

EURO-MILS

Secure European Virtualisation for Trustworthy Applications in Critical Domains

Used Formal Methods



Project number	318353
Project acronym	EURO-MILS
Project title	EURO-MILS: Secure European Virtualisation for Trustworthy Applications in Critical Domains
Start date of the project	1 st October, 2012
Duration	36 months
Programme	FP7/2007-2013
Project website	www.euromils.eu



TECHNIKON
TECHNIKON



AIRBUS
GROUP

SYSGO
EMBEDDING INNOVATIONS

UNIVERSITÉ
PARIS
SUD

OPENSYNERGY

••T••Systems•

Open Universiteit
www.ou.nl



THALES TU/e

Technische Universiteit
Eindhoven
University of Technology

JEMM
research

Editors/Authors:

Holger Blasum, Oto Havle, Sergey Tverdyshev (SYSGO AG)

Contributors (ordered according to beneficiary numbers):

Sergey Tverdyshev, Oto Havle, Holger Blasum (SYSGO AG)

Bruno Langenstein, Werner Stephan (Deutsches Forschungszentrum für künstliche Intelligenz / DFKI GmbH)

Abderrahmane Feliachi, Yakoub Nemouchi, Burkhard Wolff (Université Paris Sud)

Cyril Proch (Thales Communications & Security SA)

Freek Verbeek (Open University of The Netherlands)

Julien Schmaltz (Technische Universiteit Eindhoven)

Further information on the EURO-MILS Project: <http://www.euromils.eu>



The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement number 318353.

Used Formal Methods

Whitepaper 2015

Executive Summary

This document consists of three chapters:

- ✓ **Chapter 1** describes how Isabelle/HOL works and how to use it in a certification process in a sound way.
- ✓ **Chapter 2:** Style Guide. It describes how to write Isabelle theories so that they are suitable for collaborative work and human readers in a certification context.
- ✓ **Chapter 3:** Compliance statement. We state how, in the EURO-MILS project, the developed theories are compliant with (1) and (2).



Contents

1	Using Isabelle/HOL in Certification Processes: A System Description and Mandatory Recommendations	1
1.1	Introduction	1
1.2	Common Criteria: Normative Context	2
1.2.1	Certification Level: Different Use of Formal Methods	3
1.2.2	Requirements Addressed by Formal or Semiformal Models	4
1.2.3	Formal Methods: Other Requirements Impacted	4
1.3	Isabelle/HOL: Architecture, Language and Methodology	5
1.3.1	The Isabelle System Architecture	5
1.3.2	Isabelle and its Meta-Logic	7
1.3.3	Foundations of HOL and its Specification Constructs	8
1.3.4	Isabelle Proofs	18
1.3.5	Isabelle/HOL System Features	20
1.4	Methodological Recommendations for the Evaluator	21
1.4.1	On the Use of SML	22
1.4.2	Axioms and Bogus-Proofs	23
1.4.3	On the Use of External Provers	23
1.5	Extensions of Isabelle: Guidelines for the Evaluator	24
1.5.1	Example: Isabelle/Simpl	24
1.5.2	Example: The HOL-TestGen Test-Generation System	25
1.5.3	By the Way: Test vs. Proof	26
1.6	Recommendations for CC Certifications	26
1.6.1	A Refinement Based Approach for CC Evaluation	26
1.7	Summary	27
1.7.1	Background References	27
1.7.2	Concluding Remarks and a Summary	28
2	Style Guide	30
2.1	Introduction	30
2.2	Rules	30
2.2.1	Basics	30
2.2.2	Modeling Style	31
2.2.3	Formal Content	32
2.2.4	Layout Principles	34
2.3	Conclusion	35
3	Compliance Statement	36
3.1	Compliance to Section 1.7.2	36
3.2	Compliance to Section 2	37
	Bibliography	38

Chapter 1

Using Isabelle/HOL in Certification Processes: A System Description and Mandatory Recommendations

Chapter Authors: Yakoub Nemouchi, Abderrahmane Feliachi, Burkhart Wolff and Cyril Proch

Abstract: Interactive theorem proving is a technology of fundamental importance for mathematics and computer science. It is based on expressive logical foundations and implemented in a highly trustable way. Applications include very large mathematical proofs and semi-automated verifications of complex soft- and hardware systems. The architecture of contemporary interactive provers such as Coq, Isabelle, or the HOL family goes back to the influential LCF system from 1979, which has pioneered key principles like *correctness by construction* for primitive inferences and definitions, *free programmability* in userspace via SML, and *toplevel command interaction*.

The Isabelle System developed into one of the top 5 systems for the logically consistent development of formal theories. In particular the instance of the Isabelle system with higher-order logic called Isabelle/HOL is therefore a natural choice as a formal methods tool as required by the Common Criteria on the higher assurance levels EAL5 to EAL7.

The purpose of this paper is to give a brief introduction into the system, an overview over the methodology and its tool support, and high-level mandatory guidelines for evaluators of certifications using Isabelle. This paper is intended to be a complement of a similar text by French certification authorities [Jae08].

Keywords: Formal Methods; Certification; Compliance; Common Criteria; Isabelle; HOL

1.1 Introduction

Formal methods describe a set of mathematically based techniques and tools for specification, analysis and verification of computer systems. They are mainly used to describe and to verify, in a logically consistent way, some properties of these systems. The formal specification and verification approaches usually rely on some underlying logic. The logical foundation of theorem provers makes them a very convenient basis of any formal development, where the specification and the verification activities can be gathered in one formal environment.

Interactive theorem proving is a technology of fundamental importance for mathematics and computer science. It is based on expressive logical foundations and implemented in a highly trustable way. Applications include huge mathematical proofs and semi-automated verifications of complex hard- and software systems. The architecture of contemporary interactive provers such as Coq [Wie06, §4], Isabelle [Wie06, §6] or the HOL family [Wie06, §1] goes back to the

influential LCF system [MW79] from 1979, which has pioneered key principles like correctness by construction for primitive inferences and definitions, free programmability in userspace via SML, and toplevel command interaction.

Recently, theorem provers have been widely used in the area of computer systems security and certification and, for instance, in Common Criteria. The Common Criteria (CC) [Mem06] is a well-known and recognized computer security certification standard. The standard is centered around the role of the *developer*, who provides implementation but also “artefacts of compliance with the level of security targeted”, while the *evaluator* “confirms the compliance of the information supplied” as well as determines “completeness, accuracy and quality” of the deliverables.

Especially wrt. “completeness, accuracy and quality” of specifications and proofs, formal methods and especially mechanically proof checking techniques can push the trust and the reproducibility of the results to levels not obtainable by a human certification expert alone. This explains why at its higher assurance levels, the CC requires the use of formal methods for specification and verification. A well-established formal specification formalism must be used to model the system and the different security policies. A reliable theorem prover is needed to prove and verify different properties of the specification. Recent theorem provers offer rich and powerful formal environments that are very suitable for both activities.

Among the important number of theorem provers available nowadays, we concentrate on the Isabelle theorem prover¹. Following [Hal08], the Isabelle System, developed into one of the top five systems for the logically consistent development of formal theories. In particular the instance of the Isabelle system with higher-order logic called Isabelle/HOL is therefore a natural choice as a formal methods tool as required by the Common Criteria on the higher assurance levels EAL5 to EAL7.

The purpose of this paper is to bring together a body of system information that is generally known in the Isabelle community, but largely scattered in system documentations and papers. This includes a brief introduction into the system, a general overview over the methodology and covers certain aspects of the tool support. The paper culminates in some high-level mandatory guidelines and recommendations for both developers and evaluators of certification documents using Isabelle. It attempts to be a complement to the previous document written by Jaeger [Jae08].

The paper proceeds as follows: at first in Sec. 1.2, we give some general information from Common Criteria standard about formal methods, modeling and associated requirements. In Sec. 1.3, we provide a guided tour over the Isabelle system, while in Sec. 1.4, we refer to methodological issues of Isabelle/HOL leading to recommendations for evaluators. In Sec. 1.5 we chose two major extensions of Isabelle, one for code-verifications, one for model-based testing, and discuss their advantages and limits in a high-level certification process. The final discussion contains a little survey on publications on the topic as well as a summary for evaluators.

1.2 Common Criteria: Normative Context

For high levels of certification (i.e. for EAL5 to EAL7) in the Common Criteria [Mem06] some requirements introduce the use of formal methods at diverse phases of the design process. Regarding to the level of security target required, the utilisation of formal methods match different objectives.

¹At time writing, the current version is Isabelle2013-2.

1.2.1 Certification Level: Different Use of Formal Methods

The next table resumes for each level the main requirements and impacts from formal methodology point of view.

level EAL	Objective
EAL5	This EAL represents a meaningful increase in assurance from EAL4 (methodically designed, tested, and reviewed) by requiring semiformal design descriptions, a more structured (and hence analysable) architecture, and improved mechanisms and/or procedures that provide confidence that the TOE will not be tampered with during development.
EAL6	This EAL represents an important increase in assurance from EAL5 (semi-formally designed and tested) by requiring more comprehensive analysis, a structured representation of the implementation, more architectural structure (e.g. layering), more comprehensive independent vulnerability analysis, and improved configuration management and development environment controls.
EAL7	This EAL represents a meaningful increase in assurance from EAL6 (semi-formally verified design and tested) by requiring more comprehensive analysis using formal representations and formal correspondence, and comprehensive testing.

In addition to the normative definitions of the security levels, the CC standard defines the possibility of intermediate levels of security when a requirement is evaluated at a higher level than required by the level targeted. The addition of the symbol "+" represents this kind of evaluation (for example EAL4+).

With regard to high level certifications, the requirements on formal methods are more and more intrusive and the models are more and more detailed (from a high level architecture for EAL5 to a structured formal design for EAL7).

The Common Criteria defines two different roles the *developer* and *evaluator*. These two different roles shall comply different requirements of CC or more precisely, each requirement of CC is declined in different actions for developer and evaluator. From a general point of view, the developer shall:

- realize the design, the documentation, the implementation and the validation of the target.
- provide artifacts and elements of compliance with the level of security targeted.

In another hand the evaluator shall:

- confirm the compliance of the information supplied (by the developer) with requirements of the security level,
- determine the completeness, the accuracy and in a general manner the quality of the deliverables.

This document is intended to detail these two tasks of an evaluator with respect to implementations and validations done with the Isabelle/HOL system.

1.2.2 Requirements Addressed by Formal or Semiformal Models

With regard to high level certifications, the main requirements addressed by the use of formal methods are:

- ADV_SPM.1 requires a formal TOE *security policy model* (SPM for short). This model is generally a high level model which captures the main security properties and abstract behavior of the target.
- ADV_FSP.6 requires a semi-formal *functional specification* (FSP for short) with an additional formal specification. This second constraint concerns an intermediate functional level of design and it is considered as pertinent from a formal point of view, to manage a formal model (and not a semi-formal) which is a refinement of the initial model defines for ADV_SPM.1. The use of this intermediate formal model is efficient to define a formal specification.
- ADV_TDS.6 requires a complete semi-formal and modular design with high-level (*TOE*) *design specification* (TDS). This final design requirement introduces the architecture of the target and the notions of modules and interfaces. The main objective is to define and simply specify the structure of the design and provide a proof of correspondence between specifications of the subsystems and the functional specification.

This simple overview of some CC requirements implies that a formal approach based on a formal refinement definition is compliant to assure consistency between the different models and the diverse views and objectives considered in these requirements.

1.2.3 Formal Methods: Other Requirements Impacted

Other requirements for high level certification (EAL7) are not directly connected with formal methods but they can be addressed by the use of formal methods (see chapter 1.5.2):

- ATE_FUN.2: The objectives are for the developer to demonstrate that the tests in the test documentation are performed and documented correctly, and to ensure that testing is structured such as to avoid circular arguments about the correctness of the interfaces being tested. Although the test procedures may state pre-requisite initial test conditions in terms of ordering of tests, they may not provide a rationale for the ordering. An analysis of test ordering is an important factor in determining the adequacy of testing, as there is a possibility of faults being concealed by the ordering of tests.
- AVA_VAN.5: A methodical vulnerability analysis is performed by the evaluator to ascertain the presence of potential vulnerabilities. The evaluator performs penetration testing, to confirm that the potential vulnerabilities cannot be exploited in the operational environment for the TOE. Penetration testing is performed by the evaluator assuming an attack potential of High.

The generation of test cases from a formal model can be an interesting approach to optimize the efforts of the modeling and the formal proof: this formal specification-based testing approach is not a classic approach in industrial world but seems compliant with the two previous CC requirements.

1.3 Isabelle/HOL: Architecture, Language and Methodology

In the following, we will discuss the two questions:

How is Isabelle built?

How should Isabelle be used?

In the context of certifications of critical hard- and software systems, an understanding of its architecture and the underlying methodology may help to understand why Isabelle, if correctly used, can be trusted to a significantly higher extent than conventional software, even more than other automated theorem provers (in fact, Sascha Böhme’s work on proof reconstruction [BW10] inside Isabelle revealed errors the SMT solver Z3[MB08] that is perhaps the most tested conventional system currently on the market ...). Of course, Isabelle as software “contains errors”. However, its architecture is designed to exclude that errors allow to infer logically false statements, and methodology may help to exclude that correctly inferred logical statements are just logical artifacts, or logically trivial statements, which can be impressive stunts without any value.

1.3.1 The Isabelle System Architecture

We will describe the layers of the system architecture bottom-up one by one, following the diagram Fig. 1.1.

The foundation of system architecture is still the Standard ML (SML,[MTM97]) programming environment; the default PolyML implementation

<http://www.polyml.org> supports nowadays multi-core hardware which is heavily used in recent versions for parallel and asynchronous proof checking when editing Isabelle theories.

On top of this, the logical kernel is implemented which comprises type-checking, term-implementations and the management of global contexts (keeping, among many other things, signature information and basic logical axioms). The kernel provides the abstract data-types `thm`, which is essentially the triple (Γ, Θ, ϕ) , written $\Gamma \vdash_{\Theta} \phi$, where Γ is a list of *meta-level assumptions*, Θ the *global context*, containing, for example, the signature and core axioms of HOL and the signature of group operators, and a *conclusion* ϕ , i. e. a formula that is established to be derivable in this context (Γ, Θ) . Intuitively, a `thm` of the form $\Gamma \vdash_{\Theta} \phi$ is stating that the kernel certifies that ϕ has been derived in context Θ from the assumptions Γ .

There are only a few operations in the kernel that can establish `thm`’s, and the system correctness depends *only* on this trusted kernel. On demand, these operations can also log proof-objects that can be checked, in principle, independently from Isabelle; in contrast to systems like Coq, proof objects do play a less central role for proof checking which just resides on the inductive construction of `thm`’s by kernel inferences shown, for example, in [PP10].

On the next layer, proof procedures were implemented - advanced tactical procedures that search for proofs based on higher-order rewriting like `simp`, tableau provers such as `fast`, `blast`, or `metis`, and combined procedures such as `auto`. Constructed proofs were always checked by the inference kernel.

The next layer provides major components — traditionally called *packages* — that implement the *specification constructs* such as *constant definitions* *type abbreviations*, *type definitions*, etc., as discussed in Sec. 1.3.3 in more detail. Packages may also yield connectors to external provers (be it via the `sledgehammer` interface or via the `smt` interface to solvers such

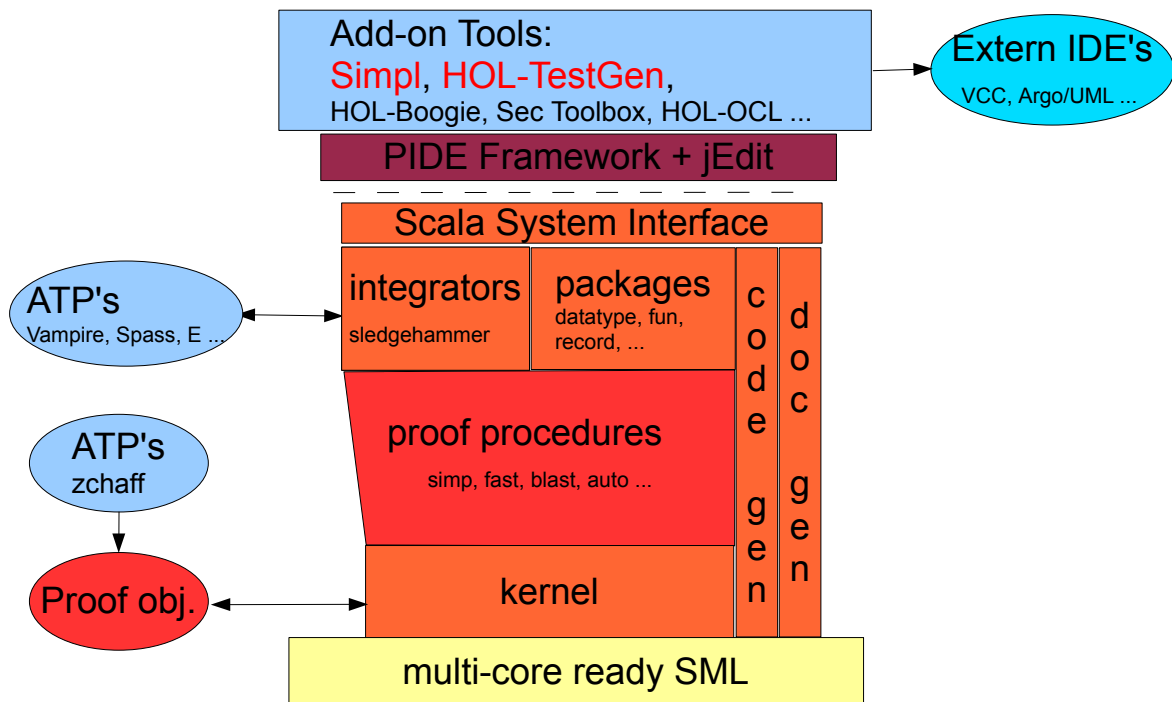


Figure 1.1: The diagram shows the different layers like execution environment, kernel, tactical level and proof-procedures, component level (providing external prover integration like Z3, specification components, and facilities like the code generator, the Scala API to the system bridging to the JVM-World, and the Prover-IDE (PIDE) layer allowing for asynchronous proof and document checking.

as Z3), machinery for (semi-trusted) code-generators as well as the Isar-engine that supports structured-declarative and imperative “apply style” *proofs* described in Sec. 1.3.4.

The Isar - engine [Wen02] parses specification constructs and proofs and dispatches their treatment via the corresponding packages. Note that the Isar-Parser is configurable; therefore, the syntax for specification constructs like constant definitions can be modified and adapted, as well as the automated proofs that derive from them the characterizing properties of a datatype (distinctness and injectivity of the constructors, as well as induction principles) as thm ’s available in the global context Θ thereafter. As we will see in Sec. 1.3.3, specification constructs represent the heart of the methodology behind Isabelle: new specification elements were only introduced by “conservative”; i.e. logically safe mechanisms that maintain the logical consistency of the theory by construction. Packages provide the technical support for these specification constructs; internally these constructs introduce declarations and axioms of a particular form and enforce the user to provide proofs for methodological side-conditions (like the non-emptiness of the carrier set defining a new type or termination for a recursive function definition).

We mention the last layer mostly for completeness: Recent Isabelle versions possess also an API written in Scala, which gives a general system interface in the JVM world and allows

to hook-up Isabelle with other JVM-based tools or front-ends like the jEdit client. This API, called the “Prover IDE” or “PIDE” framework, provides an own infrastructure for controlling the concurrent tasks of proof checking. The jEdit-client of this framework is meanwhile customized as default editor of Isabelle *sessions*; thus jEdit became the default user-interface the user has primarily access to. PIDE and its jEdit client manage collections of theory documents containing sequences of specification constructs, proofs, but also structured text, code, and machine-checked results of code-executions. It is natural to provide such theory documents as part of a CC evaluation documentation.

1.3.2 Isabelle and its Meta-Logic

The Isabelle kernel natively supports minimal higher-order logic called *Pure*. It supports for just one logical type *prop* the meta-logical primitives for implication $_ \Longrightarrow _$ and universal quantification $\bigwedge x. P\ x$. The meta-logical primitives can be seen as the constructors of *rules* for various logical systems that can be represented inside Isabelle; a conventional “rule” in a logical textbook:

$$\frac{A_1 \cdots A_m}{C} \quad \text{where } x_1 \dots x_n \text{ are free variables} \quad (1.1)$$

can be directly represented via the built-in quantifiers \bigwedge and the built-in implication \Longrightarrow as follows in the Isabelle core logic *Pure*:

$$\bigwedge x_1 \dots x_n. A_1 \Longrightarrow \dots \Longrightarrow A_m \Longrightarrow C \quad (1.2)$$

... where the variables x_1, \dots, x_n are called *parameters*, the premises A_1, \dots, A_m *assumptions* and C the conclusion; note that \Longrightarrow binds to the right. Also more complex forms of rules as occurring in natural deduction style inference systems like:

$$\frac{\begin{array}{c} [A] \\ \vdots \\ B \end{array}}{A \rightarrow B} \quad (1.3)$$

can be represented by $(A \Longrightarrow B) \Longrightarrow A \rightarrow B$. Thus, the built-in logic provided by the Isabelle Kernel is essentially a language to describe (systems of) logical rules and provides primitives to instantiate, combine, and simplify them. Thus, Isabelle [NPW02] is a *generic* theorem prover. New object logics can be introduced by specifying their syntax and natural deduction inference rules. Among other logics, Isabelle supports first-order logic, Zermelo-Fraenkel set theory and the instance for Church’s higher-order logic HOL. Moreover, Isabelle is also a generic system framework (roughly comparable with Eclipse) which offers editing, modeling, code-generation, document generation and of course theorem proving facilities; to the extent that some users use it just as programming environment for SML or to write papers over checked mathematical content to generate L^AT_EX output. Many users know only the theorem proving language *isar!* for structured proofs and are more or less unaware that this is a particular configuration of the system, that can be easily extended. Note that for all of the aforementioned specification constructs and proofs there are specific syntactic representations in *isar!*.

Higher-order logic (HOL) [Chu40, And86, And02] is a classical logic based on a simple type system. It is represented as an instance in Pure. HOL provides the usual logical connectives like $_ \wedge _$, $_ \rightarrow _$, $\neg _$ as well as the object-logical quantifiers $\forall x. P\ x$ and $\exists x. P\ x$; in contrast to first-order logic, quantifiers may range over arbitrary types, including total functions $f :: \alpha \Rightarrow \beta$. HOL is centred around extensional equality $_ = _ :: \alpha \Rightarrow \alpha \Rightarrow \text{bool}$. HOL is more expressive than first-order logic, since, e. g., induction schemes can be expressed inside the logic. Being based on a polymorphically typed λ -calculus, HOL can be viewed as a combination of a programming language like SML or Haskell, and a specification language providing powerful logical quantifiers ranging over elementary and function types.

Isabelle/HOL is the session based on the embedding of HOL into Isabelle/Pure. Note that the simple-type system as conceived by Church for HOL has been extended by Hindley/Milner style polymorphism with type-classes similar to Haskell [WB89, Wen97].

1.3.3 Foundations of HOL and its Specification Constructs

The core of the logic is done via an axiomatization of the core concepts like equality, implication, and the existence of an infinite set, the rest of the library is derived from this core by logically safe (“conservative”) extension principles which are syntactically identifiable specification construction in Isabelle files. In the following, we will briefly describe the axiomatic foundation of Isabelle/HOL and describe the most common conservative extension principles.

Since the core of the HOL logic is given by an axiomatization, the question of its consistency must be settled. The first “authoritative” book that gives a consistency proof of HOL is Gordon/Melham’s book [GM93]. In chapter 15 (page 193 ff), a semi-informally described model of a Universe U is presented which must be closed under non-empty-subsets, products, functions, power sets, existence of an infinite set and choice. It is informally (but convincingly) argued that U can be built with ZFC minus the replacement axiom.² For the core of the HOL logic, an interpretation function is given in [GM93], pages 196 ff, into U thus providing a model for the HOL axiomatization (and thus its consistency proof relative to ZFC minus replacement). This interpretation provides the modern semantic understanding of Standard-HOL, i. e. standard models for HOL with parametric polymorphism. Note, however, that Isabelle type-classes are not covered by this model (the more powerful Isabelle locale’s, however, is a specification construct that can be reduced to basic HOL).

The text [GM93] appeared meanwhile online in a slightly revised version:

Norrish, M., Slind, K., et al.: The HOL System: Logic, 3rd edn., <http://hol.sourceforge.net/documentation.html>

The Gordon/Melham book [GM93] also contains first (semi-formal) proofs about the conservativity of machine-supported conservative extension schemes based on the HOL-model in U . Four “classic” core constructions were defined and analyzed:

- constant definition
- constant specification
- type definition
- type specification.

²This part of the text is attributed to Andrew Pitts.

Logical conservativity means: whenever a theory T was consistent, its theory extension T' is also consistent.

Meanwhile, there had been a number of followup papers attempting to formalize these proofs of conservativity of the classical extension schemes. The most notable references are:

- John Harrison: *Towards self-verification of HOL Light*. In: International Joint Conference on Automated Reasoning (IJCAR). Volume 4130 of the series Lecture Notes in Computer Science. Springer (2006).
- A more recent paper (addressing the problem of formal proof of conservativity of the above extension schemes) is: Ramana Kumar, Rob Arthan, Magnus O. Myreen, and Scott Owens: *HOL with Definitions: Semantics, Soundness, and a Verified Implementation*. In: International Conference on Interactive Theorem Proving (ITP), Volume 8558 of the series Lecture Notes in Computer Science pp 308-324. Springer (2014).

The methodology of HOL systems is based on the idea that apart from the core of the HOL logic, which is necessarily axiomatized, the library (comprising a theory of products, sets, lists, maps, number-theories, etc.) is built by solely by the above-listed conservative extensions, for which special machine-support is provided. This machine support is extended to data-types, inductive definitions, and (extensible) records, which internally use the four classic ones. This machine support is represented in Fig. 1.1 under the label “packages”.

The Presentation of the core HOL in Isabelle.

In the entire library (so the Isabelle session “HOL” which is also referred to as “Main” in theory imports), there are only 11 axioms in form of foundational axioms of the HOL-logic. These 11 axioms are equivalent to the classical ones of [GM93]:

1. The equality symbol is axiomatized as an equality, i.e. it is reflexive, extensional, and satisfies the Leibniz-property (equals can be replaced by equals in any context P). The Hilbert-Operator is bound to choose the value characterized by equality:

```
axiomatization where
  refl: t = (t::α) and
  subst: s = t  $\implies$  P s  $\implies$  P t and
  ext: ( $\bigwedge$  x::α. (f x ::β) = g x)  $\implies$  ( $\lambda$ x. f x) = ( $\lambda$ x. g x) and
  the_eq_trivial: (THE x. x = a) = (a::'a)
```

2. The following axioms establish a relation between implication and rule formation, and between implication and equality, as well as True, $\forall x. P x$ and False and (which are abbreviations for $((\lambda x::\text{bool}. x) = (\lambda x. x))$, $(P = (\lambda x. \text{True}))$ and $(\forall P. P)$, respectively):

```
axiomatization where
  impI: (P  $\implies$  Q)  $\implies$  P  $\rightarrow$  Q and
  mp: P  $\rightarrow$  Q  $\implies$  P  $\implies$  Q and
  iff: (P $\rightarrow$ Q)  $\rightarrow$  (Q $\rightarrow$ P)  $\rightarrow$  (P=Q) and
  True_or_False: (P=True)  $\vee$  (P=False)
```

3. Finally, a type `ind` is postulated to have an interpretation by an infinite carrier set. Instead of the more common form to state the axiom of infinity: $\exists f::\text{ind} \Rightarrow \text{ind}$. `injective(f) \wedge \neg surjective(f)`, this axiom comes in two parts over two constants `Zero_Rep` and `Suc_Rep`:

```
axiomatization Zero_Rep :: ind and Suc_Rep :: ind  $\Rightarrow$  ind where
  Suc_Rep_inject: Suc_Rep x = Suc_Rep y  $\Rightarrow$  x = y and
  Suc_Rep_not_Zero_Rep: Suc_Rep x  $\not\leq$  Zero_Rep
```

On this basis, the type of natural numbers is constructed via an inductive definition, the integer and rational numbers via quotient constructions, etc.

4. A further axiom is devoted for another form of the Hilbert-Choice operator:

```
axiomatization Eps :: ('a  $\Rightarrow$  bool)  $\Rightarrow$  'a where
  someI: P x  $\Rightarrow$  P (Eps P)
```

An Isabelle/HOL version coming from a trusted distribution site should *only* have these axioms. Note that in the "src/HOL" folder containing the system libraries, there are many example theories and sub-sessions that actually state their own axioms; a prudent evaluator should make sure that none of these sessions were included.

Conservative Extensions.

Besides the logic, the instance of Isabelle called Isabelle/HOL offers support for specification constructs mapped to conservative extensions schemes, i. e. a combination of type and constant declarations as well as (internal) axioms of a very particular form. We will briefly describe here *type abbreviations*, *type definitions*, *constant definitions*, *datatype definitions*, *primitive recursive definitions*, *well-founded recursive definitions* as well as Locale constructions. We consider this as the "methodologically safe" core of the Isabelle/HOL system.

Using solely these conservative definition principles, the entire Isabelle/HOL library is built which provides a *logically safe language base* providing a large collection of theories like sets, lists, Cartesian products $\alpha \times \beta$ and disjoint type sums $\alpha + \beta$, multi-sets, orderings, and various arithmetic theories which only contain rules derived from conservative definitions.

Type Abbreviations (Synonyms).

For example, typed sets are built in the Isabelle libraries via type synonyms on top of HOL as functions to bool; consequently, the constant definitions for set comprehension and membership are as follows³:

type synonym	α set	$= \alpha \Rightarrow \text{bool}$	
definition	Collect	$:: (\alpha \Rightarrow \text{bool}) \Rightarrow \alpha \text{ set}$	— set comprehension
where	Collect S	$= S$	
definition	member	$:: \alpha \Rightarrow \alpha \text{ set} \Rightarrow \text{bool}$	— membership test
where	member $s S$	$= S s$	

Isabelle's powerful syntax engine is instructed to accept the notation $\{x \mid P\}$ for `Collect $\lambda x. P$` and the notation $s \in S$ for `member $s S$` . As can be inferred from the example, constant definitions are axioms that introduce a fresh constant symbol by some closed, non-recursive expressions; these types of axioms are logically safe since they work like an abbreviation. The syntactic side-conditions of the axioms are mechanically checked, of course. It is straightforward to express the usual operations on sets like \cup , \cap : $\alpha \text{ set} \Rightarrow \alpha \text{ set} \Rightarrow \alpha \text{ set}$ as definitions, too, while the rules of typed set-theory are derived by proofs from them.

³To increase readability, the presentation is slightly simplified.

Datatypes.

Similarly, a logical compiler is invoked for the following statements introducing the types option and list:

$$\begin{aligned} \text{datatype } \alpha \text{ option} &= \text{None} \mid \text{Some } \alpha \\ \text{datatype } \alpha \text{ list} &= \text{Nil} \mid \text{Cons } \alpha (\alpha \text{ list}) \end{aligned} \quad (1.4)$$

Here, $[]$ and $a\#l$ are alternative syntax for Nil and Cons $a l$; moreover, $[a, b, c]$ is defined as alternative syntax for $a\#b\#c\#[]$. Similarly, the option type shown above is given a different notation: $\alpha \text{ option}$ is written as α_{\perp} , None as \perp , and Some X as $\lfloor X \rfloor$. Internally, recursive datatype definitions are represented by type- and constant definitions. Besides the *constructors* None, Some, Nil and Cons, the statement above defines implicitly the match-operation $\text{case } x \text{ of } \perp \Rightarrow F \mid \lfloor a \rfloor \Rightarrow G$ respectively $\text{case } x \text{ of } [] \Rightarrow F \mid (a\#r) \Rightarrow G$ a r . From the internal definitions (not shown here) many properties are automatically derived like distinctness $[] \neq a\#t$, injectivity of the constructors or induction schemes.

Well-founded Recursive Function Definitions.

Finally, there is a parser for primitive and well-founded recursive function definition syntax. For example, the sort-operation can be defined by:

$$\begin{aligned} \text{fun } \text{ins} &:: [\alpha :: \text{linorder}, \alpha \text{ list}] \Rightarrow \alpha \text{ list} \\ \text{where } \text{ins } x [] &= [x] \\ \text{ins } x (y\#ys) &= \text{if } x < y \text{ then } x\#y\#ys \text{ else } y\#(\text{ins } x \text{ } ys) \end{aligned} \quad (1.5)$$

$$\begin{aligned} \text{fun } \text{sort} &:: (\alpha :: \text{linorder}) \text{ list} \Rightarrow \alpha \text{ list} \\ \text{where } \text{sort } [] &= [] \\ \text{sort } (x\#xs) &= \text{ins } x (\text{sort } xs) \end{aligned} \quad (1.6)$$

which is again compiled internally to constant and type definitions. Here, $\alpha :: \text{linorder}$ requires that the type α is a member of the *type class* linorder. Thus, the operation sort works on arbitrary lists of type $(\alpha :: \text{linorder}) \text{ list}$ on which a linear ordering is defined. The internal (non-recursive) constant definition for the operations ins and sort is quite involved and requires a termination proof with respect to a well-founded ordering constructed by a heuristic. Nevertheless, the logical compiler will finally derive all the equations in the statements above from this definition and makes them available for automated simplification.

The theory of partial functions is of particular practical importance. Partial functions $\alpha \rightarrow \beta$ are then defined as functions $\alpha \Rightarrow \beta \text{ option}$ supporting the usual concepts of domain $\text{dom } f \equiv \{x \mid f \ x \neq \text{None}\}$ and range $\text{ran } f \equiv \{x \mid \exists y. f \ y = \text{Some } x\}$. Partial functions can be viewed as “maps” or dictionaries; the *empty* map is defined by $\emptyset \equiv \lambda x. \text{None}$, and the update operation, written $p(x \mapsto t)$, by $\lambda y. \text{if } y = x \text{ then } \text{Some } t \text{ else } p \ y$. Finally, the override operation on maps, written $p_1 \oplus p_2$, is defined by $\lambda x. \text{case } p_1 \ x \text{ of } \text{None} \Rightarrow p_2 \ x \mid \text{Some } X \Rightarrow \text{Some } X$.

Type definitions.

Type definitions allows for a safe introduction of a new type. Other specification constructs, for example datatype, are based on it. The underlying construction is simple: any non-empty subset of an existing type can be turned into new type. This is achieved by defining an isomorphism

between this set and the new type; the latter is introduced by two fresh constant symbols (representing the abstraction and the concretization function) and three internally generated axioms. As a simple example, consider the definition of type containing three elements. This type is represented by the first three natural numbers:

```
typedef three = {0::nat, 1, 2}
  apply (rule_tac x = 0 in exI)
  apply blast
done
```

(1.7)

In order to enforce that the representing set on the right hand side is non empty, the package requires for this new type a proof of non-emptiness:

```
typedef three = {0::nat, 1, 2}
1.       $\exists x. x \in \{0, 1, 2\}$ 
```

(1.8)

To use this new type we need to finish the proof of non empty set started by the use of **typedef** which can be done differently. For example we can finish the proof using existing theorems on the logical operator \exists in Isabelle/HOL. To see all Isabelle's theorems related to \exists we use the Isabelle command `find_theorems`. The query searches for theorems whose name contains an "ex" substring. One of the results is:

```
find_theorems name: ex
HOL.exI:       $?P \ ?x \implies \exists x. ?P \ x$ 
```

(1.9)

The searched theorems is applied in the following. In our case, the Isabelle proof method `rule_tac` is used, a resolution step, which unifies the theorem `HOL.exI` against the first proof goal in a resolution step:

```
apply (rule_tac x = 0 in exI)
apply blast
done
```

Its application in the proof allows to replace the schematic variable $?x$ by the constant 0 in our proof; this is specified by the key word **in** followed by the name of the theorem. The other the schematic variable $?P$ is automatically filled in (using higher-order unification), which is possible since only one solution remains. The remainder of the proof consists of a call to the highly automated method `blast`, which does the trick for the necessary set-theoretic proof.

It remains to point out that the same proof can be done by different proof-style called *structured proof* or *Isar-proof*. The same proof can be represented in this style as follows:

```
typedef three = {0::nat, 1, 2}
proof
show      1  $\in \{0::nat, 1, 2\}$ 
by      blast
qed
```

(1.10)

After finishing the proof about the definition of this new type, many theorems will be deduced automatically by Isabelle. We can check the new deduced theorems related to this new type

by using the command **find_theorems**. In the concrete example, there are 82 new theorems deduced that were related to this type definition.

```

find_theorems name: three
searched for:
name: "three"
found 82 theorem(s) (40 displayed)

```

(1.11)

Inductively defined predicates.

This section is dedicated to the most important definition principle after recursive functions and datatypes: inductively defined predicates. We will introduce a simple example: the set of even (natural) numbers. For more complicated examples see [TN]. The specification construct allows for building the *least* set which is closed under a given collection of introduction rules; in our case: one rule that states 0 is an even number, and the other one rule that states that if we add 2 to every even number we will get an even number. Using the keyword **inductive**, we declare the constant *even* to be a predicate that allows us to get the set of natural numbers with desired properties. (Note that sets and boolean functions are treated the same.)

```

inductive even :: nat ⇒ bool
where
  zero[intro!]: even 0
  step[intro!]: even n ⇒ even(Suc(Suc n))

```

(1.12)

Note that the declaration of the rules comes, as usual in many places in the Isar-language, with an instrumentation: for both rules, the names *zero* and *step* were introduced, and with a number of *attributes* it can be stated *how* the given rule or *thm* should be *used* in proofs: the keyword `[intro!]` indicates that they should be used as introduction rules in proof search. After the inductive statement, Isabelle generates a fixed point definition for *even* and proves theorems about it. These theorems include the introduction rules specified in the declaration, an elimination rule for case analysis and an induction rule for the global judgment. To inspect these theorems we can again use **find_theorems** which results in:

```

find_theorems name: even
three.even.cases:
even ?a ⇒ (?a = 0 ⇒ ?P) ⇒
( ∧ n. ?a = Suc (Suc n) ⇒ even n ⇒ ?P ) ⇒ ?P
three.even.induct:
even ?x ⇒ ?P 0 ⇒
( ∧ n. even n ⇒ ?P n ⇒ ?P (Suc (Suc n)) ) ⇒ ?P ?x
three.even.zero: even 0
three.even.step: even ?n ⇒ even (Suc (Suc ?n))

```

(1.13)

We can refer to these theorems by automatically-generated names, for example: *three.even.cases*, *three.even.induct* ...

Type classes.

We will introduce another important concept. Type-classes can be seen as a simple modularization concept (similar locales, but with less expressive power), which is particularly well

integrated into the type system. Similar to Haskell[WB89, Wen97], type classes κ restrict type variables to belong to a particular class of types having common properties. The introduction of a new class `plus` and its operation \oplus is done by this Isabelle/Isar fragment:

```
class plus =
fixes plus ::  $\alpha \Rightarrow \alpha \Rightarrow \alpha$  (infixl  $\oplus$  70)
```

(1.14)

```
instantiation nat :: plus
begin
primrec plus_nat :: nat  $\Rightarrow$  nat  $\Rightarrow$  nat
where
  (0 :: nat)  $\oplus$  n = n
  Suc m  $\oplus$  n = Suc (m  $\oplus$  n)
instance proof qed
end
```

(1.15)

The type of the operation \oplus carries a class constraint $'\alpha :: \text{plus}$ on its type variable, meaning that only types of class `plus` can be instantiated for $'\alpha$. To locally instantiate a type-class by an other existing type we use the command **instantiation**. For example to instantiate `plus` on `nat` we write the key word **instantiation** and the name of existing type that we want to instantiate, and the type of the operation which is in our case `nat \Rightarrow nat \Rightarrow nat`. Now we define the \oplus function and give a semantic to it on type `nat`. Note that all function names are written by the combination of the name of the class operation and the name of the instance which the class operation will be applied on (example `plus_nat`). In case of uncertainty, these names may be inspected using the command **print_context** as follow:

```
instantiation nat :: plus
begin
print_context
...
Isabelle output will be
nat :: three.plus
plus_nat  $\equiv$  three.plus.classe.plus :: nat  $\Rightarrow$  nat  $\Rightarrow$  nat
```

(1.16)

In general, assumptions were assumed in the context of the instantiations, proofs for those assumptions are mandatory in instantiations. Such proofs are done using the command **instantiation** before the end of the context of the instantiations. In our example, the proof is a standard phrase necessary for technical reasons. We can also add many instantiations for the operations. For example the operation \oplus of the class `plus` can be applied on the type of products:

```
instantiation prod :: (plus, plus)plus
begin
fun plus_prod ::  $\alpha^*\beta \Rightarrow \alpha^*\beta \Rightarrow \alpha^*\beta$ 
where
  (x, y)  $\oplus$  (w, z) = (x  $\oplus$  y, w  $\oplus$  z)
instance proof qed
end
```

(1.17)

Now, in a term $(3, 4, 5) \oplus (1, 2, 3)$, the type inference will infer that there is actually a series of instantiations that define this product on triples ... More in depth explanations for type classes are in [TN].

While type-classes have a strictly weaker expressive power than Isabelle's Locales to be discussed in the sequel, they have the advantage that the types can be inferred completely automatic; their annotation can therefore be omitted in most cases. Furthermore, the lack of dependent types (a concept existing in Coq) can in some practical cases be compensated by type-classes; it is, for example, perfectly possible to define the bit vector type "32 word" and "64 word" inside a word-library providing types " α word" (here, "32" is a syntactic synonym for a type-class of types that are representable by 32 bits). Thus, type-classes can establish dependencies of types from values which is impossible in a standard Hindley-Milner type-system.

Locales.

Locales are Isabelle's approach for dealing with parametric theories [Bal10]. They have been designed as a module system that can adequately represent the complex inter-dependencies between structures found in abstract algebra, but have proven fruitful also in other applications. We will briefly discuss major features of locales.

As a prerequisite, recall that the general format of a *logical rule* represented in Isabelle/Pure is:

$$\bigwedge x_1 \dots x_n . A_1 \implies \dots \implies A_m \implies C$$

On the level of the Isar-language, a rule of this form can equivalently be represented as:

$$\begin{array}{l} \textbf{fixes } x_1 \dots x_n \\ \textbf{assumes } A_1 \\ \textbf{and } \dots \\ \textbf{and } A_m \\ \textbf{show } C \end{array} \quad (1.18)$$

Parameters and assumptions together form a *local context*. A formula C is a *theorem in a local context* if it is a *conclusion*. A *locale* is just local context that has been made persistent. As a particular feature, they allow for introducing local syntax for the x_i and individual prover instrumentation for the assumptions. To the user, however, they provide powerful means for declaration, combination, and for reuse of theorems proved in them. The following example is the formalization of partial order with locale `partial_order`.

$$\begin{array}{l} \textbf{locale } \text{partial_order} = \\ \textbf{fixes } \text{less_equal} :: \text{nat} \Rightarrow \text{nat} \Rightarrow \text{bool} \text{ (infixl } \lesssim 50) \\ \textbf{assumes } \text{refl [intro,simp]: } x \lesssim x \\ \textbf{and } \text{anti_sym[intro]: } x \lesssim y \implies y \lesssim x \implies x = y \\ \textbf{and } \text{trans[trans]: } x \lesssim y \implies y \lesssim z \implies x \lesssim z \end{array} \quad (1.19)$$

In this **locale** the parameter is `less_equal`, which is a binary predicate with infix syntax \lesssim . The parameter syntax is available in the subsequent assumptions, which correspond to the familiar partial order "axioms". Isabelle recognizes unbound names as free variables. In locale assumptions, they are implicitly universally quantified. That is, $x \lesssim y \implies y \lesssim z \implies x \lesssim z$ in fact means $\bigwedge x y z. x \lesssim y \implies y \lesssim z \implies x \lesssim z$. There are two Isar commands to inspect

a locale: **print_locale** lists the names of all locales of current theory; **print_locale** α prints the parameters and assumptions of locale α ; the variation **print_locale!** α additionally outputs the conclusions that are stored in the locale. For the Isar command:

$$\mathbf{print_locale} \text{ partial_order} \quad (1.20)$$

the system produces the output:

$$\begin{aligned} &\mathbf{locale} \text{ partial_order} \\ &\mathbf{fixes} \text{ less_equal} :: \text{nat} \Rightarrow \text{nat} \Rightarrow \text{bool} \\ &\mathbf{assumes} \text{ partial_order} \text{ op } \lesssim \end{aligned} \quad (1.21)$$

Analogously, for:

$$\mathbf{print_locale!} \text{ partial_order} \quad (1.22)$$

the output

$$\begin{aligned} &\mathbf{locale} \text{ partial_order} \\ &\mathbf{fixes} \text{ less_equal} :: \text{nat} \Rightarrow \text{nat} \Rightarrow \text{bool} \\ &\mathbf{assumes} \text{ partial_order} \text{ op } \lesssim \\ &\mathbf{notes} \text{ partial_order_axioms} = \\ &\quad (\text{partial_order} \text{ op } \lesssim) [\text{attribute} < \text{attribute} >] \\ &\mathbf{notes} \text{ refl} = (?x \lesssim ?x) [\text{HOL.intro, simp}] \\ &\mathbf{and} \text{ anti_sym} = (?x \lesssim ?y \Longrightarrow ?y \lesssim ?x \Longrightarrow ?x = ?y) [\text{HOL.intro}] \\ &\mathbf{and} \text{ trans} = (?x \lesssim ?y \Longrightarrow ?y \lesssim ?z \Longrightarrow ?x \lesssim ?z) [\text{trans}] \end{aligned} \quad (1.23)$$

is produced. Here, the keyword **notes** denotes a conclusion element. There is two conclusions, which were added automatically. Instead there is only one assumption, namely $\text{partial_order}(\text{op } \lesssim)$. The locale declaration has introduced the predicate partial_order to the theory. This predicate is called the *locale predicate*. Its definition may be inspected by the command:

$$\text{partial_order_def} \quad (1.24)$$

corresponding to the output:

$$\begin{aligned} &\text{partial_order } ?\text{less_equal} \equiv \\ &\quad (\forall x. ?\text{less_equal } x \ x) \wedge \\ &\quad (\forall x \ y. ?\text{less_equal } x \ y \longrightarrow ?\text{less_equal } y \ x \longrightarrow x = y) \\ &\quad (\forall x \ y \ z. ?\text{less_equal } x \ y \longrightarrow ?\text{less_equal } y \ z \longrightarrow ?\text{less_equal } x \ z) \end{aligned} \quad (1.25)$$

Each conclusion has *foundational theorem* as counterpart in the theory. Technically, this is simply the theorem composed of local context and conclusion. For the transitivity, for example, we have the output:

$$\begin{aligned} &\text{partial_order } ?\text{less_equal} \Longrightarrow ?\text{less_equal } ?x \ ?y \Longrightarrow ?\text{less_equal } ?y \ ?z \\ &\Longrightarrow ?\text{less_equal } ?x \ ?z \end{aligned} \quad (1.26)$$

The specification of a locale is fixed, but its list of conclusions may be extended through Isabelle commands that take a *target* argument. In the following, two examples on two Isabelle

commands that accept a target. The first example on the command **definition** and the second example is on the command **lemma**.

```
definition (in partial_order)
  strict_less :: nat  $\Rightarrow$  nat  $\Rightarrow$  bool
  where x  $\lesssim$  y = x  $\lesssim$  y  $\wedge$  x  $\neq$  y
```

(1.27)

The strict order `strict_less` with infix syntax \lesssim is defined in terms of the locale parameter `less_equal` and the general equality of the object logic we work in. The definition generates a constant `partial_order.strict_less` with definition `partial_order.less_def` :

```
partial_order?less_equal  $\implies$ 
partial_order.less ?less_equal ?x ?y =
(?less_equal ?x ?y  $\wedge$  ?x  $\neq$  ?y)
```

(1.28)

The context of a locale can be extended by a block of commands, delimited by **begin** and **end**, like when we start a new theory. The main restriction when we use a block of commands, is that the block refer to the same target (the same locale). If the block of commands follows a locale declaration, that makes this locale the target. In other cases, the target for a block may be given with the **context** command. In the example below, we will introduce two new definitions for the locale `partial_order`, in those new definitions we will introduce the notion of infimum and supremum for partial orders.

```
context partial_order
begin
definition is_inf
  where is_inf x y i =
    (i  $\lesssim$  x  $\wedge$  i  $\lesssim$  y  $\wedge$ 
     ( $\forall z. z \lesssim x \wedge z \lesssim y \longrightarrow z \lesssim i$ ))
definition is_sup
  where is_sup x y s =
    (x  $\lesssim$  s  $\wedge$  y  $\lesssim$  s  $\wedge$ 
     ( $\forall z. x \lesssim z \wedge y \lesssim z \longrightarrow s \lesssim z$ ))
end
```

(1.29)

ML Code.

It is possible inside Isabelle documents to directly access the underlying ML-layer of the system architecture, and even extend the environment of the underlying ML interpreter/compiler. One can include the fragment:

```
ML{* fun fac x = if x = 0 then 1 else x * fac(x-1); *
```

in a document and then later on evaluate:

```
ML{* fac 20; *}

```

Since Isabelle itself sits as a collection of ML modules in this SML environment, it is possible to access its kernel and tactical functions:

```
ML{* open Tactic;
  fun mis x = res_inst_tac [(x, x)] {@thm exI} 1*}
```

which defines a new tactic that applies just the existential-introduction rule of HOL. This is the key to build large and own tactic procedures and even tools inside the Isabelle environment. Note that the fragment `{@thm exI}` is called an *antiquotation*; it is expanded before being passed to the SML compiler with code that accesses the `thm exI` (see section Sec. 1.3.3, pp8.) in the Isabelle database for theorems. By additional SML-code, this tactic can be converted into a *Isar-method*, which can be bound to own syntax inside the Isar-language. Thus, the proof language is technically extensible by own, user-defined proof-commands (see [Wen] for the details).

1.3.4 Isabelle Proofs

In addition to types, classes and constants definitions, Isabelle theories can be extended by proving new lemmas and theorems. These lemmas and theorems are derived from other existing theorems in the context of the current theory. Isabelle offers various ways to construct proofs for new theorems, we distinguish two main categories: forward and backward proofs. In addition to Isabelle proofs, some external proofs can be integrated – in a logically safe way – and compiled into an Isabelle proof.

Local forward proofs.

The goal of a forward proof is to derive a new theorem from old ones. This is done either by instantiating some unknowns in the old theorems, or by composing different theorems together.

The instantiation can be done using the **of** and **where** operators as follows: `thm[of inst1 inst2 ...]` or `thm[where var1=inst1 and var2=inst2 ...]`. If we consider for example the existential introduction theorem called `exI` and given by $?P \ ?x \implies \exists x. \ ?P \ x$. The unknown variable `x` can be instantiated with a fixed variable `a` using the following command `exI[of _ a]` which is equivalent to `exI[where x=a]`. Note that when using **of** the instances of the variables appear in the same order of appearance of the unknown variables in the theorems. Consequently, we can avoid instantiating a variable by giving a dummy value in the position of its corresponding instance.

The second way of deriving theorems is by composing different theorems together using the **OF** or **THEN** operators. The first operator **OF** is used to compose one theorem to others. For a theorem `th1` given by $A \implies B$ and a theorem `th2` given by A' , the theorem `th1[OF th2]` results from the unification of A and A' and thus instantiating the unknowns in B . Theorems with multiple premises can be composed to more than one theorem given as arguments to the **OF** operator. For example, given the conjunction introduction theorem `conjI` given by $?P \implies ?Q \implies ?P \wedge ?Q$ and the reflexivity theorem `ref` given by $?x = ?x$, the composition of these theorem `conjI[OF refl[of a] refl[of b]]` results in the following theorem $a = a \wedge b = b$. In a similar way, the **THEN** operator is used to compose different theorems together. The theorem `th1[THEN th2]` is obtained by applying the rule `th2` to the theorem `th1`. For example, composing a theorem `th1` given by $a = b$ with the symmetry rule `sym` given by $?s = ?t \implies ?t = ?s$ is written `th1[THENSym]` and the result is $b = a$.

Global backward proofs.

The usual and mostly used proof style is the backward or goal-directed proof style. First, a proof goal is introduced then the proof is performed by simplifying this goal into different subgoals and, finally, prove the resulting subgoals from existing theorems. The proofs are build

using natural deduction by applying some existing (proved) inference rules. For each logical operator, two kinds of rules are defined: introduction and elimination rules.

The backward proofs can be structured in two different ways:

1. Apply style proofs, where the proof goal is simplified using a succession of rules applications. This results in a so-called apply-script, describing the proof steps. An example of such a proof is given in the following:

```

lemma conj_rule:  $\llbracket P; Q \rrbracket \implies P \wedge (Q \wedge P)$ 
  apply (rule conjI)
  apply assumption
  apply (rule conjI)
  apply assumption
  apply assumption
done
    
```

(1.30)

Although this proof style is easy to apply, long apply-scripts can become unreadable and hard to maintain. A more structured and safe way to write the proofs is by using the Isar language.

2. Structured Isar proofs allow for writing sophisticated and yet still fairly human-readable proofs. The Isar language defines a set of commands and shortcuts that offer more control on the proof state. An example of a structured induction proof is given in the following:

```

lemma
  fixes n :: nat
  show  $2 * (\sum i = 0..n. i) = n * (n + 1)$ 
proof (induct n)
  case 0
  have  $2 * (\sum i = 0..n. i) = (0 :: nat)$ 
  by simp
  also have  $(0 :: nat) = 0 * (0 + 1)$ 
  by simp
  finally show ?case .
next
  case (Suc n)
  have  $2 * (\sum i = 0..Suc\ n. i) = 2 * (\sum i = 0..n. i) + 2 * (n + 1)$ 
  by simp
  also have  $2 * (\sum i = 0..n. i) = n * (n + 1)$ 
  by (rule Suc. hyps)
  also have  $n * (n + 1) + 2 * (n + 1) = Suc\ n * (Suc\ n + 1)$ 
  by simp
  finally show ?case .
qed
    
```

(1.31)

For the sake of this presentation, we appeal to an “immediate intuition” of a mathematically knowledgeable reader; for detailed introduction into the structured proof language, the reader is referred to the Isar Reference Manual of the System documentation.

Locales can be directly referred to in proofs. For example, one could in a constructivist version of HOL (see `src/HOL/ex/Higher_Order_Logic.thy`) state and prove:

```

locale          classical =
assumes         classical :  $(A \implies A) \implies A$ 

theorem         (in classical)
Peirce's_Law :    $((A \rightarrow B) \rightarrow A) \rightarrow A$ 
proof
...
qed

```

(1.32)

Thus, the effect of the “(in classical)” clause in the example above is to add additional assumptions into the local context. A skeptical evaluator might therefore insist on proofs of the existence of witnesses for the locale, i. e. a proof for $\exists x. \text{partial_order } x$. Since in a classical setting the existence of a function can be stated via the Hilbert-operator, that decides for a Turing machine that it terminates for a given input, a very skeptical evaluator might even insist on a constructive witness for these existence proofs.

1.3.5 Isabelle/HOL System Features

Finally, Isabelle/HOL manages a set of *executable types and operators*, i. e., types and operators for which a compilation to SML, OCaml, Scala, or Haskell is possible. Setups for arithmetic types such as `int` have been done allowing for different trade-offs between trust and efficiency. Moreover any datatype and any recursive function are included in this executable set (providing that they only consist of executable operators). Of particular interest for evaluators is the use of the `Isar` command:

```
valid sort[1, 7, 3]
```

(1.33)

In the context of the definitions Fact 1.5, it will compile them via the code-generator to SML code, execute it, and output:

```
[1, 3, 7]
```

(1.34)

This provides an easy means to inspect constructive definitions and to get easy feedback for given test examples for them. See the part “Code generation from Isabelle/HOL theories” by Florian Haftmann from the Isabelle system documentation for further details.

Of particular interest for evaluators or certifications are Isabelle’s features for semantically supported typesetting: within the document element:

```
text{* This is text containing  $\lambda$ 's and  $\beta$ 's ... *}
```

for example, arbitrary LaTeX code can be inserted for using technical and mathematical notation of annotations of formal document elements. Inside a text-document, the *document antiquotation* mechanism already mentioned in Sec. 1.3.3 can be applied:

```
text{* Text containing theorems like  $\{ \text{thm exI} \}$  ... *}
```

which results in a print of theorems directly from their formal Isabelle presentation. It is possible to define new antiquotations, for example to track security requirements or security claims in theorems or tests. A detailed description of document antiquotations is found in the “Isar

Reference Manual” by Makarius Wenzel from the Isabelle system documentation. It is even possible to define own antiquotations in Isabelle for categories of the common criteria like protection profiles, security targets, requirements, security properties etc. For all these entities, be it informal or formal, declarations and applications of antiquotations can be used in text fragments that allow for a direct consistency checking over the entire document. Since a concrete setup for such mechanism offers a number of deviation points, we refrain in this document on *mandatory* recommendations and refer to a future document on *styleguide recommendations*.

Evaluators are encouraged to use the Isabelle/jedit user-interface directly (and not just the generated .pdf documentation), since it allows for an in-depth inspection and exploration of the formal content of a theory: tooltips reveal typing information, evaluations of critical expressions can often be done by the `value . . .` document item, and operator-symbols occurring in HOL-expressions were hyper-linked to referring definitions or binding occurrences. Note, however, that a user-interface is a dozen system layers away from a Isabelle inference kernel which opens the way for implementation errors in display and editing components, increasing the risk of misinterpretations. A final check of an entire document should therefore be made in the (GUI-less) build mode (which enforces also stronger checking).

1.4 Methodological Recommendations for the Evaluator

As said earlier, there are four potential dangers of a formal proof system that it wrongly accepts the desired theorem “This operating system is secure”:

1. Inherent inconsistency of the logics (e. g., HOL) or inconsistent use of the logics (introduction of inconsistent axioms by one way or the other).
2. The incorrect implementation of Isabelle the Isabelle Kernel and of the HOL instance in it.
3. The incorrect package implementation realizing advanced specification constructions like type definitions etc.
4. Since Isabelle is highly configurable, there is a certain danger of obfuscation of bogus-proofs.

Beyond the more philosophical objections⁴, the risk outlined by the by first item is in fact **minimal**: Higher-order logic (HOL) in itself is an extremely well studied object of academic interest ([And86, GM93]; compare to the discussion in Sec. 1.3.3), and while there are known limits in proving soundness and completeness inside a HOL-prover, they just stimulated a lot of recent research to come a “formal proof over HOL in HOL” as close as possible, e.g. by adding to HOL an axiom over the existence of a sufficiently large cardinal [Har06, MOK13].

The risk outlined by the second item is also **very small**. The reasons are threefold:

- A Some of the aforementioned soundness proofs cover also the implementation aspects of the core of a provers of the HOL-family (HOL-light, ...).
- B The specific architecture of provers of the LCF family (HOL4, Isabelle, HOL-light, Coq) enforces that any proof is actually checked by by this fairly small core.

⁴For example, the fundamental doubt in the existence of infinite sets[And86]...

- C These core-inferences can optionally be protocolled in an proof-object which can, in principle, in case of serious doubt be checked by another implementation of a HOL-prover. However, since these objects tend to be very large, this approach requires decent engineering. Fortunately, this should only be necessary in exceptional cases.

The risk of the third item is **minimal** as far as the described standard conservative standard extension schemes such as `type_synonym`'s, `datatype`'s, `definition`'s and `fun`'s, `typedef`'s, `specification`'s, `inductive`'s, `type-classes` and `locales` are concerned. The same holds for diagnostic commands like `type`, `term`, `valid`, etc. that do not change the global context of a theory. These are fairly well-understood schemes which have in parts been proven formally correct for similar systems such as the HOL4 system[KAMO14]. These schemes cover the largest parts of the Isabelle/HOL libraries. Here lies the main advantage of the LCF-approach and the methodology to base libraries on conservative (logically safe) definitions.

The risk is **small** as far as other standard extension schemes are concerned; since extension schemes generate internally axioms, there have been reported consistency problems with combinations of other extension schemes such as `consts` and `defs` as well as `defs (overloaded)`; the Isabelle reference manual points out that the internal checks of Isabelle do not guarantee soundness.⁵

It remains the risk of item four, which is concerned with the resulting methodology in “how to use Isabelle”. For very large theory documentations, it must be considered **non-negligible**. It is the key-issue addressed in the remainder of this section.

1.4.1 On the Use of SML

As mentioned earlier, Isabelle is an open environment that allows via

```
ML{* SML ML code *}
```

to include arbitrary SML programs, in particular programs that make direct inferences on top of the kernel. Per se, this use of Isabelle is not unsafe; critical parts of the HOL library use this mechanism. Isabelle is designed to have user land SML code extensions, and the kernel protects itself against logical inconsistencies coming from ML extensions. However, there are a few deliberate opt-outs, and furthermore, it is in principle possible to obfuscate them in Isabelle ML code such that an evaluator may be fooled by a text appearing to be an Isabelle proof but isn't in the sense of the inference kernel. Thus, besides the principle possibility that a pretty-printed theorem does not state what it appears to state by some misuse of mathematical notation (an inherent problem of any formal method), there is the possibility of fake-proofs as a consequence of ML code and (re)-configurations of the ISAR proof language.

If SML-code is accepted in an evaluation, it has to be made sure — potentially by extra justifications or external experts with Isabelle implementation expertise — that this code does not implicitly generate axioms, registers oracles and defines proof methods equivalent to **sorry** (or variants like `sorry_fun`) to be discussed in the sequel; in any case, the evaluation is substantially simpler if SML-code is strictly avoided.

⁵See Isabelle Isar-Reference Manual (Version 2013-2, pp. 103): “It is at the discretion of the user to avoid malformed theory specifications!”

1.4.2 Axioms and Bogus-Proofs

Obviously, when using the Isar `axiomatization` construct allowing to add an arbitrary axiom, it is immediately possible to bring the system in an inconsistent state. The immediate methodological consequence is to ban it from use in to be evaluated theories completely (such that it is only internally used inside specification constructs in and in the aforementioned foundational axioms coming with the system distribution) and to restrict theory building on conservative extensions. This is also common practice in scientific conferences addressing formal proof such as ITP.

However, there are more subtle ways to introduce an axiom that leads to inconsistency. First, there is a mechanism in Isabelle to register *oracles* into the system [TN, Wen]. They can be used for a particularly simple, but logically unsafe integration of external provers into Isabelle and can be used inside self-defined tactics. Logically, an oracle is a function that produces axioms on the fly. It is an instance of the axiom rule of the kernel, but there is an operational difference: The system always records oracle invocations within proof-objects of theorems by a unique tag. Of course, oracle invocations should again be avoided in a certified proof.

A particular instance of the oracle mechanism is the **sorry** proof method. This is method is always applicable and closes any (sub)-proof successfully, and a useful means in top-down proof developments in Isabelle. Unnecessary to repeat that no **sorry** statements should remain in a proof document underlying certification. By the way, the system is by default in a mode in which it refuses to generate proof documents containing **sorry**'s, only by explicitly putting it in a mode called `quick_and_dirty` this can be overcome. There are several ways to activate `quick_and_dirty`, by it by explicit ML statements like `quick_and_dirty := true`, be it in the `ROOT.ML`-files (till version 2013-1), or be it in the session-configuration files `ROOT-files` (since version 2013).

Oracles and **sorry**'s are particularly dangerous in methodological foundation proofs (type or type-class is non-empty, recursions well-founded), since the use of the the oracle-tag inside the corresponding proof-objects gets lost on the level of type expressions. Thus, a **sorry** could introduce inconsistent types whose “effects” could be used in bogus-proofs depending on them.

We will discuss this a little more in detail: Recall that deduction in Isabelle/HOL is centered around the requirement that types and type-classes are non-empty. This is a consequence of the fact that the β -reduction rule $((\lambda x :: \tau. E) E' \rightarrow E[x := E'])$ is executed pervasively during deduction, be in in resolution or rewriting steps. It is well-known however, that β -reduction is unsound in the presence of empty types⁶. Thus, an obfuscated **sorry** in a methodological proof leaves no other than very local traces in the proof objects and can be exploited much later via an inconsistent type in a proof based on this type definition; the exploit could again be obfuscated by another self-defined proof-method, say `auto'` which will be hard to detect by inspection. The only systematic way to rule out obfuscated bogus-proof is either by ruling out ML-constructs or by checking *all* proof objects of the entire theory.

1.4.3 On the Use of External Provers

The Isabelle distribution comes with a number of external provers, namely:

- **sledgehammer** : its use is uncritical, since it remains completely extern to proof documentations and is only used for the generation of high-level Isabelle proofs, that were

⁶Consider the case of τ having a semantic interpretation into an empty set $I(\tau) = \emptyset$: then the semantic interpretation of the function $(\lambda x :: \tau. E)$ must be in the function space: $\emptyset^D = \emptyset$ where D is the space of interpretations for the type τ' of E . Obviously, there is no possible result for the application ...

certified by the kernel.

- `blast`, `metis`: these are internal devices but also uncritical, since their results were used via a proof object certification.
- `smt`: this method uses, for example, the external SMT-solver Z3. The integration is carefully made and uses no oracles - instead, a form of tactical proof re-construction mechanism is used [BW10] that is logically safe.

Other external provers have to be considered carefully; in particular integrations using the oracle-mechanism should be ruled out.

1.5 Extensions of Isabelle: Guidelines for the Evaluator

As said earlier, the ML code should be considered harmful in theories to be evaluated. There are, however, a number of add-ons on Isabelle, which can be considered as tools in their own right and which heavily use ML code inside.

1.5.1 Example: Isabelle/Simpl

Isabelle/Simpl is an verification environment built conservatively on Isabelle/HOL. It supports a sequential imperative programming language, for which it defines its syntax, semantics, Hoare Logics and a verification condition generator (again derived), which form together a complete verification environment. Together with an (untrusted) parser that compiles C programs into Isabelle/Simpl[GAK12], this particular environment follows a similar program verification technique like Frama-C/Why/AltErgo ([CKK⁺12, FP13], <http://alt-ergo.lri.fr>) or VCC/Boogie/Z3[BMSW10].

The entire environment is part of the Isabelle-oriented “Archive of formal Proofs”, see <http://afp.sourceforge.net/> in general and <http://afp.sourceforge.net/entries/Simpl.shtml> in particular.

The environment has been used for one of the most ambitious code-verification projects recently, the verification of the L4-Microkernel (cf. <http://www.ertos.nicta.com.au/research/l4.verified/>, [KHS09]).

In itself, Isabelle/Simpl can be considered nearly as “trustable” as Isabelle/HOL itself : the library is built upon conservative extensions of the HOL -kernel, and the ML extensions are done by Isabelle developers themselves and stood the test of the time. Program verification proofs establishing that a Simpl-program is correct with respect its (pre-post-condition) specifications can be handled by the same evaluation procedures as any other Isabelle development.

However, as in any process involving the verification of C programs, the C parser and its transition from “real C” to the idealized imperative language Simpl has to be considered with a wise dose of skepticism. Here is are whole spectrum of different glimpses possible: since the C parser defines a semantics-by-translation for its fragment of C, the question remains unproven that this semantics is faithful to the semantics of the real C compiler generating production-level code (which involves questions on compiler correctness, semantic faithfulness of the execution environment, correctness of compilation optimizations, hardware-correctness, etc.). The problem has been addressed via particular validation techniques of the parsing process [GAK12], but is, in full generality, unsolvable.

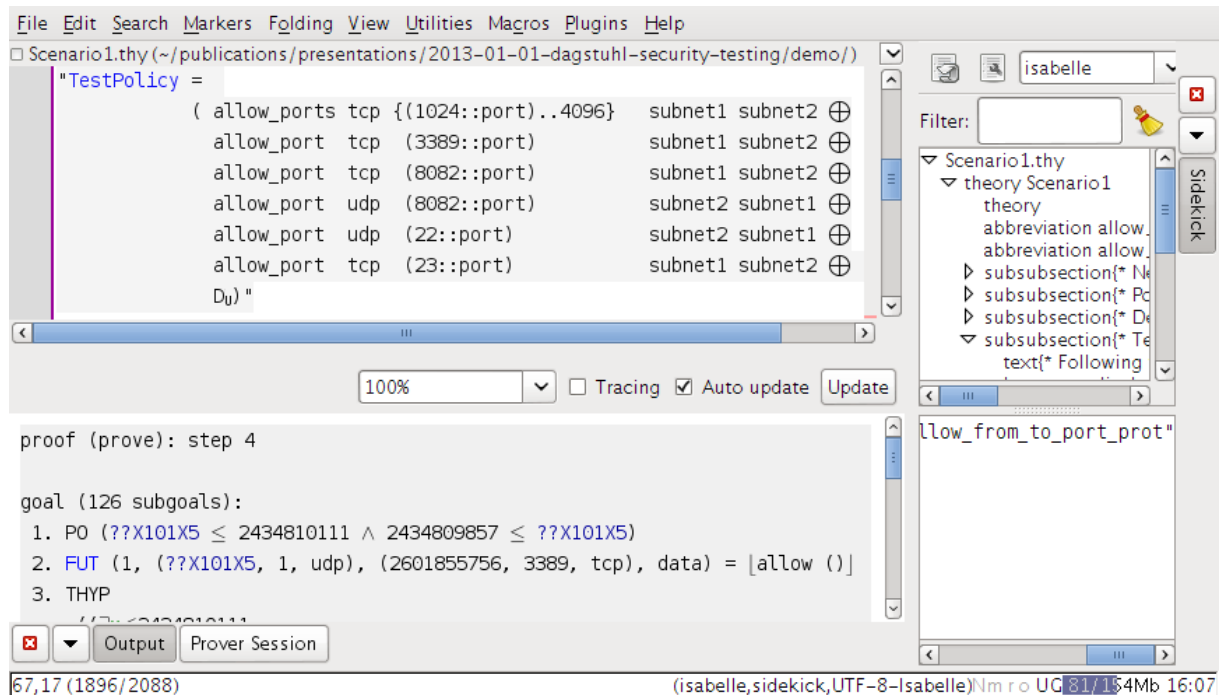


Figure 1.2: An Isabelle session showing the jEdit client as Isabelle Interface. The upper-left sub-window allows one to interactively step through a test theory comprising test specifications while the lower-left sub-window shows the corresponding system state of the spot marked in blue in the upper window.

1.5.2 Example: The HOL-TestGen Test-Generation System

hol!-TESTGEN (see Fig. 1.2) is an interactive, i. e., semi-automated, test generation tool for specification-based tests built upon Isabelle/HOL. Instead of using Isabelle/HOL as “proof assistant,” it is used as modeling environment for the domain specific background theory of a test (the *test theory*), for stating and logically transforming test goals (the *test specifications*), as-well as for the test generation method implemented by Isabelle’s tactic procedures. In a nutshell, the test generation method consists of:

1. a *test case generation* phase, which is essentially a equivalence partitioning procedure of the input/output relation based on a **cnf!**-like normal form computation,
2. a *test data selection* phase, which essentially uses a combination of constraint solvers using random test generation and the integrated SMT-solver Z3 [MB08],
3. a *test execution* phase, which reuses the Isabelle/HOL code-generators to convert the instantiated test cases to test driver code that is run against a system under test.

A detailed account on the symbolic computation performed by the test case generation and test selection procedures is contained in [BW13]. The test case generation method is basically an *equivalence partitioning* combined with a variable splitting technique that can be seen as an (*abstract*) *syntax testing* in the sense of the ISO 29199 specification [Int12, Sec. 5.2.1 and 5.2.4].

The equivalence partitioning separates the input/output relation of a program under test (*PUT*), usually specified by pre- and post-conditions, into classes for which the tester has reasons to believe that *PUT* will treat them the same.

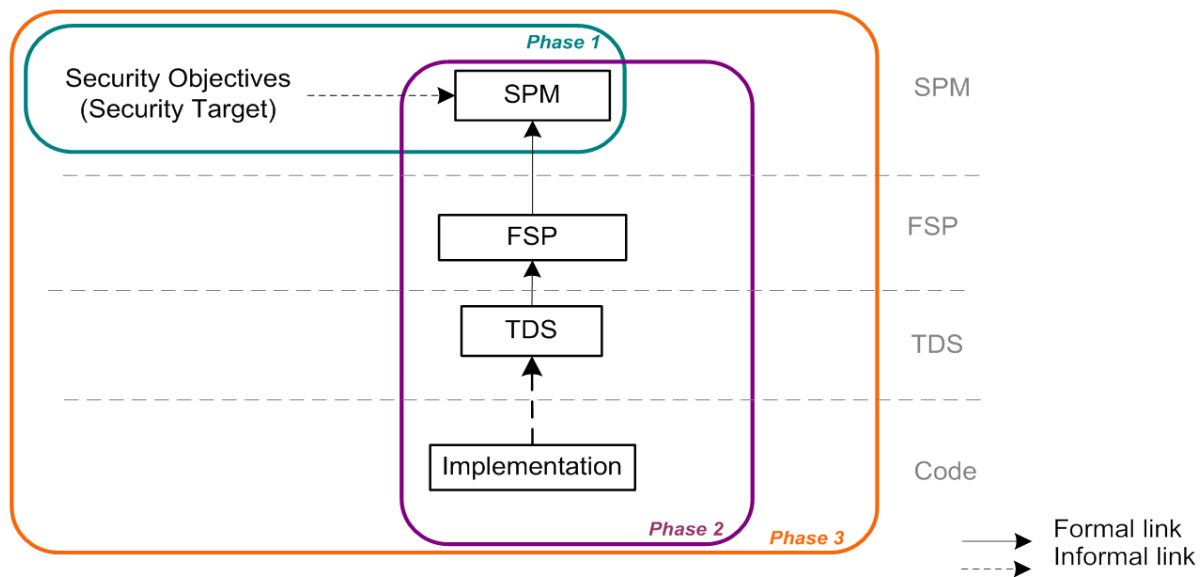


Figure 1.3: Refinement steps for a formal development approach compliant to CC

1.5.3 By the Way: Test vs. Proof

Of course, the **hol!**-TESTGEN approach inherits all glory, but also all limitations of a testing approach: The entire specification is reduced via specific *test purposes* and underlying *test hypothesis* (“pick one out of the equivalence class, and it’s going to be ok for all class members”) to a *finite* number of tests to be checked. These purposes and hypotheses may be difficult to justify and need careful inspection, more difficult than having just a universal statement over the entire input/output relation. On the other hand, testing can establish confidence over the **real system**, and makes no modeling assumptions — like the Simpl-approach — over the underlying hardware, the correct modeling of behavior of hardware components such as sensors, the compiler, and the equivalence of the assumed operational semantics of the used programming language(s) with the actually executed one. For this reason, it can be safely stated that for certifications of the highest-levels, a suitable *combination* of test and proof techniques will be necessary. Proofs for the higher levels of the models establishing the desired security properties in an TOE, tests for establishing that the assumptions made in the lower levels of the models correspond to the reality in the TOE.

1.6 Recommendations for CC Certifications

1.6.1 A Refinement Based Approach for CC Evaluation

Figure 1.3 presents a refinement scheme which implements different refinement steps from security policy model SPM to implementation. With this approach, the properties demonstrated on an abstract SPM are formally preserved down to the levels of the functional specification model FSP and a TOE specification design model, the TSD. At each level of abstraction the dedicated model and its associated proofs demonstrate the security properties and are compliant with the CC requirement. The use of a formal refinement methodology demonstrates the consistency between each refined model and preserves the properties demonstrated at high level of abstraction.

The evaluation of this kind of approach can be conducted in three different phases by the evaluator:

- Phase 1: Verification of the proof of the SPM formal specification. On the initial abstract model, a verification shall be conducted to check the relevance of the security objectives modeling in the formal model with the informal specification. A second point is the verification of the model soundness to assure that the model is not inconsistent (refers to chapter 1.4.2).
- Phase 2: Refinement of the SPM formal specification. A first step of this phase is the verification of the refinement process and methodology. On each refinement, verifications on the properties and on the soundness of the model are conducted. From the initial abstract model, on each intermediate concrete model, the evaluator checks the traceability (i.e. the traceability of the requirements) between models. An informal link can be considered between the last formal model of the TDS and the implementation. A bi-directional detailed traceability of the security requirements shall be managed between these two different artefacts to verify the implementation of the security requirements and that the implementation contains only desired requirements⁷.
- Phase 3: General and transverse activities. This last phase consists mainly of the verification on the proofs and on justifications on the tools used as support for development and design. The complete traceability from the security target to the implementation is verified including traceability between each refinement steps of formal models. During this phase, the evaluator replays the proofs and checks the consistency of the formal properties and assumptions defined on the environment and the context (see 1.3.5 last paragraph for details of facilities supplied by Isabelle/HOL). The use of keywords to report the proof of parts of the proof obligations is forbidden (for example the use of the **sorry** proof method, see chapter 1.4.2 for details).

When formal methods are used, some practices should be applied to facilitate the work of the evaluator and be more efficient.

- Formal models should be defined in accordance with some naming convention information and is a huge help for traceability.
- Formal models should be defined in accordance with “coding” rules ([Jae08]). The proofs associated can be replayed.
- Documentation and deliveries should respect templates and integrate traceability with requirements or elements from input specifications. From this point, the use of the Isabelle interface should be interesting with regard to its functionalities, refers to 1.3.5.

1.7 Summary

1.7.1 Background References

The most notable text describing the scientific history behind the LCF-family of HOL provers is done by Mike Gordon [Gor00]. It covers the beginning of the entire research programme

⁷to check that no parts of the code violate the security properties by side effects.

from 1972 to the mid-80ies, ranging from foundational issues of the logic over contributions to type-systems (as the “Hindley-Milner-Polymorphism”)[Mil78] to the issue of the practical, safe implementation of rewrites and decision procedures [Pau99].

The LCF research programme was in parallel to another notable source of nowadays interactive theorem proving technologies: the Automath-project. In 1968, N.G. de Bruijn designs the first computer program to check the validity of general mathematical proofs, using typed λ -calculi as a direct means to represent proof objects as such. The emphasis of this programme was initially on proof-checking; de Bruijn’s system Automath eventually checked every proposition in a primer that Landau had written for his daughter on the construction of real numbers as Dedekind cuts. A descendant of this family, which also has deeply influenced the Isabelle kernel design (proof objects, core inferences) is the Coq system (see <http://coq.inria.fr>).

Another notable survey on research programme is contained in the papers contained in *A Special Issue on Formal Proof* distributed by the American Mathematical Society (see <http://www.ams.org/notices/200811/>, but also [Hal08]), which presents nicely the relevance of modern ITP technology for purely mathematical problems (an argument, which has been strengthened recently by the formal proof of the Feit-Thompson theorem, whose precise formulation has haunted mathematicians for decades [Gon13], and the formal proof of the Kepler-conjecture, which is a known mathematical problem for about 400 years).

1.7.2 Concluding Remarks and a Summary

We have presented the Isabelle/HOL system and pointed out the essential arguments, why by a particular combination of system-architecture and methodology, the system is suited to give the currently highest possible guarantee on a formal proof in particular and a logical theory development in general. In a sense, Isabelle/HOL offers the same guarantees for logical systems as Coq [Jae08], and in some sense better guarantees than, for example, the B method or model-checkers like FDR. Isabelle/HOL is therefore a natural choice for evaluations in the higher certification levels EAL5 to EAL7 in the Common Criteria (CC) [Mem06].

If the methodological side-conditions are respected which can be reduced essentially to a number of syntactic checks, the formal consistency of the entire certification document containing formal specifications, proofs of consistency and the proofs of security properties, refinement-proofs between the different abstraction layers, and finally test-case generations as well as test-results can be guaranteed, and the evaluator can therefore concentrate on the more fundamental questions: does the model represent the right thing? are the modeling assumptions justified?

As “take-home-message” we would summarize these side-conditions as follows:

1. Use a trusted, unmodified Isabelle version from the distribution.
2. Check the restriction to definitional axioms only, enforce the use of “safe” specification constructs discussed here (Section 1.4).
3. Rule out axiomatization, **sorry**, their variants or disguised equivalents (such as oracle declarations).
4. In particular **sorry**’s or equivalent constructions in methodological proofs have to be ruled out.
5. Check the quick-and-dirty mode status.

6. Exploring a TOE interactively, for example by jEdit, which allows for inspecting theories and definitions, their animation, the checking of types and of proof details, is a great means to increase confidence for an evaluator. However, the final check should be done in a non-interactive mode (pretty-printing and display machinery is actually quite far from the kernel and can be erroneous in itself).
7. The main theorem in an CC evaluation is presumably of the form: “the security property X stated in the context of the security model Y is satisfied for the functional model Z under some conditions A in some locale B”. A skeptical evaluator may insist on proofs that A and B are actually satisfiable, under circumstances even in a constructive sense.
8. A conservative evaluator should restrict or ban ML-statements (with the possible exception of declarations of antiquotations), otherwise inspect ML-statements with particular care.

The internal code generator (also used in `code`-antiquotations and `value`-statements) stood the test of the time, but enjoys not quite the same level of trust as the proof facilities. The generation of proof objects for a complete theory is in principle possible, but should not be necessary except in case of a concrete suspicion of a fraudulent proof attempt.

Chapter 2

Style Guide

Chapter Authors (ordered according to beneficiary numbers): Sergey Tverdyshev, Oto Havle, Holger Blasum, Bruno Langenstein, Werner Stephan, Abderrahmane Feliachi, Yakoub Nemouchi, Burkhart Wolff, Freek Verbeek, Julien Schmaltz

2.1 Introduction

In contrast to issues related to consistency and the proper use of the HOL logic, its methodology and the Isabelle implementation (which were pointed out in the *mandatory part* in Section 1), we will present in some detail the specific pragmatics in our use of Isabelle/HOL within the EURO-MILS project. While a description of pragmatics (or: a Style Guide) is more open-ended and the alternatives are much less clear-cut, we will keep the focus much more on *our* choices and much less on possible alternatives; this document is therefore a significantly shorter one.

2.2 Rules

2.2.1 Basics

Rule: Follow Mandatory Guide and the AFP Style Guide Rules

The mandatory rules (captured in Section 1) concern the logical and methodological aspects of proper use of Isabelle/HOL as a logic and a formal methods environment. The necessity of their use is considered to be self-understood.

Additionally, we apply the so-called AFP style-guide, i.e. a set of rules imposed by the Isabelle Archive of Formal Proofs (AFP), <http://afp.sourceforge.net/submitting.shtml>, as submission guidelines. These rules can be seen as a set of best-practice rules established by the Isabelle developer team resulting from the rich experience in porting Isabelle theories from one Isabelle version to the next. These are in particular:

1. No use of the command `back`.
2. Instantiations must not use Isabelle-generated names such as `xa` use `Isar` or `rename_tac` to avoid such names.
3. No use of the command `smt`. The result of this command depends on external tools that are not under our control and may stop working in the future.

4. Apply scripts should be indented by subgoal as in the Isabelle distribution. If an apply command is applied to a state with $n+1$ subgoals, it must be indented by n spaces relative to the first apply in the sequence.
5. We prefer structured Isar proofs over apply style, but do not mandate them.
6. If there are proof steps that take significant time, i.e. longer than roughly 1 min, please add a short comment to that step, so maintainers will know what to expect.

Rule: No Obfuscation of Standard Isabelle Notation

Isabelle/HOL is a particularly flexible framework that allows on numerous levels the redefinition of syntax, proof- and toplevel commands. While this may be useful in many cases, it may also create confusion and misinterpretation by evaluators with respect to the *interpretation* of terms and proofs. We adopt therefor the general rule not to override Isabelle/HOL standard library syntax as described currently in Tobias Nipkow's Paper *What's in Main* ("Main" is the session name for Isabelle/HOL) coming with the Isabelle documentation and currently found under <http://isabelle.in.tum.de/website-Isabelle2013-2/dist/Isabelle2013-2/doc/main.pdf> for Isabelle version 2013-2. This paper contains a catalogue on standard syntax for the Logic, Orderings, Lattices, Sets, Lists, etc.

For example, for constant symbols of the basic logic it defines:

- $x = y, \text{True}, \text{False}, \neg P, P \wedge Q, P \vee Q, P \rightarrow Q, \forall x.P, \exists x.P, \exists!x.P, \text{THE } x.P.$
- `undefined :: 'a`
- `default :: 'a`

Additionally, the following syntactic abbreviations:

- $x \neq y \equiv \neg(x = y)$
- $P \leftrightarrow Q \equiv x = y$
- `if x then y else z` $\equiv \text{If } x \text{ } yz$
- `let x = e1 in e2` $\equiv \text{Let } e1(\lambda x. e2)$

Overloading of operators and syntactic abbreviations of the HOL-library as described in *What's in Main* should be avoided. New *alternative* syntactic notations may be helpful, but should be chosen with care. Generally speaking, the non-obfuscation principle holds also for the command and method language.

2.2.2 Modeling Style

Rule: Curried Style

Prefer the style of curried notation whenever possible; the major reason for this is that libraries and the proof engine are geared to it (un-tupling pairs is often a show-stopper for automatic proof derivations).

There are two major exceptions to this rule: relations and result-types of functions. Here, the library offers conversions for both curried formats and Cartesian products, more commonly used in mathematics.

Note that a function like $f :: \alpha \rightarrow (\beta \times \gamma) \text{ option} \rightarrow \delta$ is already in curried form; the Cartesian product is embedded in an option type.

Rule: Use Records

Use Records whenever possible (no “fst(fst(snd(fst(” mumbo jumbo), and attempt to structure incremental extensions of state spaces by extensible records.

Rule: Use Locales

Locales are “Functors” on formal theories. They should be used to factor out common parts, but also to specify high-level security concepts as in CC. However, over-generalizations should be avoided (it is well-known that this is a difficult balance).

Locale Instantiations ARE formal proofs of consistency. Note however that this does not necessarily mean constructivity: one can specify in HOL the Halting-Function via the Hilbert-Choice; this does mean that the function exists, however, alas, this does not mean that it is computable.

Some readers may have a clear leaning towards constructivism when it comes to instantiating locales. However, since Isabelle is a classical framework and we do use underspecification for the purpose of specification conciseness, a clear-cut line to constructivity is very difficult to establish. Inside the functor, there may be non-constructive parts, even if the instantiation itself is constructive.

The only strong *guarantee* for constructivity could possibly be given if the entire functional model is finally instantiated in a way that the code-generator can generate executable code.

Rule: Avoid Abstract Types if Possible

Abstract types (introduced by `typedef1 \<tau>`) should be used with care and reservation; they complicate the argument of constructivity drastically and represent an obstacle to automated test-case generation. On the other hand, using only parametric polymorphism instead drastically complicates the signatures of key functions which may also obscure the dependencies inside the model. As a consequence, the project decided to admit them at some places.

2.2.3 Formal Content

Isabelle is not just an interactive theorem proving environment. Beyond rich logical libraries as well as infrastructure for formal definitions and semi-automated proofs, it offers support for a particular *document model* comprising consistency-checking of the document’s content (called *formal content*) and type-setting. A document consists of a collection of files, called *session*, which can be of different file type, typically `*.thy`, `*.ML`, `*.tex`, and `*.sty` files.

Albeit a setup comprising even the `*.c`-sources of a larger project is in principle possible, the file types used in the EURO-MILS project are just `*.thy` and `*.tex`, so documents written in the Isabelle/Isar language and LaTeX. Since `*.ML` code can be used for obfuscation and intransparent use of Isabelle, we avoided them, which also excludes the effective use of other file types.

We describe *.thy-files a little more in detail. Generally speaking, it consists of a *header* declaring name and import relation to other files of the session, and then a sequence of (*toplevel*) *commands*. The latter were typically introduced by a keyword such as **typedef**, **definition**, **fun**, **lemma**, etc.

Commands can be arbitrarily interwoven with unstructured comments (text inside (* ... *)) or ... -- ... at the end of the line), for which we will adopt no particular style-convention, and structured comments which may have the syntax:

- header <text>, chapter <text>, section <text>, subsection <text>, subsubsection <text>, text <text> and text_raw <text>.

These were the text commands that can be used as *toplevel* Isar commands in Isabelle documents. In particular, LaTeX paragraphs can be generated by text_raw {*\paragraph{Paragraph name}*}.

- sect <text>, subsect <text>, subsubsect <text>. These sub-commands were used inside the Isar proof language which is used inside lemma ... commands.

Note that <text> can have the format "... " or {* ... *}; the latter is the richer environment and therefore the preferred format.

An important aspect of formal content is the possibility to give *references*, i.e. links to entities of various type, a formal status. This means that instead of “(see section 3.1)” (as text) or “theorem ext” (as text) inside a <text> of a structured comment, it is preferable to give these references a *formal status* as well which means that it can be machine checked during document generation.

Rule: Maximize Formal Content

As a general rule, if a comment does not serve as temporary remark intended to improve the development, it should be attempted to turn it into formal content.

Rule: Formalize External References by LaTeX Macros

References to external entities can be:

- References to entities of the CC document process, which were not necessarily part of an Isabelle document, for example *assumptions*, *security requirements*, *features*, *characteristics*, etc., stated inside the *Target of Evaluation* (TOE) document of the process.
- References to entities in the implementation, so to type declarations, variables or system calls in the TOE implementation.

A particular type of document-internal link is also supported by a LaTeX Macro-mechanism : References to headers, sections, subsections, figures and tables should be referenced by the LaTeX hyperref package (see <http://en.wikibooks.org/wiki/LaTeX/Hyperlinks>). For example, a reference for a subsection can be declared in the command:

```
subsection{* Foo Section \label{ssec:foo} *}
```

which can be referenced later inside formal text by:

```
text{* Later in the text, we refer to \autoref{ssec:foo} *}
```

Rule: Formalize Document-Internal Links by Isabelle Antiquotations

With respect to formal entities such as types, terms, theorems and definitions, but also file-names referring to files belonging to a document, etc. Isabelle offers a particular easy and flexible mechanism to support checking of these references during the document editing process integrated in the Isabelle IDE. It is called text-antiquotation. For the following formal entities, there are predefined antiquotations:

- `@{typ \<tau>}`. Check and print type τ .
- `@{const c}`. Check existence of c and print it.
- `@{term t}`. Type-check and print term t . Example:

```
text {* This sentence demonstrates quotations and
      antiquotations: @{term "%x y. x"} is a
      well-typed term.
*}
```

The highlighting in the Isabelle IDE allows to identify defined constants, bound and free variables by colour-codes.

- `@{prop \<phi>}`. Print and type-check proposition ϕ . Variants: `@{prop [display] \<phi>}`: print large proposition ϕ (with linebreaks). `@{prop [source] \<phi>}`: check proposition ϕ , print its input.
- `@{thm a}`. Print fact a . Variants: `@{thm a [no_vars]}`: Print fact a , fixing schematic variables. `@{thm [source] a}`: Check availability of fact a , print its name.
- `@{text s}`. Uninterpreted text s printed emphasized.
- `@{file foo}`. Print file-name foo ; check its existence inside the session.

See the “The Isabelle/Isar Reference Manual” coming with the Isabelle distribution for more details, in particular with respect to the antiquotation *attributes* in brackets.

A hint to the evaluator: by inserting a `typ \<tau>`, `term t`, `thm n`, formal entities of these categories can be checked inside the document with respect to existence, consistency, type-checkability, etc. If terms are defined in a constructive way, the Isabelle code-generator *can* produce a value for a given term t via the `value t` command, which can be a great help during inspection.

2.2.4 Layout Principles

Rule: Respect Blue Bar while Editing

Formal content text formulas and definitions should be kept under 100 char width per line — this length is represented by a blue bar in Isabelle/jedit. Text can be arbitrarily spaced, but exceedingly long lines produce overfull hboxes inside formula and uncontrolled layout during LaTeX generation.

Rule: Multiline Comments

Formal comments over several lines should be `text { * . . . * }`. By convention, do not use: `-- { * . . . * }` over several lines.

Please respect the 100 char per line rule also for inline comments.

2.3 Conclusion

Beyond the issues of logical consistency (which can be assured by a number of methodological rules and hopefully wise self-constraints described in the mandatory part), we adopted a number of more stylistic rules on the level of the linking between the various entities in the document leading to, as we believe, an increase of the consistency of the semi-formal aspects of the Isabelle documents.

Chapter 3

Compliance Statement

Chapter Authors (ordered according to beneficiary numbers): Sergey Tverdyshev, Oto Havle, Holger Blasum, Bruno Langenstein, Werner Stephan, Yakoub Nemouchi, Burkhart Wolff, Freek Verbeek, Julien Schmaltz

In our use of Isabelle/HOL, we claim conformance to Chapter 1 and Chapter 2.

3.1 Compliance to Section 1.7.2

1. We use Isabelle2013-2, obtained from the Isabelle distribution site [PNW13].
2. We describe the safe use of the following specification constructs:
 - Constants. The Isabelle **consts** command is not used.
 - Locales. Locales establish a scope where variable assignments and assumptions are valid. This can be used to establish general theorems without hardcoding them to a particular theory. A locale is used to allow to develop a generic (and peer-reviewable) theory (MCISK) that establishes separation properties for separation kernels. Then in Theories-step/Step_link_MCISK.thy, conformance to the locale is shown.
 - Extensible records. We use extensible records in the representation of system state.
 - Proofs. Structured Isar style is preferred. Proof automation is used whenever possible. We allow the use of **sledgehammer** for automatic proof search. However, in the generated theories we do *not* write the command **sledgehammer**; instead the output of **sledgehammer** is used in the proofs.
3. We never use the **axiom** command to insert arbitrary statements into the theories. We use conservative constructs instead:
 - Algebraic datatypes (example: *phys_address_t*).
 - Subset types (example: *kmem_consumption_f*)

Absence of **axiom** and **sorry** has been checked by textual search.

4. Methodological proofs: Soundness is assured by safe use of specification constructs, see point 2 above.
5. Absence of “quick-and-dirty” mode: Theories in [EUR15] have been compiled with *quick_and_dirty=false*.

6. We have checked theory correctness in a non-interactive mode by compiling [EUR15].
7. No statement is made that proofs are fully constructive. However, we have tried to avoid non-constructive constructs where feasible.
8. We use antiquotations, but do not use ML statements, this can be ascertained by textual search for *ml* and *ML*.

3.2 Compliance to Section 2

- Rule: Follow Mandatory Guide and the AFP Style Guide Rules
 - (Basics 1) The command **back** is not used. This has been checked by textual search.
 - (Basics 2) Isabelle-generated names (marked in brown color in the jEdit GUI) are not used by the theories nor the proofs we wrote. This has been checked by visual inspection.
 - (Basics 3) The command **smt** is not used. This has been checked by textual search.
 - (Basics 4) This rule is seen as a recommendation (“should”) and not mandatory. No compliance is claimed. Rationale: automatic tool support is under development [Kle15].
 - (Basics 5) Structured proofs are largely used (checked by visual inspection).
 - (Basics 6) There are no steps that take longer than one minute.
- Rule: No Obfuscation of Standard Isabelle Notation: We are not redefining (“overloading”) any Isabelle symbol (checked by visual inspection).
- Rule: Curried Style: Curried style is used for function definitions, for example the function in *port_waiting* in file Theories-step/Step_port.thy is defined in functional style (the definition does not take a tuple as argument, but a function).
- Rule: Use Records: Extensible records are used for definitions of state.
- Rule: Use Locales: Locales are used for MCISK.
- Rule: Avoid Abstract Types if Possible: Tuple definitions are used in a very restrained way. At all other places, records are used.
- Rule: Maximize Formal Content: Care has been taken to convert “(* *)” to document-visible comments “*text* { * }”, where possible.
- Rule: Formalize External References by LaTeX Macros: External references to PikeOS functions, parameters, and error codes are formalized as LaTeX macros.
- Rule: Formalize Document-Internal Links by Isabelle Antiquotations: Isabelle antiquotations (“@{thm x}”) have been used.
- Rule: Respect Blue Bar while Editing: We have checked that line-breaks largely are within 100-character “blue line” limit.
- Rule: Multiline Comments: We have used “-{ * ... *}” only for one-line comments.

Bibliography

- [And86] Peter B. Andrews. *An introduction to mathematical logic and type theory: to truth through proof*. Academic Press Professional, Inc., San Diego, CA, USA, 1986.
- [And02] Peter B. Andrews. *Introduction to Mathematical Logic and Type Theory: To Truth through Proof*. Kluwer Academic Publishers, Dordrecht, 2002.
- [Bal10] Clemens Ballarin. *Tutorial to Locales and Locale Interpretation*, 2010.
- [BMSW10] Sascha Böhme, Michal Moskal, Wolfram Schulte, and Burkhart Wolff. Hol-boogie - an interactive prover-backend for the verifying c compiler. *J. Autom. Reasoning*, 44(1-2):111–144, 2010.
- [BW10] Sascha Böhme and Tjark Weber. Fast LCF-style proof reconstruction for Z3. In *ITP*, pages 179–194, 2010.
- [BW13] Achim D. Brucker and Burkhart Wolff. On theorem prover-based testing. *Formal Aspects of Computing*, 25(5):683–721, 2013.
- [Chu40] Alonzo Church. A formulation of the simple theory of types. *Journal of Symbolic Logic*, pages 56–68, June 1940.
- [CKK⁺12] Pascal Cuoq, Florent Kirchner, Nikolai Kosmatov, Virgile Prevosto, Julien Signoles, and Boris Yakobowski. Frama-c: a software analysis perspective. In *International Conference on Software Engineering and Formal Methods (SEFM’12)*, pages 233–247. Springer, October 2012.
- [EUR15] EURO-MILS. Formal implementation of TOE inclusive formal proofs. Technical Report D31.3, EURO-MILS: Secure European Virtualisation for Trustworthy Applications in Critical Domains, FP7/2007-2013, 2015. Document in PDF format <https://euomils.technikon.com/Activity-A3/WP31-Assurance-Formal-Methods/D31-3-Implementation-Model/document.pdf>, Isabelle sources at <https://euomils.technikon.com/Activity-A3/WP31-Assurance-Formal-Methods/Theories-step/> and <https://euomils.technikon.com/Activity-A3/WP31-Assurance-Formal-Methods/Theories-trace/>.
- [FP13] Jean-Christophe Filliâtre and Andrei Paskevich. Why3 — where programs meet provers. In Matthias Felleisen and Philippa Gardner, editors, *Proceedings of the 22nd European Symposium on Programming*, volume 7792 of *Lecture Notes in Computer Science*, pages 125–128. Springer, March 2013.
- [GAK12] David Greenaway, June Andronick, and Gerwin Klein. Bridging the gap: Automatic verified abstraction of c. In *ITP*, pages 99–115, 2012.
- [GM93] Mike J. C. Gordon and Tom F. Melham. *Introduction to HOL*. Cambridge University Press, 1993.

- [Gon13] Georges Gonthier. Engineering mathematics: the odd order theorem proof. In *POPL*, pages 1–2, 2013.
- [Gor00] Mike Gordon. From LCF to HOL: a short history. In *Proof, Language, and Interaction*, pages 169–185. MIT Press, 2000.
- [Hal08] Thomas C Hales. Formal proof. *Notices of the AMS*, 55(11):1370–1380, 2008.
- [Har06] John Harrison. Towards self-verification of hol light. In *IJCAR*, pages 177–191, 2006.
- [Int12] ISO/IEC DIS 29119: Software and Systems Engineering—Software Testing. ISO Draft International Standard, July 2012.
- [Jae08] Éric Jaeger. Remarques relatives à l’emploi des méthodes formelles (déductives) en sécurité des systèmes d’information. 2008. 51 Boulevard de la Tour-Maubourg 75700 Paris SP 07, France.
- [KAMO14] Ramana Kumar, Rob Arthan, Magnus O. Myreen, and Scott Owens. HOL with definitions: Semantics, soundness, and a verified implementation. In *Interactive Theorem Proving - 5th International Conference, ITP 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 14-17, 2014. Proceedings*, pages 308–324, 2014.
- [KHS09] Gerwin Klein, Ralf Huuck, and Bastian Schlich. Operating system verification. *J. Autom. Reasoning*, 42(2-4):123–124, 2009.
- [Kle15] Gerwin Klein. Gerwin’s style guide for Isabelle/HOL. part 1: Good proofs, 2015. <http://proofcraft.org/blog/isabelle-style.html>, accessed 29 May 2015.
- [MB08] Leonardo Moura and Nikolaj Bjorner. Z3: An efficient SMT solver. In C.R. Ramakrishnan and Jakob Rehof, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, volume 4963 of *Lecture Notes in Computer Science*, pages 337–340. Springer Berlin Heidelberg, 2008.
- [Mem06] The Common Criteria Recognition Agreement Members. Common Criteria for Information Technology Security Evaluation. <http://www.commoncriteriaportal.org/>, September 2006.
- [Mil78] Robin Milner. A theory of type polymorphism in programming. *Journal of Computer and System Sciences*, 17(3):348 – 375, 1978.
- [MOK13] Magnus O. Myreen, Scott Owens, and Ramana Kumar. Steps towards verified implementations of hol light. In *ITP*, pages 490–495, 2013.
- [MTM97] Robin Milner, Mads Tofte, and David Macqueen. *The Definition of Standard ML*. MIT Press, Cambridge, MA, USA, 1997.
- [MW79] R. Milner and C.P. Wadsworth. *Edinburgh LCF: A Mechanized Logic of Computation*. Lecture Notes in Computer Science. Springer, 1979.

- [NPW02] Tobias Nipkow, Larry C. Paulson, and Markus Wenzel. *Isabelle/HOL—A Proof Assistant for Higher-Order Logic*, volume 2283 of *Lecture Notes in Computer Science*. Springer-Verlag, Heidelberg, 2002.
- [Pau99] Lawrence C. Paulson. A generic tableau prover and its integration with isabelle. *J. UCS*, 5(3):73–87, 1999.
- [PNW13] Larry Paulson, Tobias Nipkow, and Makarius Wenzel. Isabelle, 2013. <http://isabelle.in.tum.de/website-Isabelle2013-2/index.html>, accessed 26 May 2015.
- [PP10] Leaf Petersen and Enrico Pontelli, editors. *Proceedings of the POPL 2010 Workshop on Declarative Aspects of Multicore Programming, DAMP 2010, Madrid, Spain, January 19, 2010*. ACM, 2010.
- [TN] Markus Wenzel Tobias Nipkow, Lawrence C. Paulson. *A Proof Assistant For Higher-Order Logic*.
- [WB89] Philip Wadler and Stephen Blott. How to make ad-hoc polymorphism less ad-hoc. In *POPL*, pages 60–76, 1989.
- [Wen] Makarius Wenzel. *The Isabelle/Isar Reference Manual*.
- [Wen97] Markus Wenzel. Type classes and overloading in higher-order logic. In *TPHOLs*, pages 307–322, 1997.
- [Wen02] Markus M Wenzel. *Isabelle/Isar—a versatile environment for human-readable formal proof documents*. PhD thesis, Technische Universität München, Universitätsbibliothek, 2002.
- [Wie06] Freek Wiedijk. *The Seventeen Provers of the World: Foreword by Dana S. Scott (Lecture Notes in Computer Science / Lecture Notes in Artificial Intelligence)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006.

Secure European virtualisation for trustworthy applications in critical domains. The mission of the EURO-MILS project is to develop a solution for virtualization of heterogeneous resources and provide strong guarantee for isolation of resources by means of Common Criteria certification with usage of formal methods.

www.euromils.eu

for further information please contact the coordinator
TECHNIKON Forschungs- und Planungsgesellschaft mbH
coordination@euromils.eu

