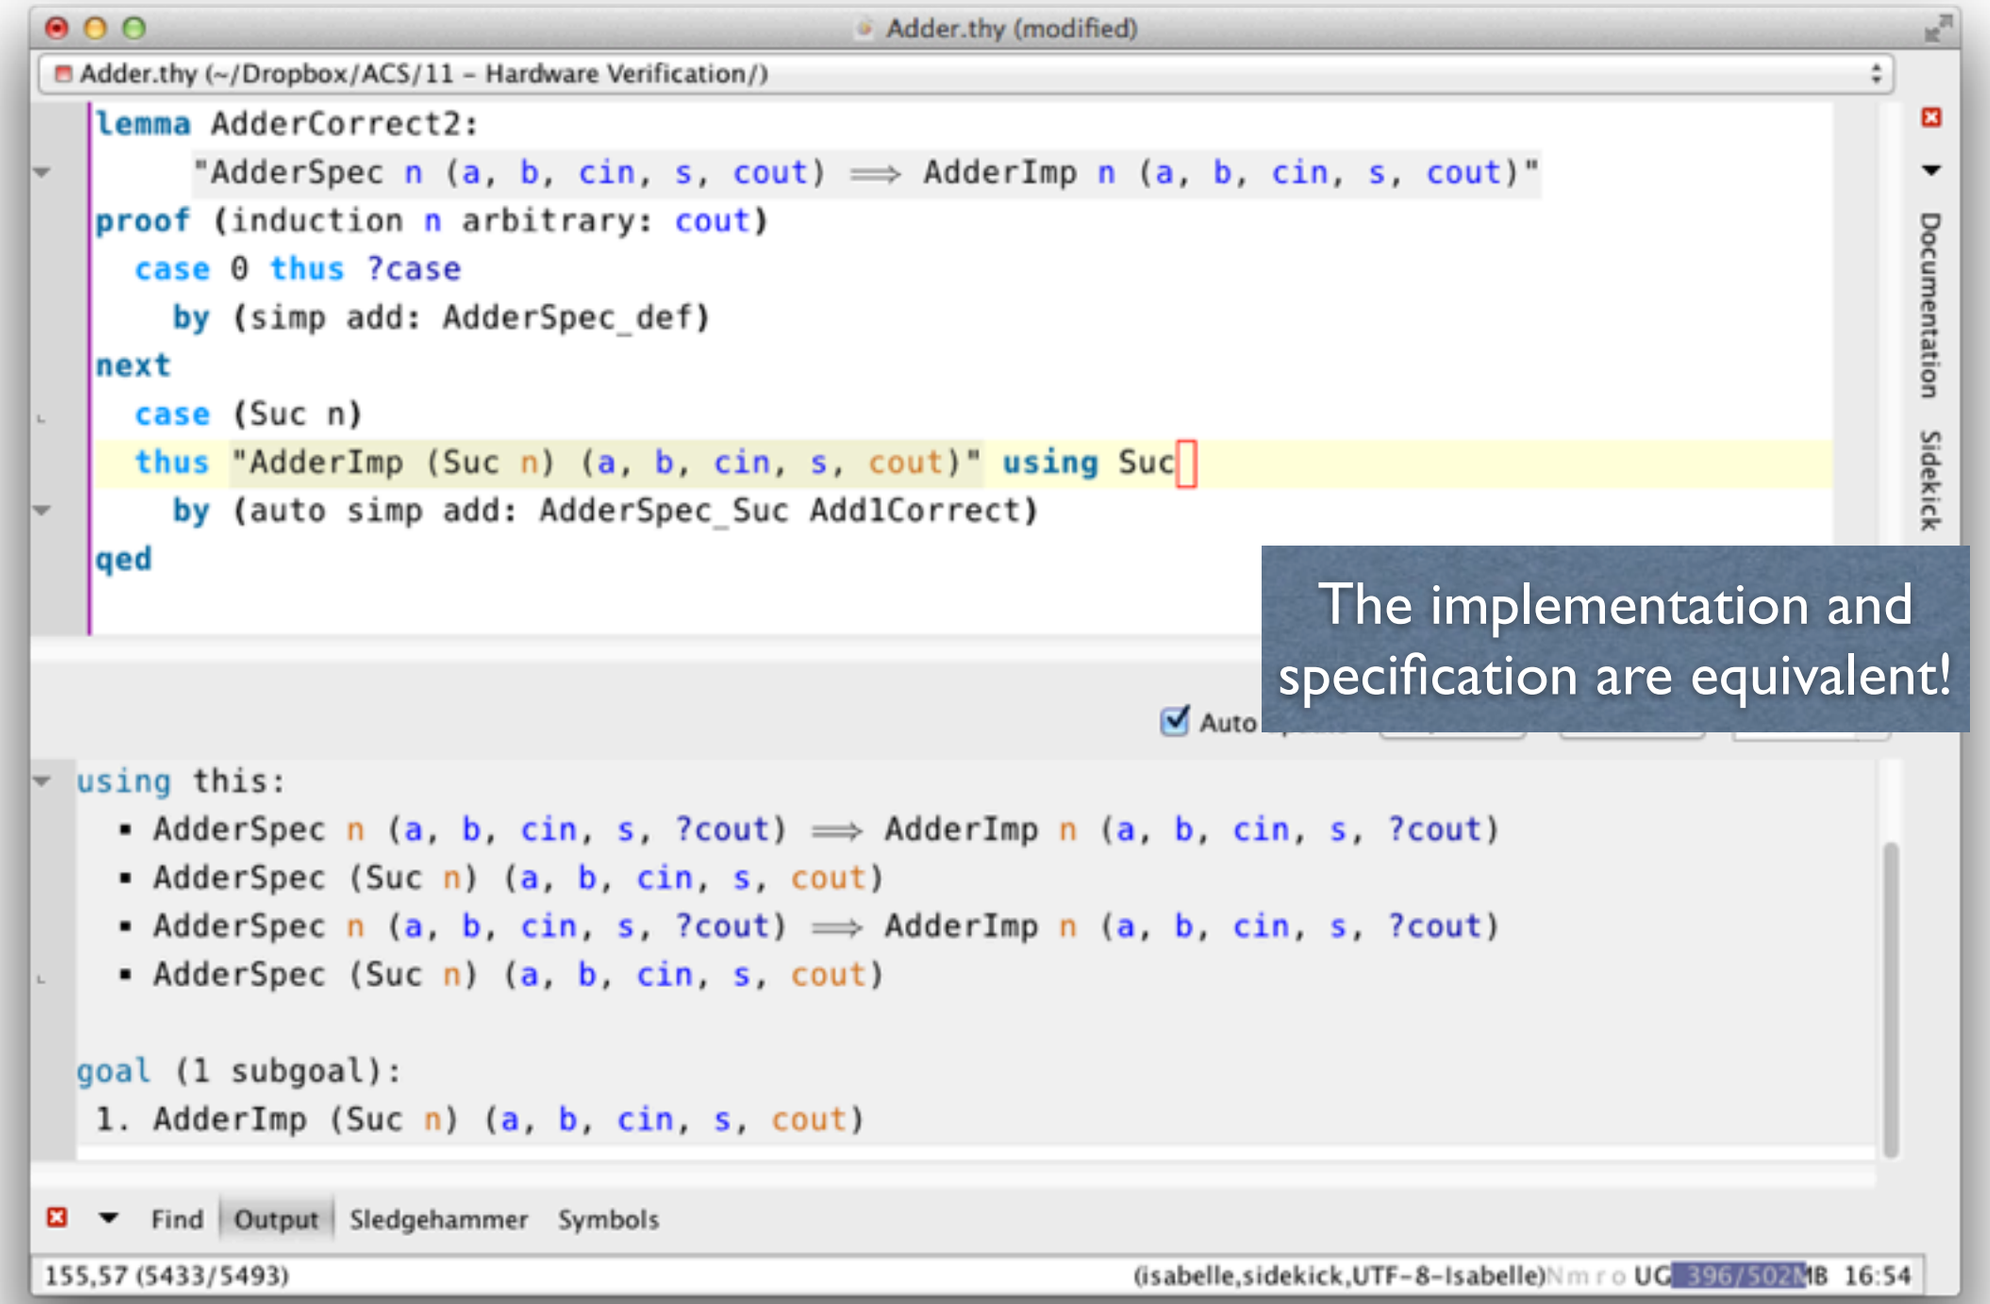

With the help of `AdderSpec_Suc`, the opposite direction of the logical equivalence is a trivial induction.

Making Instances of Theorems

- *thm* `[of` *a b c*`]`
  replaces variables by terms from left to right
- *thm* `[where` *x*=*a*`]`
  replaces the variable *x* by the term *a*
- *thm* `[OF` *thm*$_1$ *thm*$_2$ *thm*$_3$`]`
  discharges premises from left to right
- *thm* `[simplified]`
  applies the simplifier to *thm*
- *thm* [*attr*$_1$, *attr*$_2$, *attr*$_3$]
  applying multiple attributes

We proved `AdderSpec_Suc` with the help of "`using`", which inserted a crucial lemma into the proof. We needed specific instances of the lemma because Isabelle's arithmetic decision procedures cannot make use of the general formula. Fortunately, we needed only three instances and could express them using the keyword "`of`". This type of keyword is called an *attribute*. Attributes modify theorems and sometimes declare them: we have already seen attributes like `[simp]` and `[intro]` many times.

The most useful attributes are shown on the slide. Replacing variables in a theorem by terms (which must be enclosed in quotation marks unless they are atomic) can also be done using "`where`", which replaces a named variable. in the left to right list of terms or theorems, use an underscore (_) to leave the corresponding item unspecified. An example is `bits_val_less [of _ n]`, which denotes `bits_val ?f n < 2 ^ n`.

Joining theorems conclusion to premise can be done in two different ways. An alternative to `OF` is `THEN`: *thm*$_1$ [`THEN` *thm*$_2$] joins the conclusion of thm1 to the premise of thm2. Thus it is equivalent to *thm*$_2$ `[THEN` *thm*$_1$`]`. The result of such combinations can often be `simplified`. Finally, we often want to apply several attributes one after another to a theorem.

See the *Tutorial*, section **5.15 Forward Proof: Transforming Theorems**.

Interactive Formal Verification *12*: The Mutilated Chess Board

Lawrence C Paulson
Computer Laboratory
University of Cambridge

Overview

- The mutilated chessboard: a classic example in modelling a problem intuitively.
- More techniques involving Isar.
- To conclude, brief references to other Isabelle tools and capabilities.

The Mutilated Chess Board

Can this damaged board be tiled with 31 dominoes?

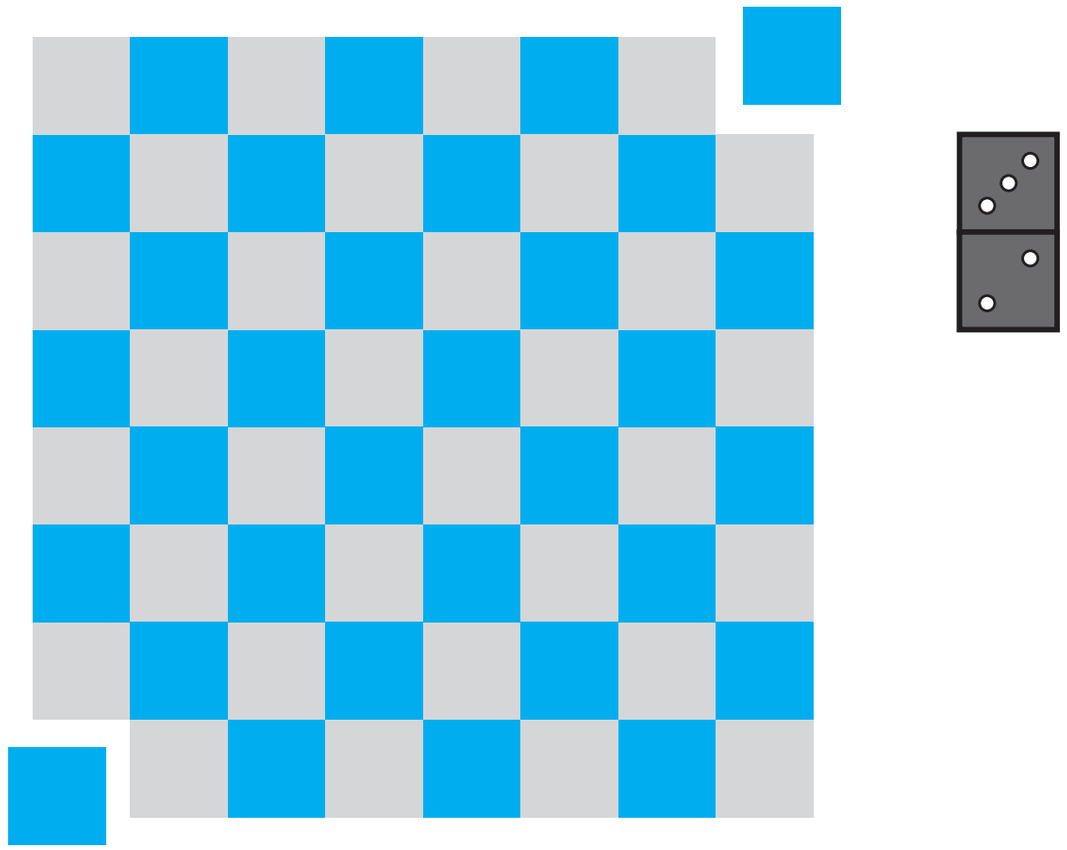

A clear proof requires an *abstract* model.

An earlier version of this formalisation is described in the paper referenced below. Comparing that version of the proof with the present one gives an indication of the progress made by Isabelle developers, especially as regards structured proof.

L. C. Paulson.
A simple formalization and proof for the mutilated chess board.
Logic J. of the IGPL **9** 3 (2001), 499–509.

`http://jigpal.oxfordjournals.org/cgi/reprint/9/3/475`

Proof Outline

- Every *row* of length $2n$ can be tiled with dominoes.
- Every *board* of size $m \times 2n$ can be tiled.
- Every tiled area has the same number of black and white squares.
- Removing some white squares from a tiled area leaves an area that cannot be tiled.
- No mutilated $2m \times 2n$ board can be tiled.

The diagram is compelling with no reasoning at all. By comparison, even the five steps shown above are more complicated than we would like. However, the Isabelle formalisation is simpler and shorter than the others that I am aware of.

An Abstract Notion of Tiling

- A *tile* is a set of *points* (such as squares).
- Given a set of tiles (such as dominoes),
  - the empty set can be tiled,
  - and so can $a \cup t$ provided
    - t can be tiled, and
    - a is a tile disjoint from t (no overlaps!)

Instead of formalising chess boards concretely, we look more abstractly at the question of covering a set by non-overlapping tiles.

Tilings Defined Inductively

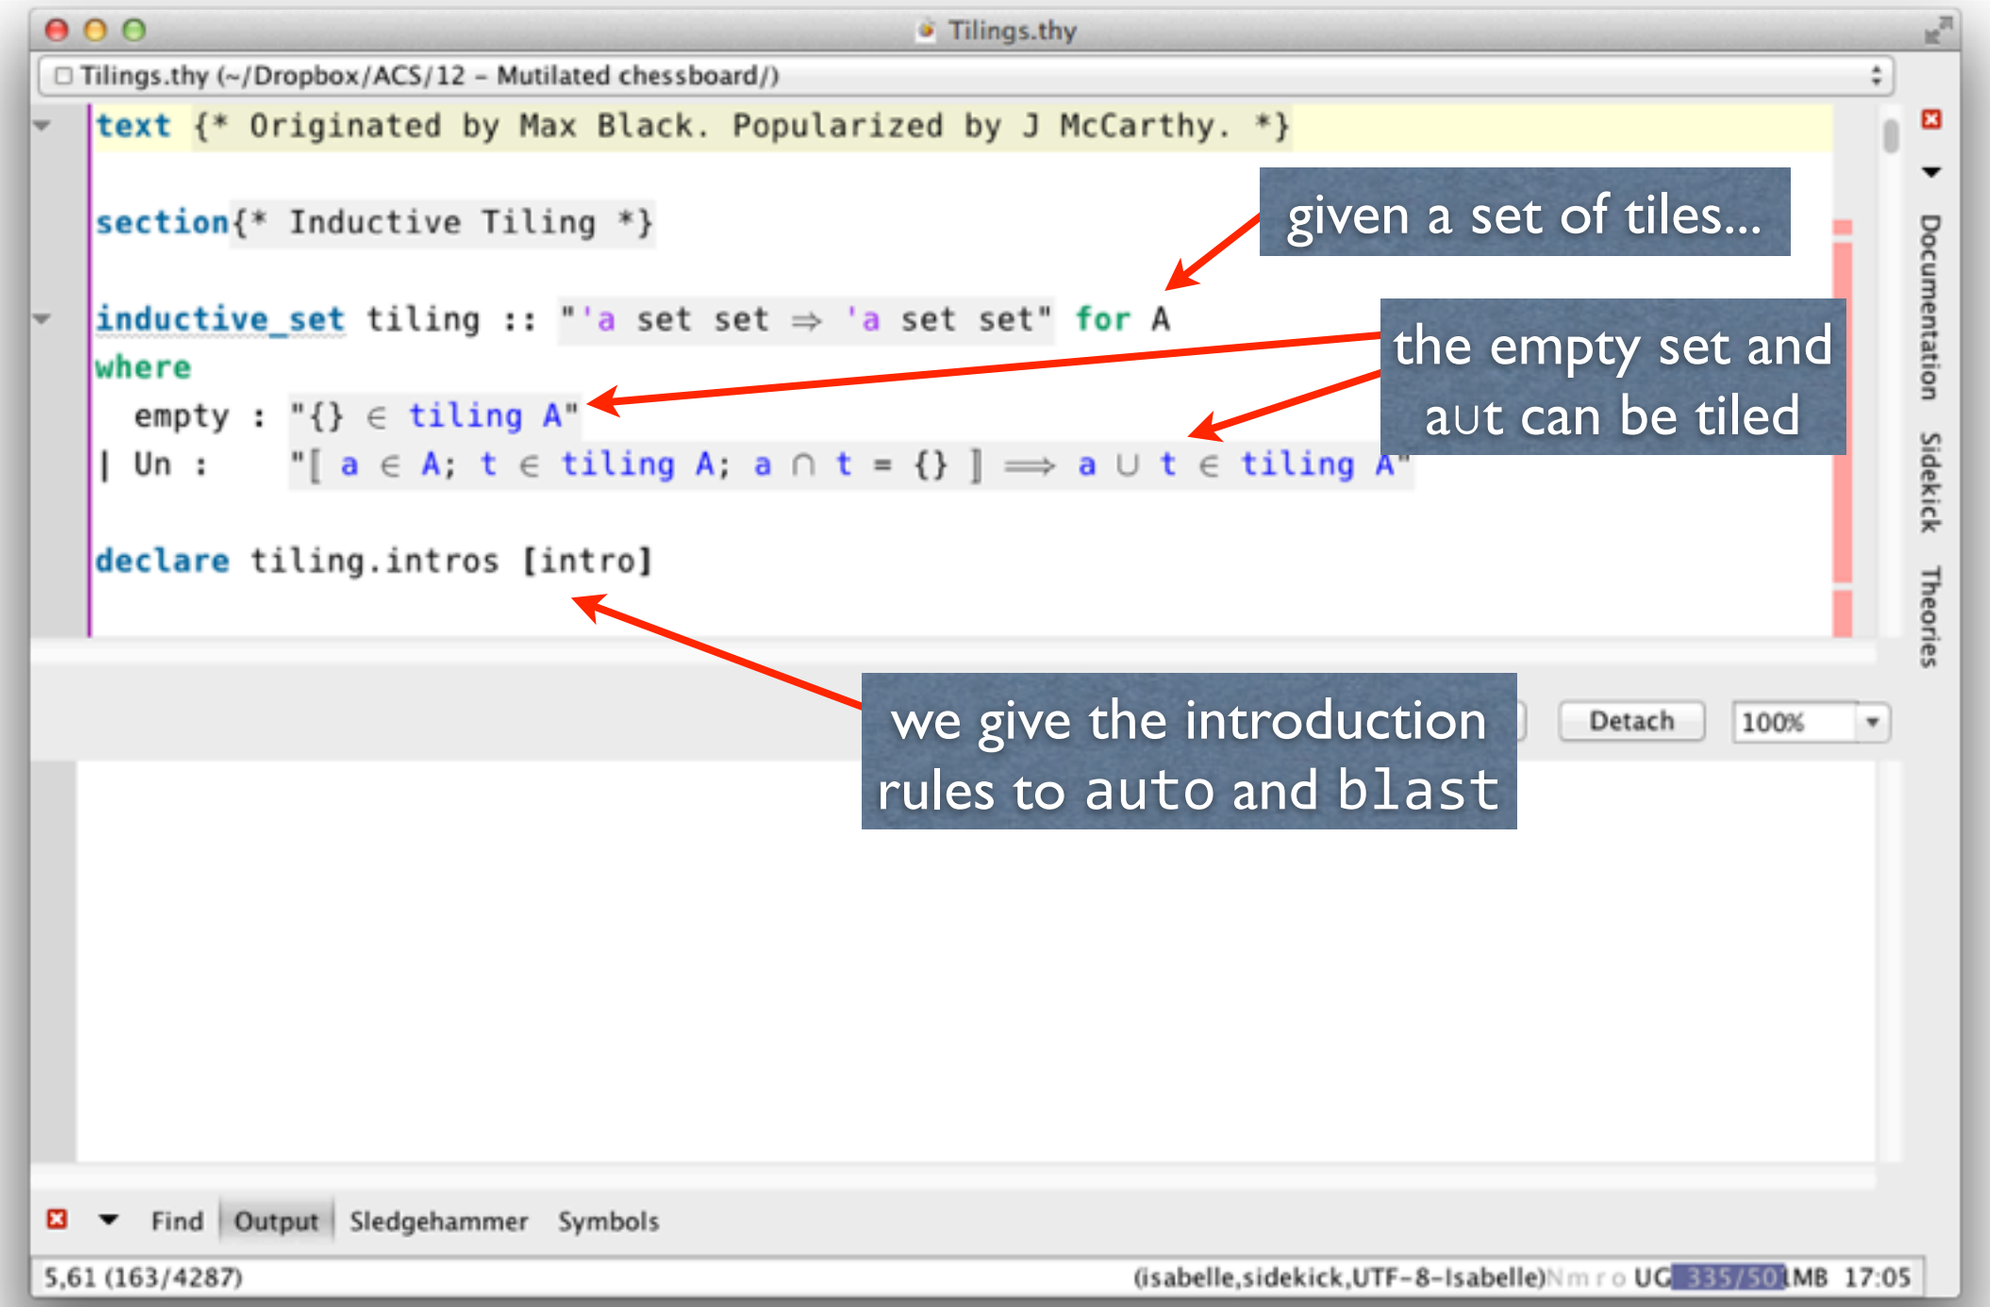

Simple Proofs about Tilings

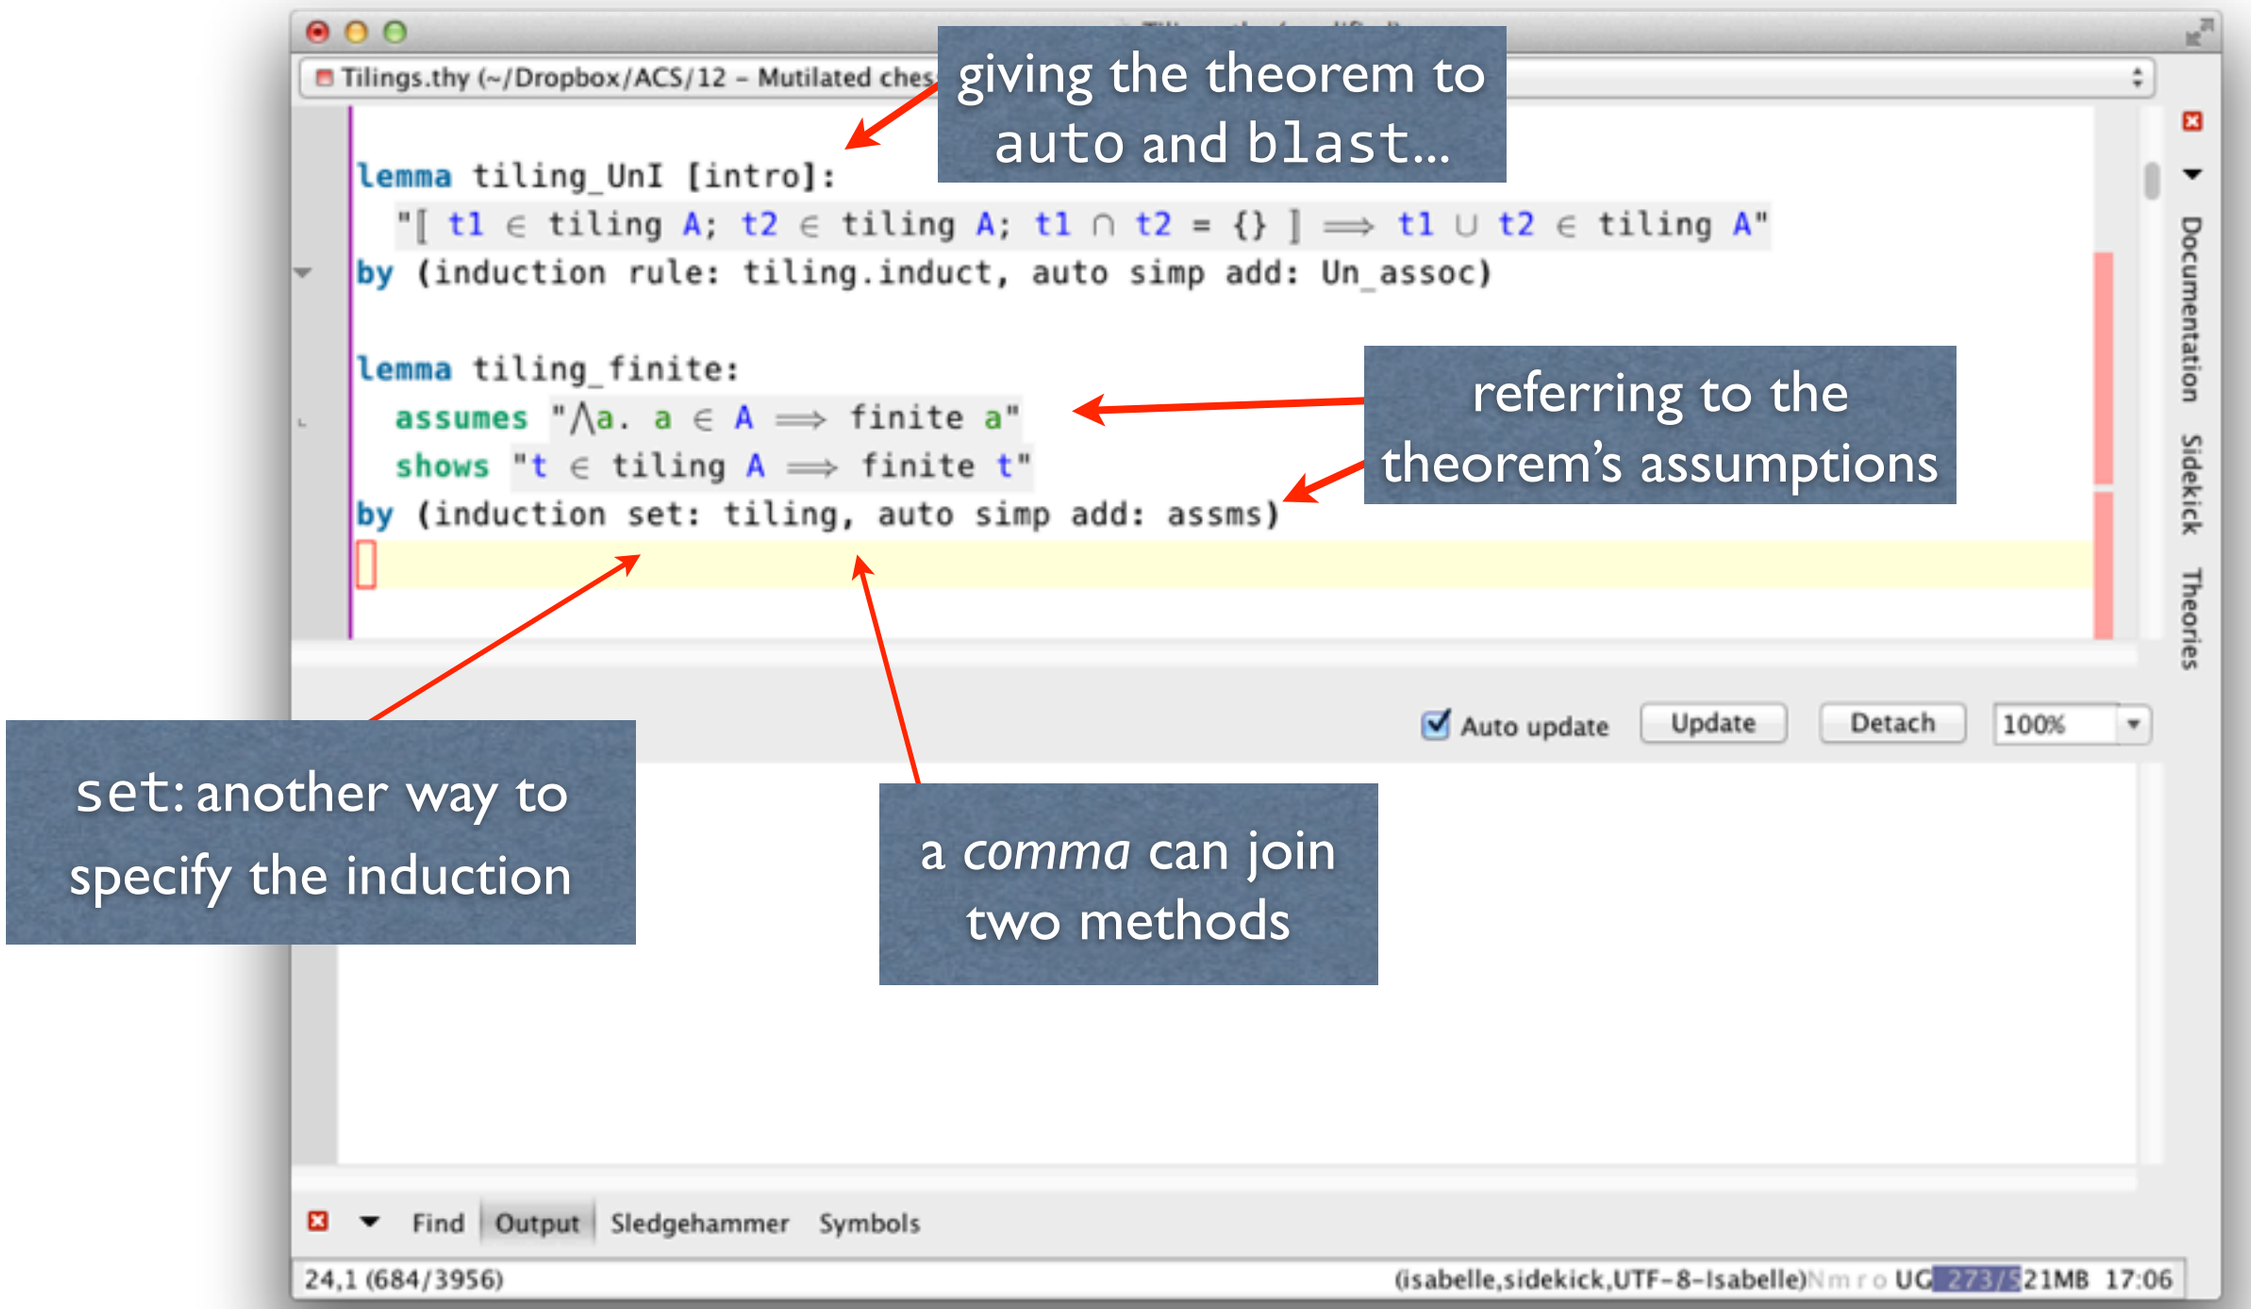

Two disjoint tilings can be combined by taking their union, yielding another tiling. The induction is trivial, using the associativity of union. Section 4 of the paper "A simple formalization and proof for the mutilated chess board" explains the proof in more detail.

If each of our tiles is a finite set, then all the tilings we can create are also finite. The induction is again trivial. Even if we have infinitely many tiles, a tiling can only use finitely many of them.

We see something new here: the identifier `assms`. It provides a uniform way of referring to the assumptions of the theorem we are trying to prove, if we have neglected to equip those assumptions with names.

Another novelty is the method `induct set: tiling`, which specifies induction over the named set without requiring us to name the actual induction rule.

Yet another novelty: we can join a series of methods using commas, creating a compound method that executes its constituent methods from left to right. Lengthy chains of methods would be difficult to maintain, but joining two or three as shown is convenient. Now the proof can be expressed using "by", because it is accomplished by a single (albeit compound) method.

Dominoes for Chess Boards

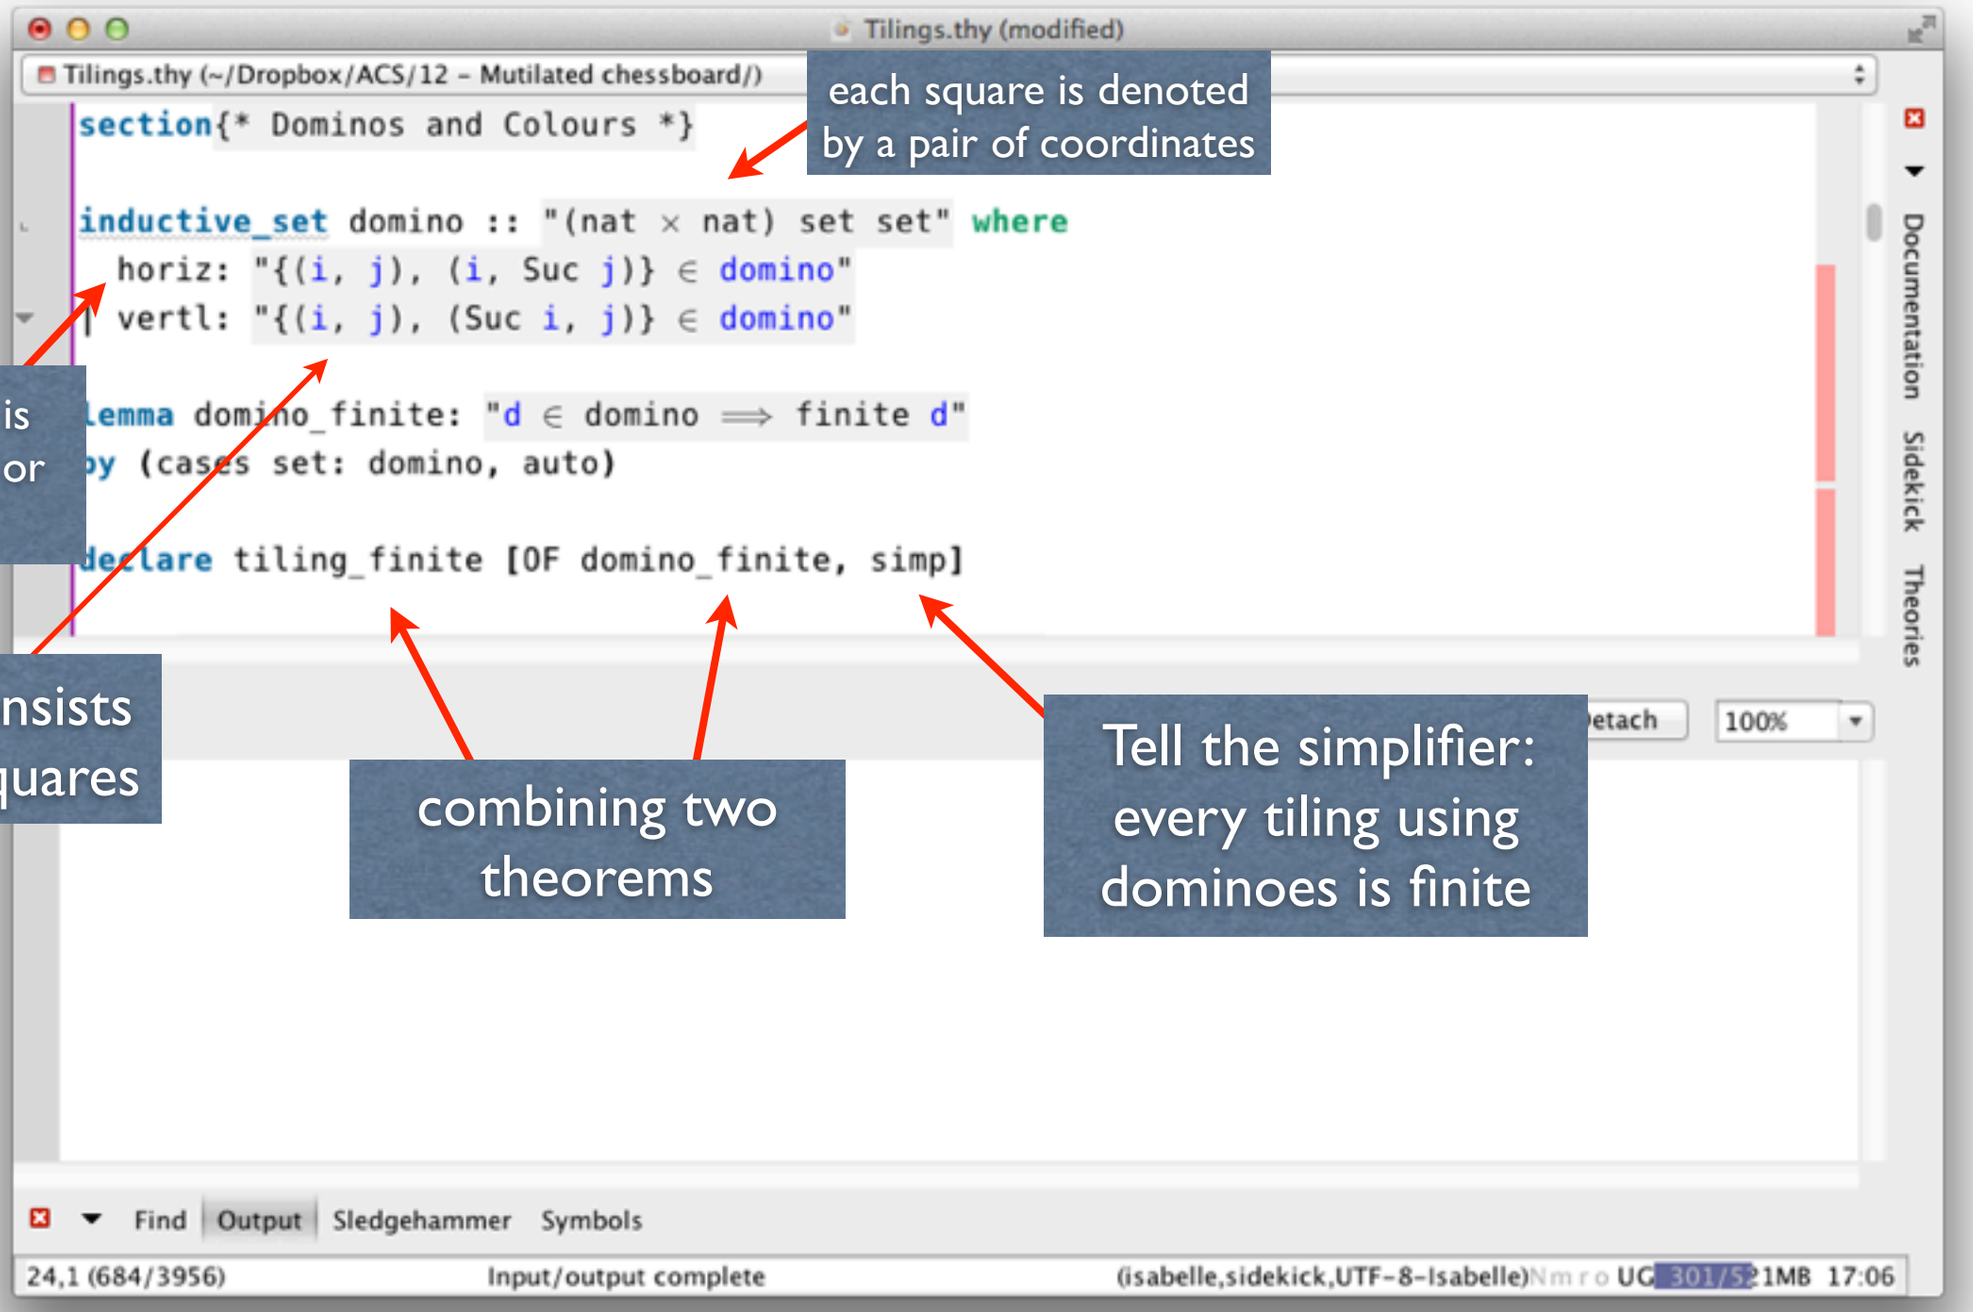

The formalisation of dominoes is extremely simple: each domino is a two element set of the form $\{(i,j), (i,j+1)\}$ or $\{(i,j), (i+1,j)\}$, expressing a horizontal or vertical orientation. The set of dominoes is not actually inductive and we could have defined it by a formula, but the inductive set mechanism is still convenient.

Because each domino contains two elements, dominoes are trivially finite. The declaration shown above combines two finiteness properties, asserting that tilings that consist of dominoes are finite, and it gives this fact to the simplifier. Concluding a series of attributes by `simp` or `intro` is common.

White and Black Squares

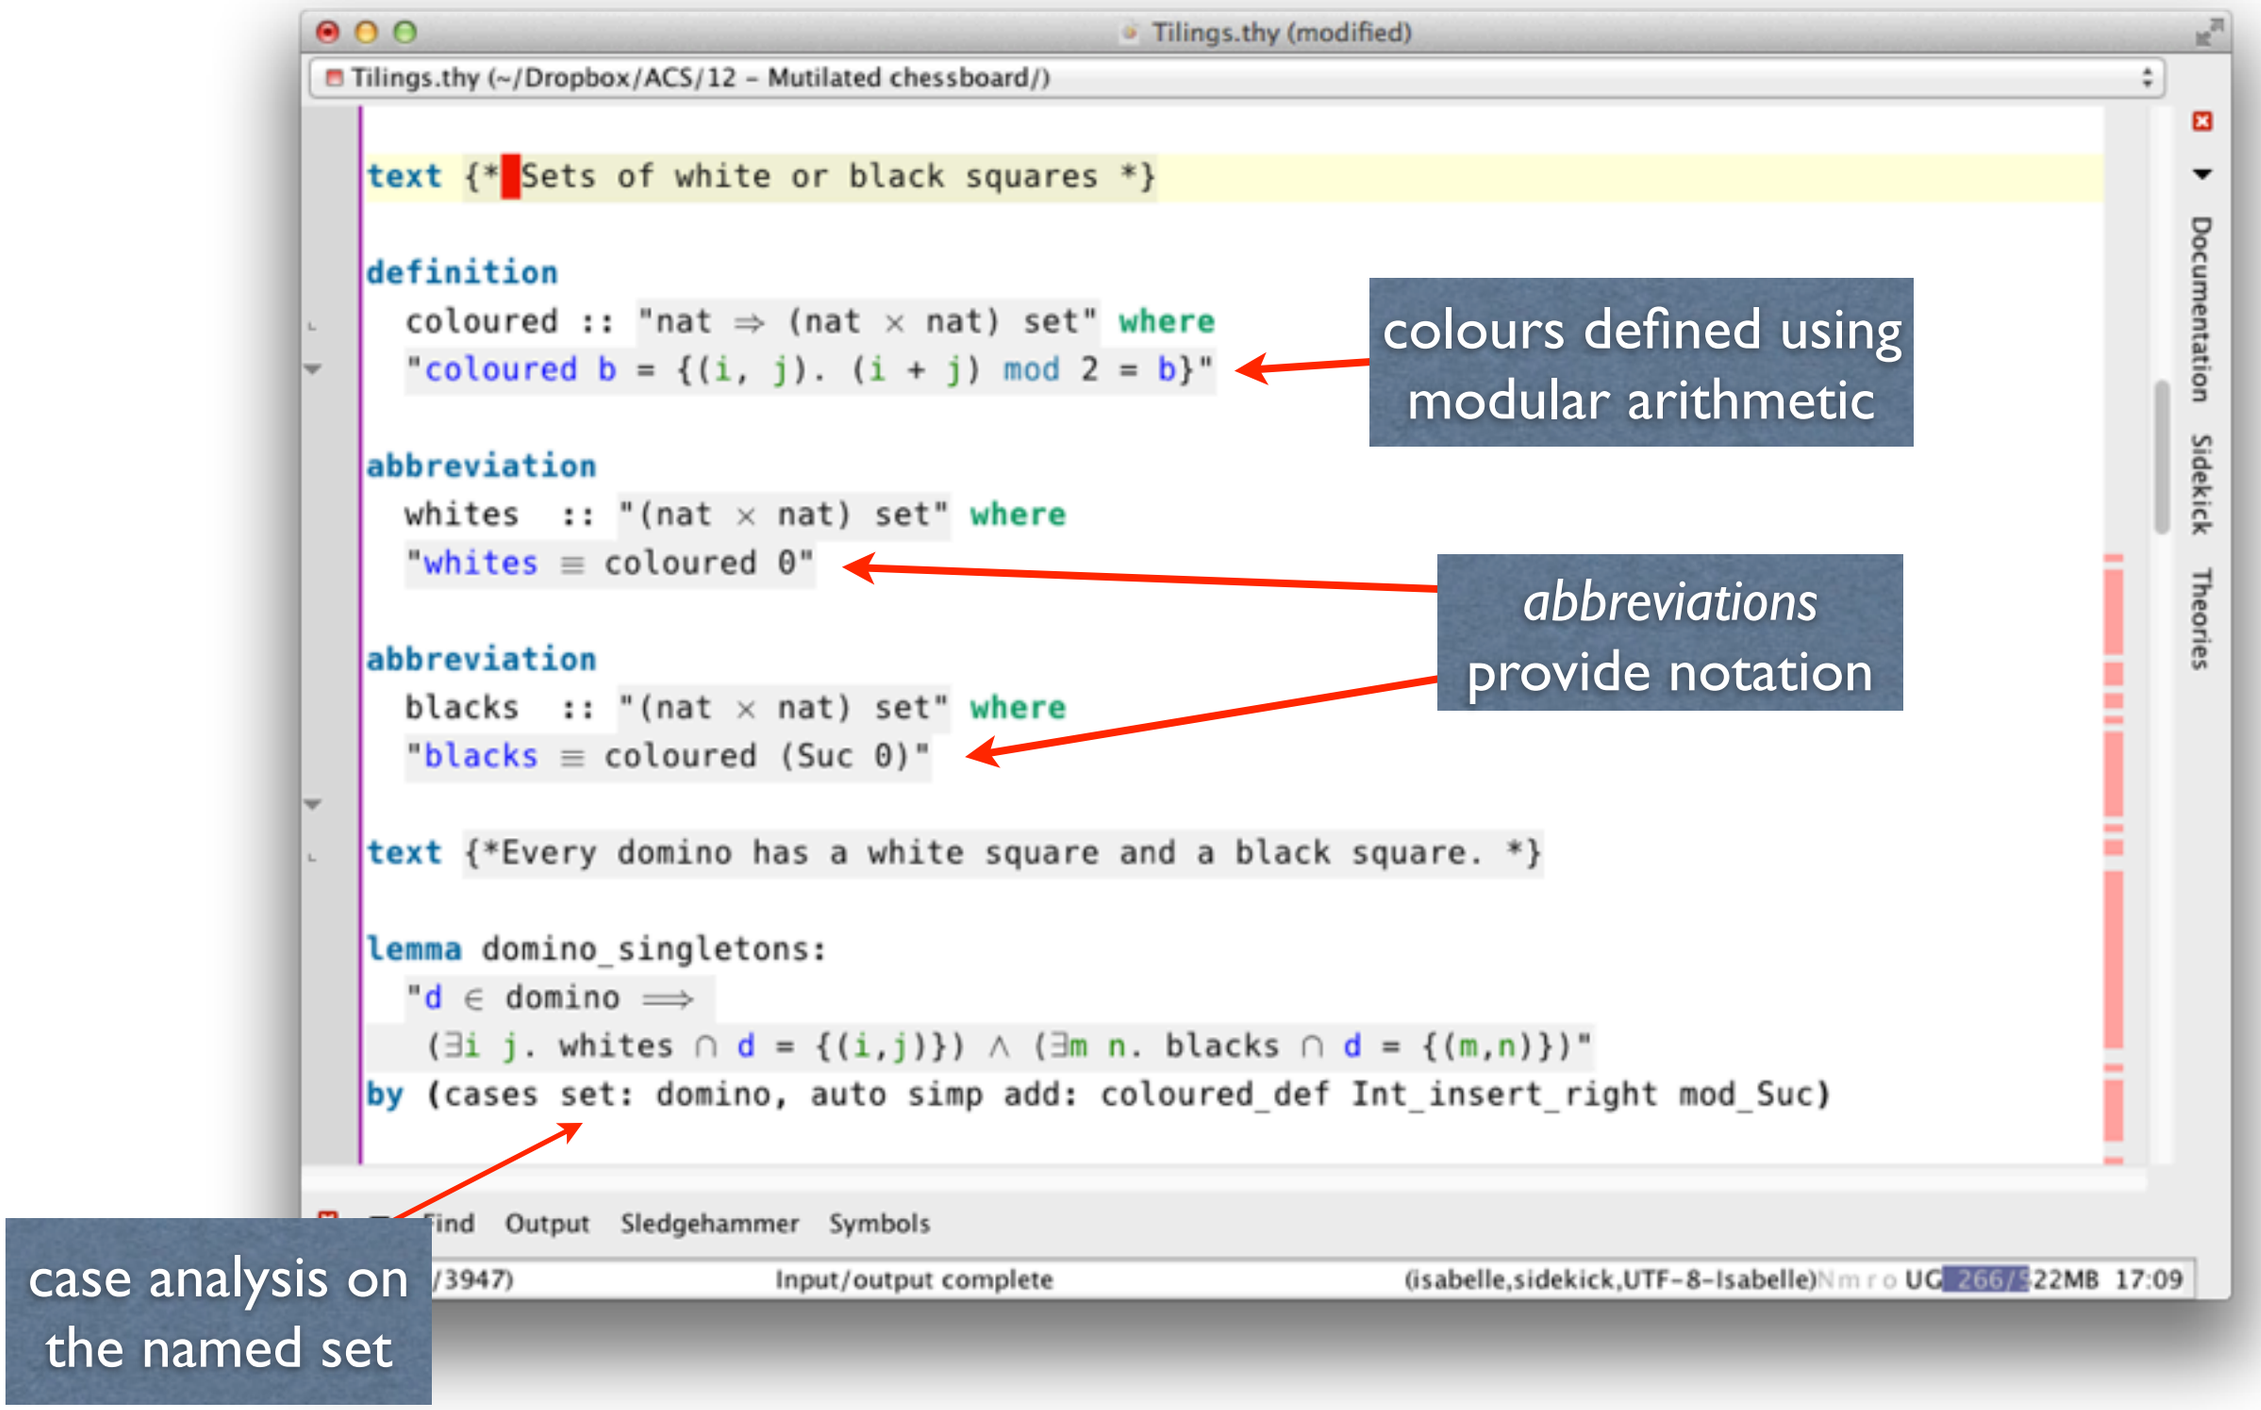

The distinction between white and black is made using modulo-2 arithmetic. The constants “whites” and “blacks” do not have definitions in the normal sense; they are declared as abbreviations, which means that these constants never occur in terms. They provide a shorthand for expressing the terms “coloured 0” and “coloured (Suc 0)”. Recall that to define a constant in Isabelle introduces an equation that can be used to replace the constant by the defining term. And this equation is not even available to the simplifier by default. With abbreviations, no such equations exist.

See the *Tutorial*, section **4.1.4 Abbreviations**, for more information. More generally, section 4.1 describes concrete syntax and infix annotations for Isabelle constants.

It is now trivial to prove that every domino has a white square and a black square, by case analysis on the two kinds of domino. The proof requires giving the simplifier some facts about intersection and the modulus function.

Rows and Columns

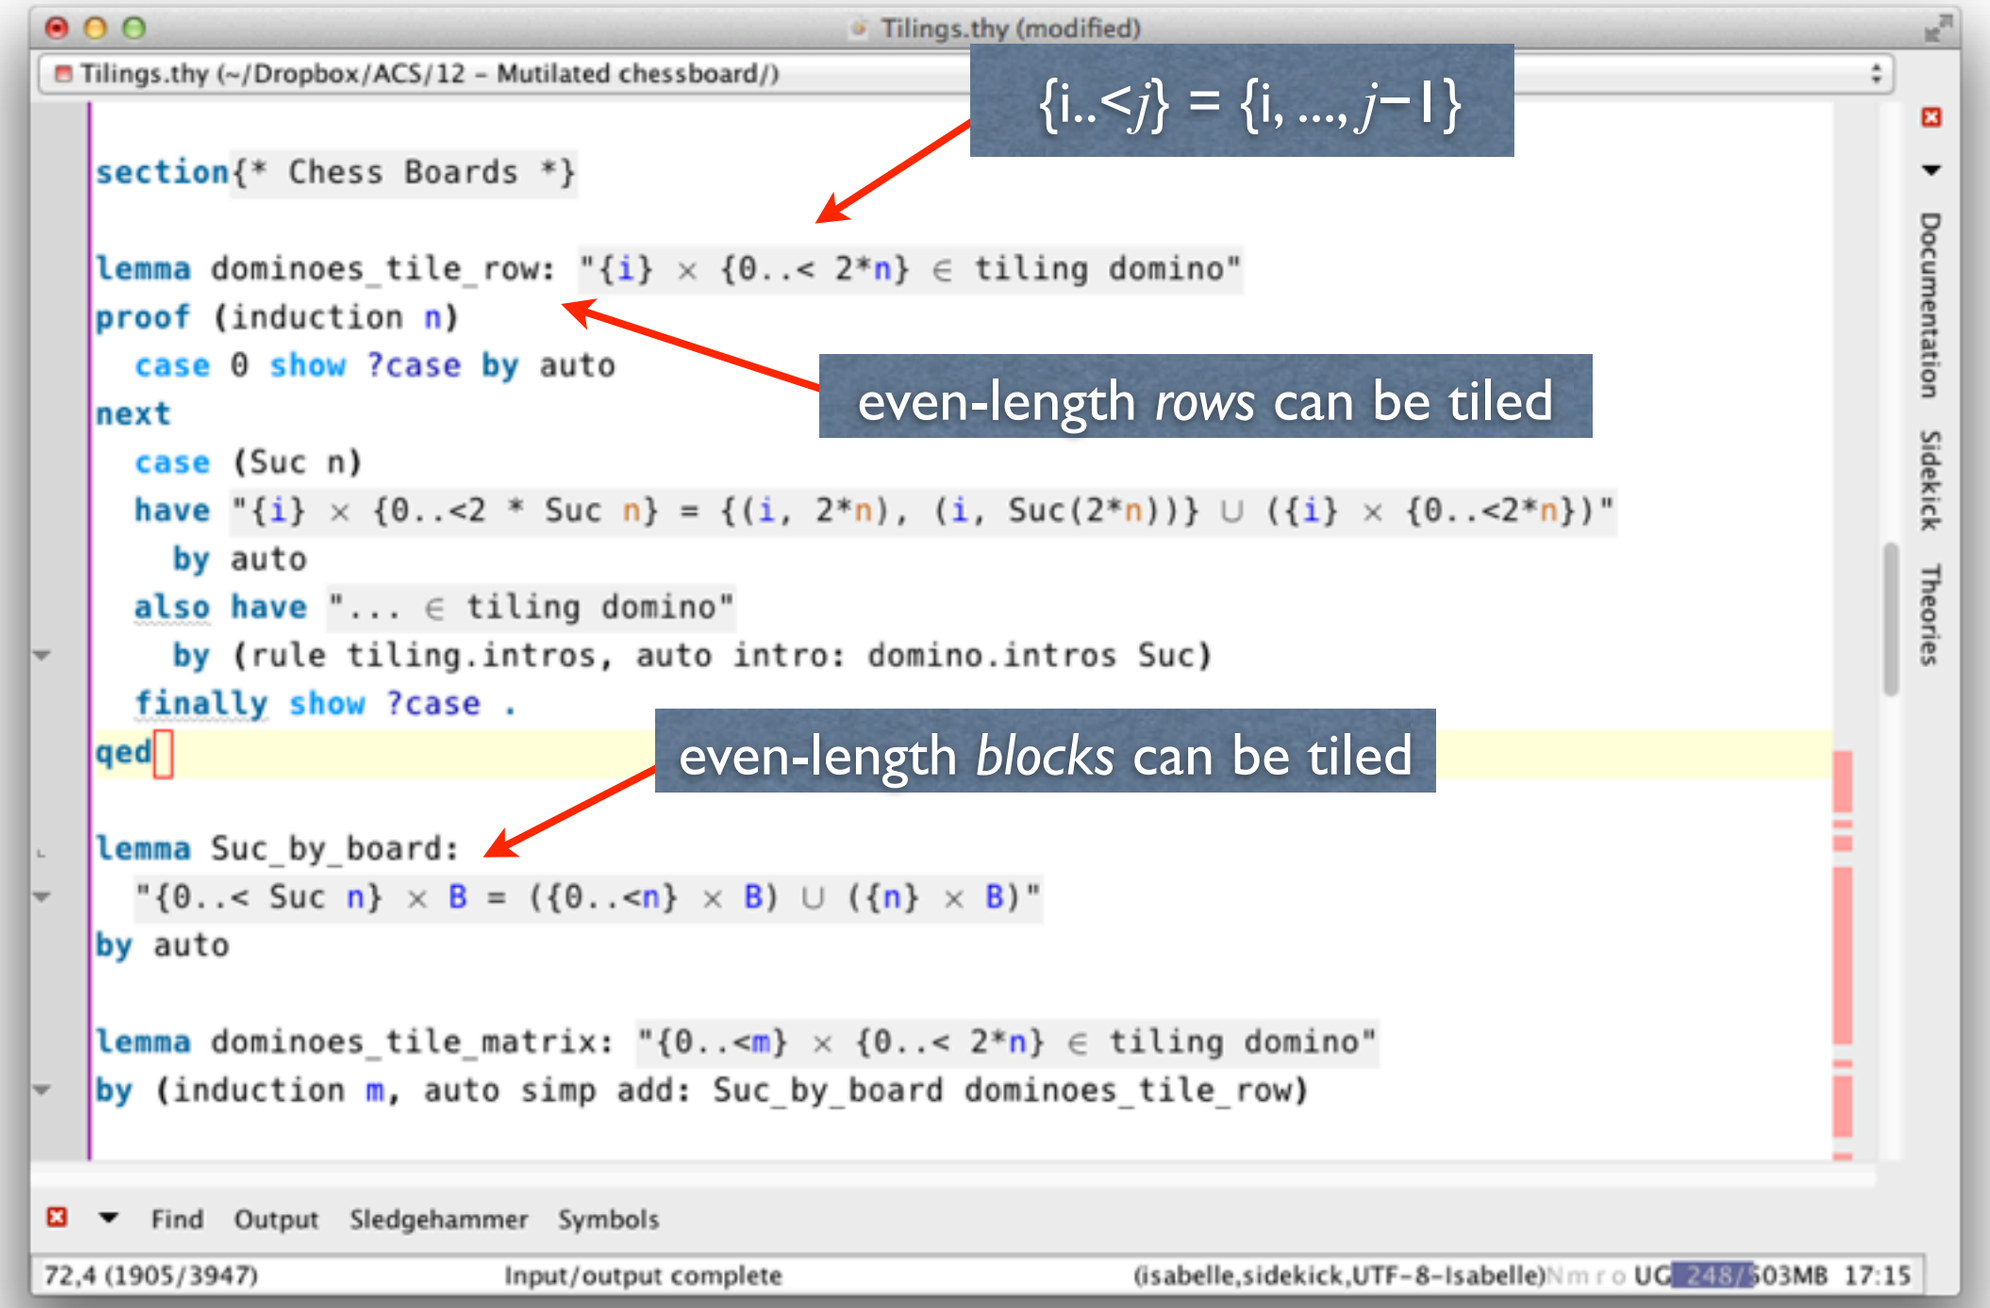

The first theorem states that any row of even length can be tiled by dominoes. In the inductive step, observe how the expression `{0..<2 * Suc n}` is rewritten to involve an explicit domino, `{(i, 2*n), (i, Suc(2*n))}`. Structured proofs make this sort of transformation easy, provided we are willing to write the desired term explicitly.

The alternative approach, of choosing rewrite rules that transform a term precisely as we wish, eliminates the need to write the intermediate stages of the transformation, but it can be more time-consuming overall. You know this other approach has been adopted if you see this sort of command:

```
apply (simp add: mult_assoc [symmetric] del: fact_Suc)
```

The theorem `mult_assoc` is given a reverse orientation using the attribute `[symmetric]`, while the theorem `fact_Suc` is removed from this simplifier call.

The induction at the bottom of this slide is an example of the alternative approach done correctly. We first prove a lemma to rewrite the induction step precisely as we wish: in other words, so that it will create an instance of `dominoes_tile_row`. The lemma is easily proved and the inductive proof is also easy.

For Tilings, #Whites = #Blacks

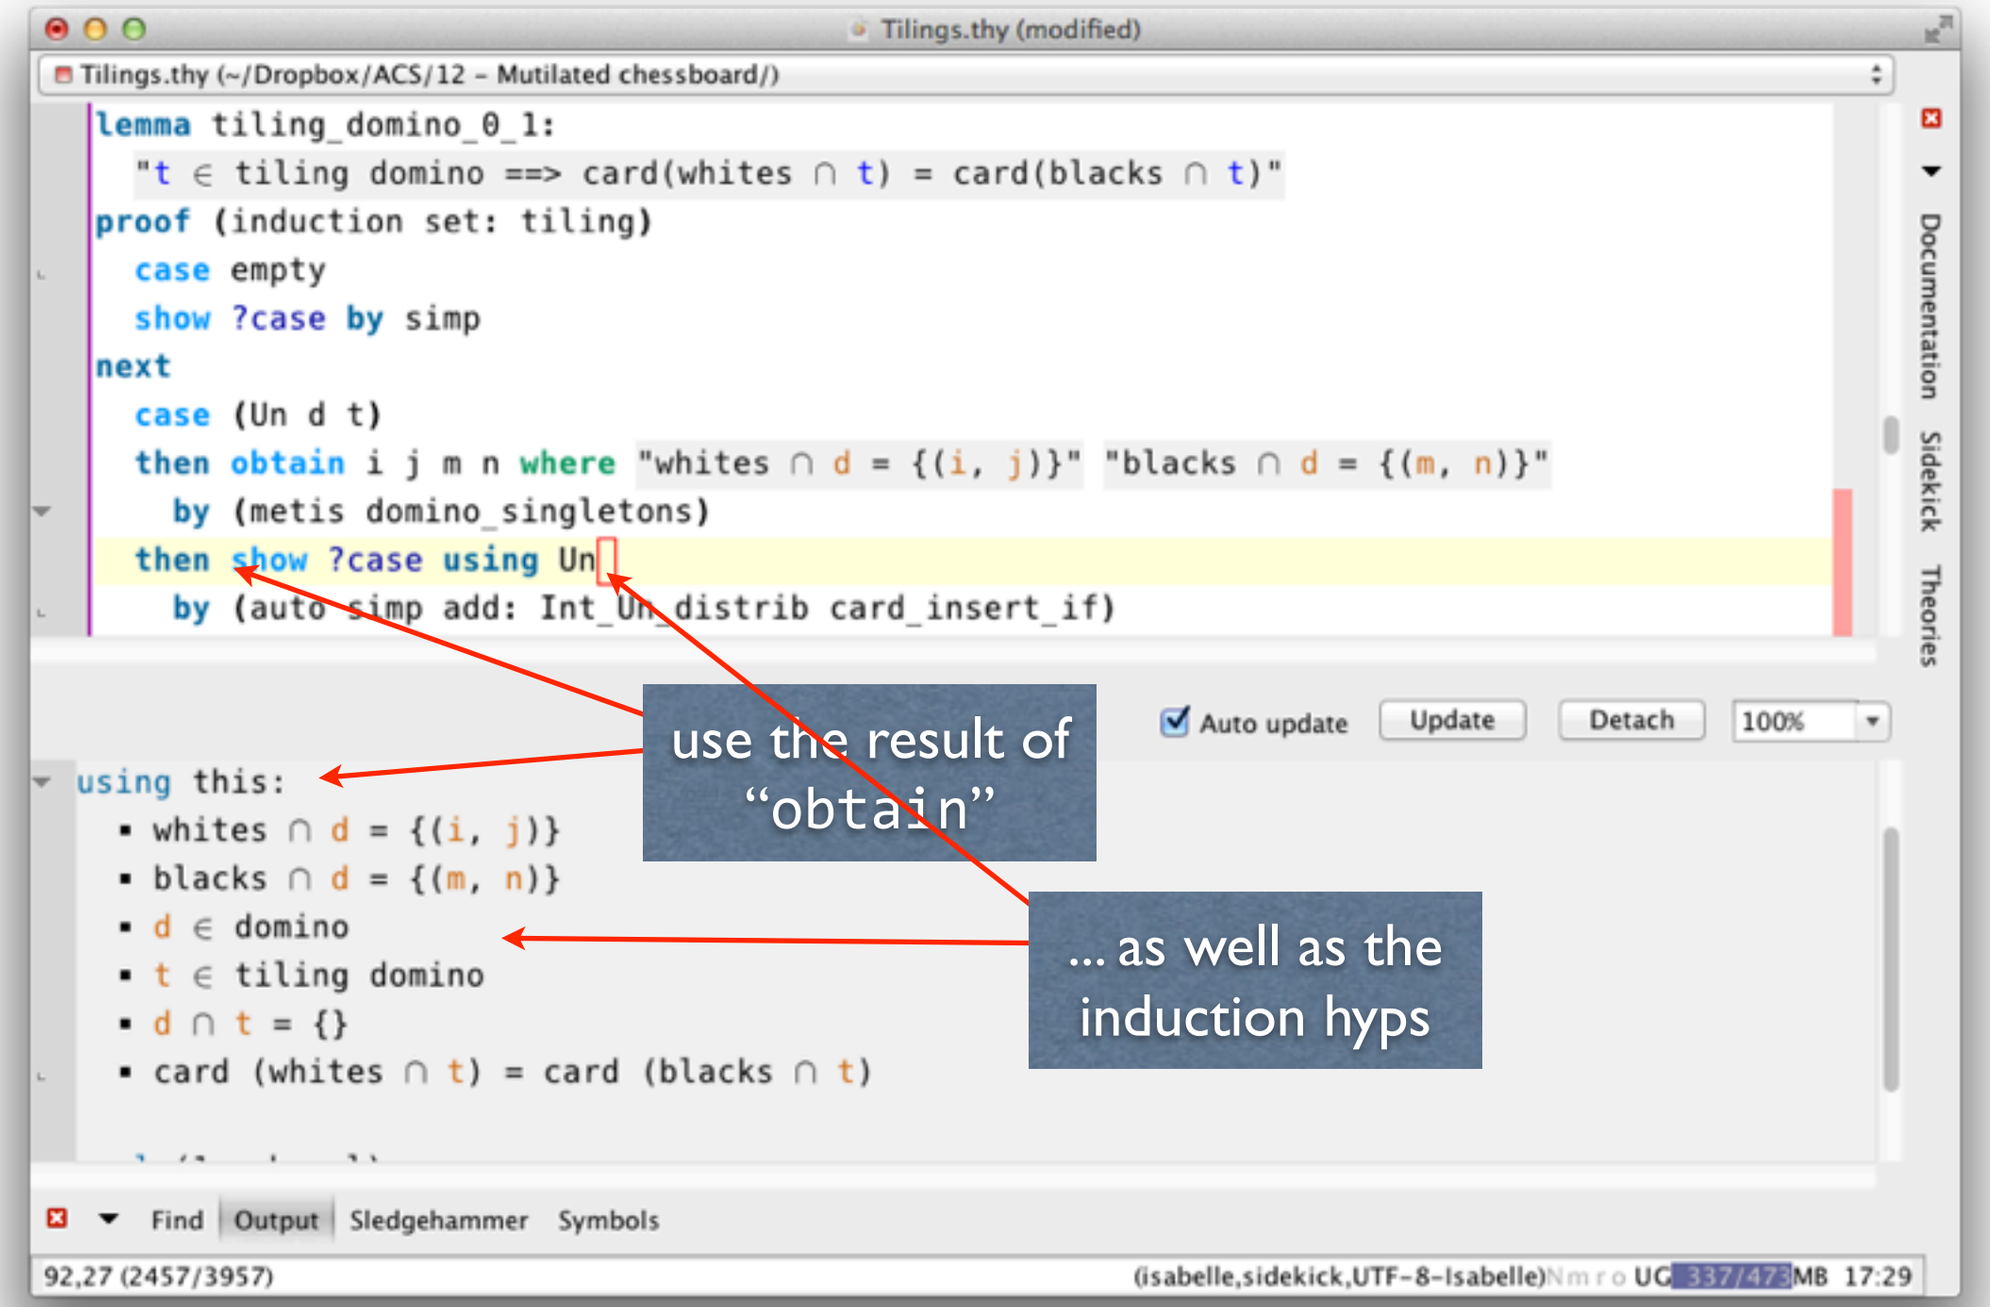

The crux of the argument is that any area tiled by dominoes must contain the same number of white and black squares. This statement is easily expressed using set theoretic primitives such as cardinality and intersection. The proof is by induction on tilings. It is trivial for the empty tiling. For a non-empty one, we note that the last domino consists of a white square and a black square, added to another tiling that (by induction) has the same number of white and black squares.

No Tilings for Mutilated Boards

default proof of a negation

accumulating some facts

Tilings.thy (modified)

Tilings.thy (~/Dropbox/ACS/12 – Mutilated chessboard/)

```
theorem gen_mutil_not_tiling:
  assumes "t ∈ tiling domino" "sqs ⊆ whites ∩ t" "sqs ≠ {}"
  shows "(t - sqs) ∉ tiling domino"
proof
  assume tm: "t - sqs ∈ tiling domino"
  have fsqs: "finite sqs" using assms
    by (metis Int_subset_iff finite_subset tiling_finite [OF domino_finite])
  hence c: "0 < card sqs" "0 < card (whites ∩ t)" using assms
    by (auto simp add: card_gt_0_iff)
  have "card (whites ∩ (t-sqs)) = card ((whites ∩ t) - sqs)"
    by (metis Int_Diff)
  also have "... < card (whites ∩ t)" using fsqs c assms
    by (auto simp add: card_Diff_subset)
  also have "... = card (blacks ∩ t)"
    by (blast intro: tiling_domino_0_1 assms)
  also have "... = card (blacks ∩ (t - sqs))"
    proof -
      have "blacks ∩ (t - sqs) = blacks ∩ t" using assms
        by (force simp add: coloured_def)
      then show ?thesis by simp
    qed
finally show False using tiling_domino_0_1 [OF tm] by auto
```

card (whites ∩ (t - sqs)) <
card (blacks ∩ (t - sqs))

Documentation Sidekick Theories

Find Output Sledgehammer Symbols

92,27 (2457/3957) (isabelle,sidekick,UTF-8-Isabelle) UG 247/486MB 17:30

The other crucial point is that if some white squares are removed, then there will be fewer white squares than black ones; although obvious to us, this proof requires the series of calculations shown on the slide. Once we have established this inequality, then it is trivial to show that the remaining squares cannot be tiled.

The Final Proof...

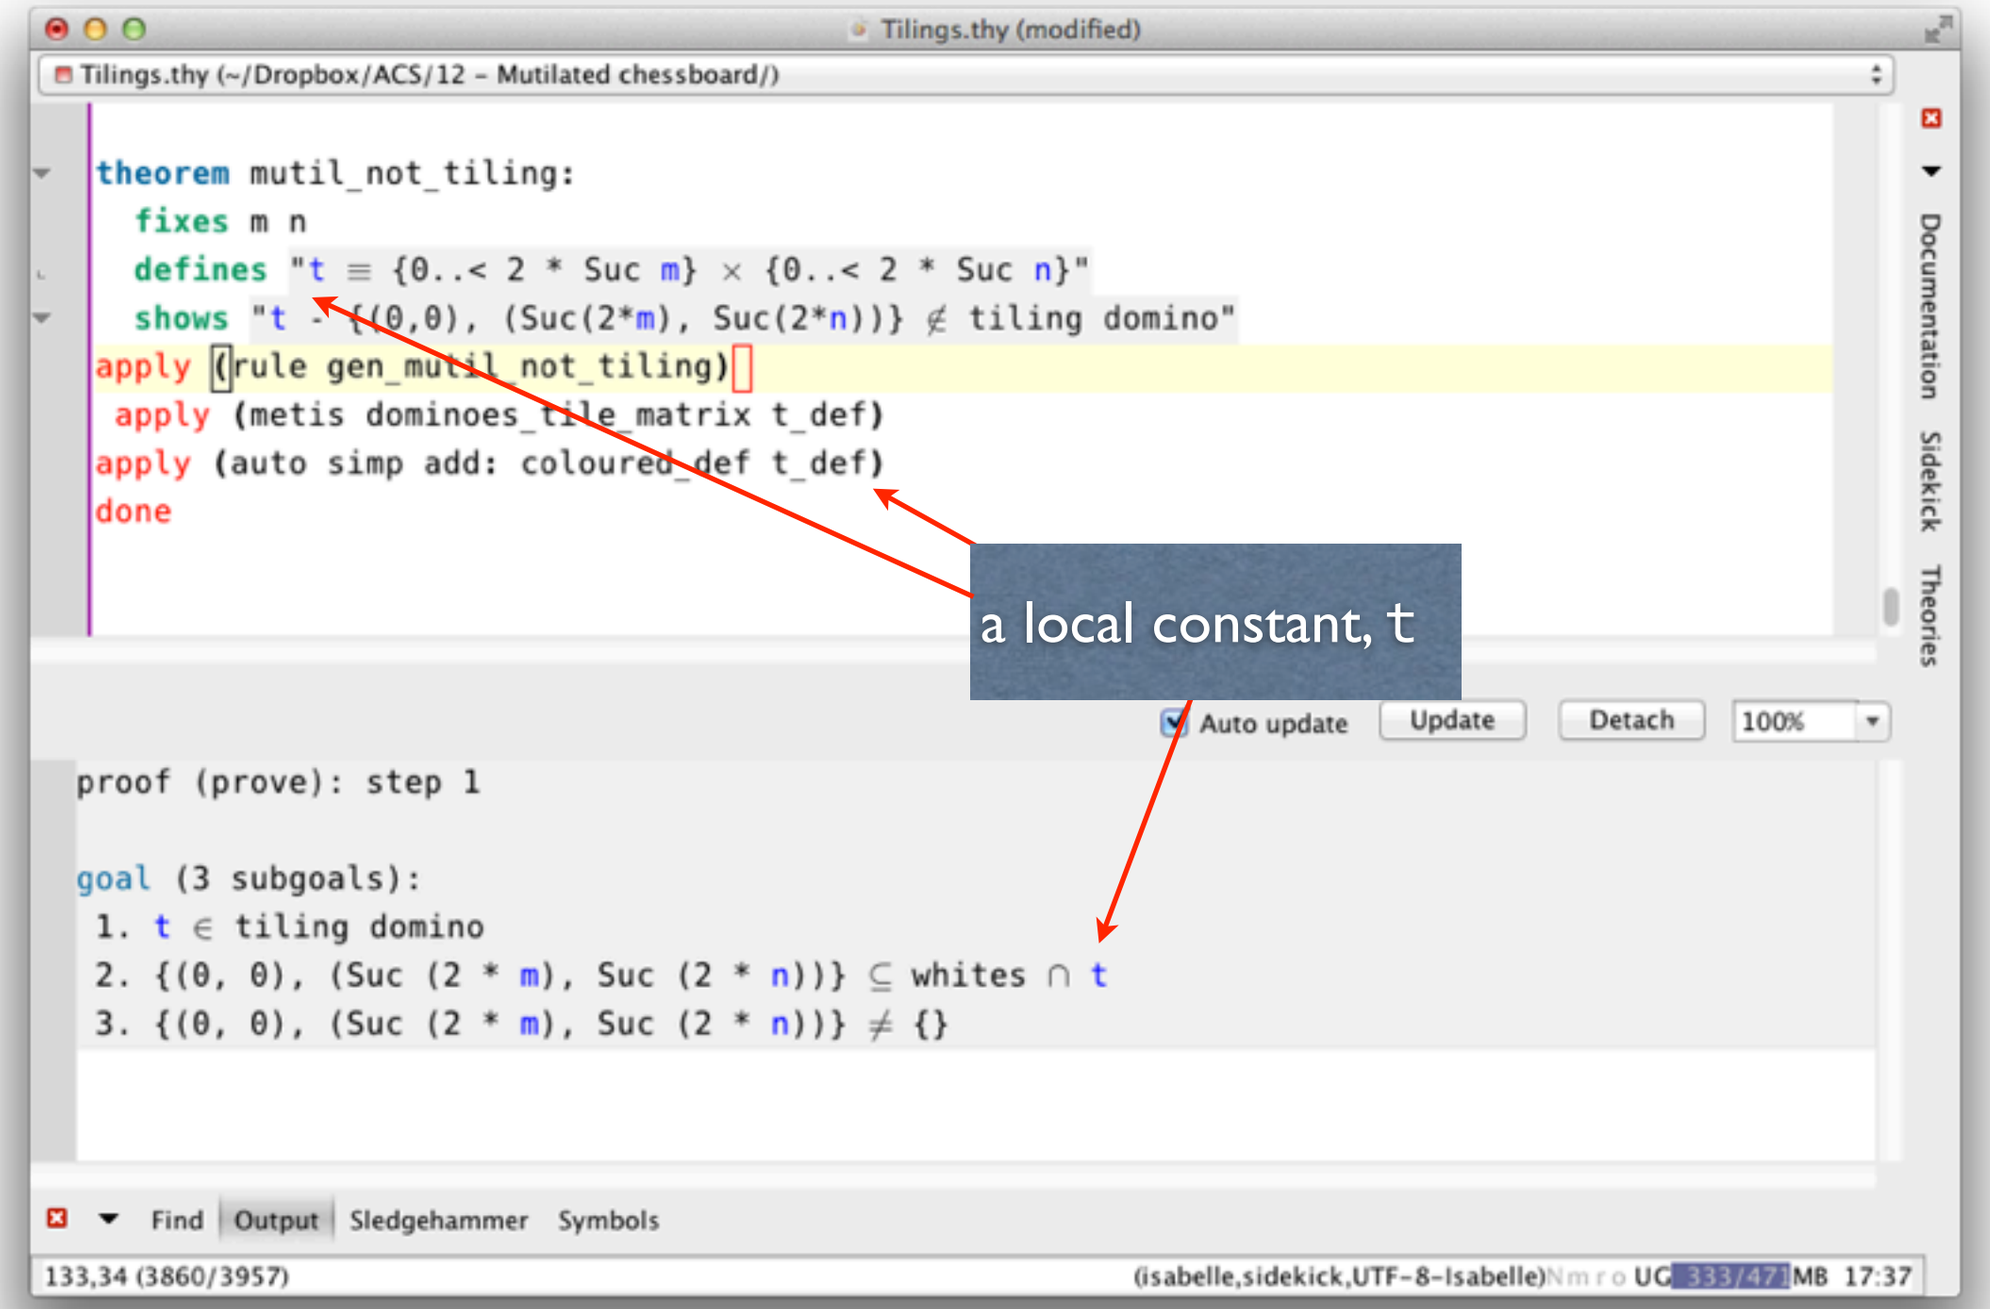

An 8 x 8 chess board can be generalised slightly, but the dimensions must be even (otherwise, the removed squares will not be white) and positive (otherwise, nothing can be removed).

Here we display yet another novelty: a “`defines`” element. Within the proof, `t` is a constant whose definition is available as the theorem `t_def`. But once the proof is finished, Isabelle stores a theorem that does not mention `t` at all.

The “fixes” element is necessary because otherwise the “`defines`” element will be rejected on the grounds that it has “hanging” variables (`m` and `n`) on the right-hand side.

The Result for Chess Boards

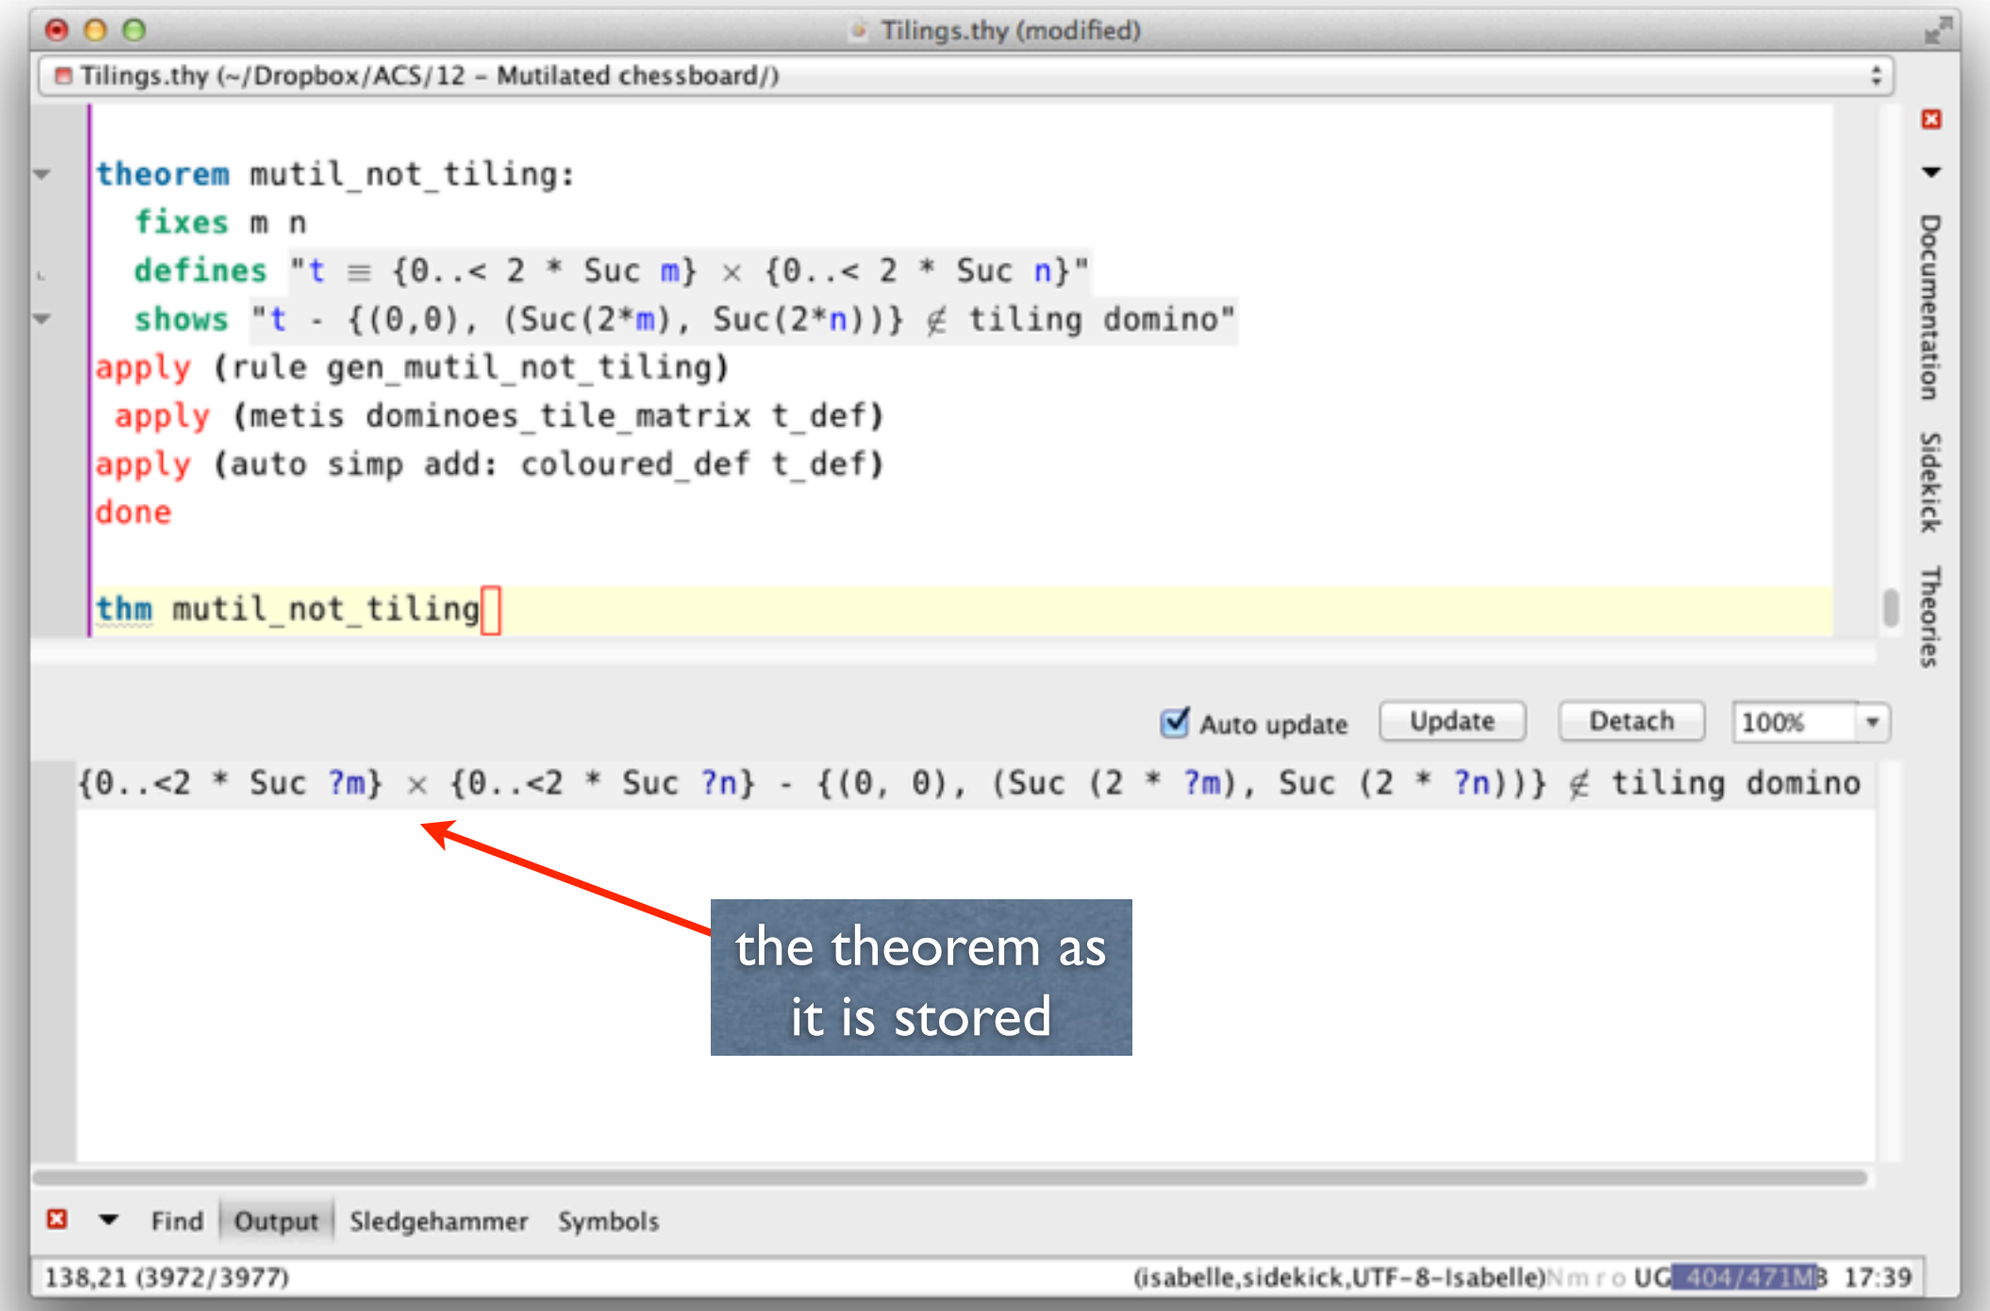

Note that the final theorem does not mention `defines` or `t`. All such definitions are expanded.

With Isabelle/jEdit, it's necessary to display the theorem explicitly using `thm`.

Finding Structured Proofs

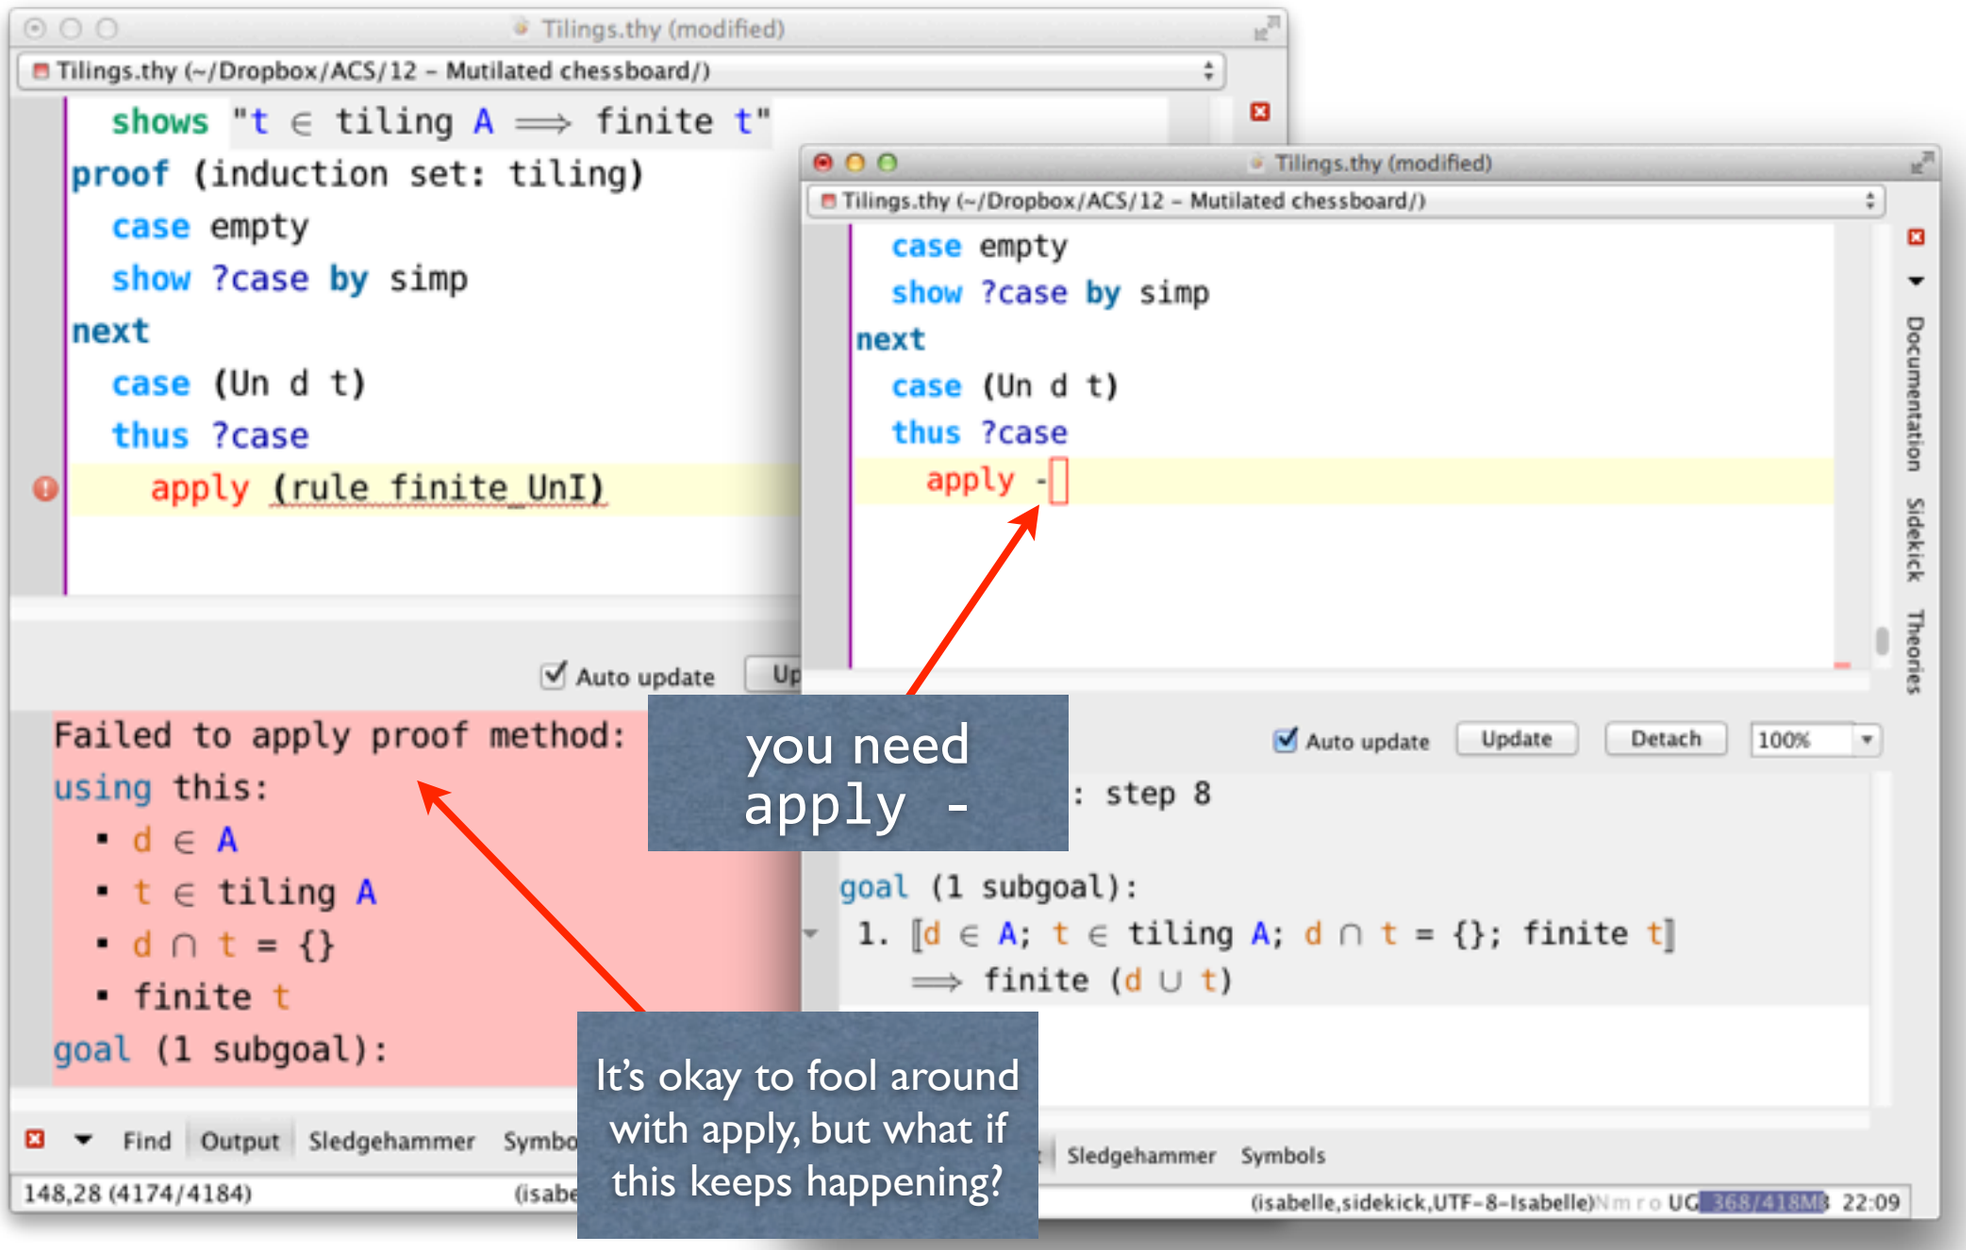

A common way to arrive at structured proofs is to look for a short sequence of `apply`-steps that solve the goal at hand. If successful, you can even leave this sequence (terminated by "`done`") as part of the proof, though it is better style to shorten it to a use of "`by`". Sometimes however almost everything you try produces an error message. The problem may be that you are piping facts into your proof using `then`/`hence`/`thus`/`using`. Some proof methods (in particular, "`rule`" and its variants) expect these facts to match a premise of the theorem you give to "`rule`". The simplest way to deal with this situation is to type `apply -`, which simply inserts those facts as new assumptions. It would be very ugly to leave `-` as a step in your final proof, but it is useful when exploring.

Other Facets of Isabelle

- *Document preparation*: you can generate LaTeX documents from your theories.
- *Axiomatic type classes*: a general approach to polymorphism and overloading when there are shared laws.
- *Code generation*: you can generate executable code from the formal functional programs you have verified.
- *Locales*: encapsulated contexts, ideal for formalising abstract mathematics.

See the *Tutorial*, section **4.2**, for an introduction to document preparation.

Locales are documented in the “Tutorial to Locales and Locale Interpretation” by Clemens Ballarin, which can be downloaded from Isabelle’s documentation page.

Axiomatic Type Classes

- Controlled overloading of operators, including + − × / ^ ≤ and even gcd
- Can define concept hierarchies abstractly:
  - Prove theorems about an operator from its axioms
  - Prove that a type belongs to a class, making those theorems available
- Crucial to Isabelle's formalisation of arithmetic

Axiomatic type classes are inspired by the type class concept in the programming language Haskell, which is based on the following seminal paper:

Philip Wadler and Stephen Blott. How to make ad-hoc polymorphism less ad hoc. In *16th Annual Symposium on Principles of Programming Languages*, pages 60–76. ACM Press, 1989.

A very early version was available in Isabelle by 1993:

Tobias Nipkow. Order-sorted polymorphism in Isabelle. In Gérard Huet and Gordon Plotkin, editors, *Logical Environments*, pages 164–188. Cambridge University Press, 1993.

More recent papers include the following:

Markus Wenzel. Type Classes and Overloading in Higher-Order Logic. *In*: Elsa L. Gunter and Amy P. Felty, *Theorem Proving in Higher Order Logics*. Springer Lecture Notes In Computer Science 1275 (1997), 307 - 322.

Lawrence C. Paulson. Organizing Numerical Theories Using Axiomatic Type Classes. *J. Automated Reasoning* **33** 1 (2004), 29–49.

Full documentation is available: see "Haskell-style type classes with Isabelle/Isar", which can be downloaded from Isabelle's documentation page, `http://www.cl.cam.ac.uk/research/hvg/Isabelle/documentation.html`

Code Generation

- Isabelle definitions can be translated to equivalent ML and Haskell code.
- Inefficient and non-executable parts of definitions can be replaced by equivalent, efficient terms.
- Algorithms can be verified and then executed.
- The method `eval` provides *reflection*: it proves equations by execution.

See “Code generation from Isabelle/HOL theories”, by Florian Haftmann; it can be downloaded from Isabelle’s documentation page.

The End

You know my methods. Apply them!

Sherlock Holmes