

MECHANIZING NONSTANDARD REAL ANALYSIS

JACQUES D. FLEURIOT AND LAWRENCE C. PAULSON

Abstract

This paper first describes the construction and use of the hyperreals in the theorem-prover Isabelle within the framework of higher-order logic (HOL). The theory, which includes infinitesimals and infinite numbers, is based on the hyperreal number system developed by Abraham Robinson in his nonstandard analysis (NSA). The construction of the hyperreal number system has been carried out strictly through the use of definitions to ensure that the foundations of NSA in Isabelle are sound. Mechanizing the construction has required that various number systems including the rationals and the reals be built up first. Moreover, to construct the hyperreals from the reals has required developing a theory of filters and ultrafilters and proving Zorn's lemma, an equivalent form of the axiom of choice.

This paper also describes the use of the new types of numbers and new relations on them to formalize familiar concepts from analysis. The current work provides both standard and nonstandard definitions for the various notions, and proves their equivalence in each case. To achieve this aim, systematic methods, through which sets and functions are extended to the hyperreals, are developed in the framework. The merits of the nonstandard approach with respect to the practice of analysis and mechanical theorem-proving are highlighted throughout the exposition.

1. *Introduction*

In the early 1960's, Abraham Robinson finally provided a rigorous foundation for the use of infinitesimals in analysis by developing the new concept of *nonstandard analysis* (NSA) [29]. The idea was to introduce a new number system known as the *hyperreals*, which contains not only the real numbers but also infinitesimals and infinite numbers. The notions of infinitesimals and other nonstandard numbers introduce many subtleties into the theory that need to be dealt with.

In this paper, we first describe the constructions of Robinson's hyperreals in Isabelle. Our approach is purely definitional, to ensure that infinitesimals and other nonstandard numbers have a sound foundation in the system. To reach our goal has required constructing the various number systems leading to the reals, and then going one step further to define the hyperreals by working on sequences of reals. The hyperreals have considerable intrinsic interest since they exhibit many new properties. Moreover, as a tool, they are of great value to the formalization of analysis — an aspect that will be described as we report on the mechanization of nonstandard real analysis.

Received 12 November 1999, revised 14 April 2000; published 30 June 2000.

2000 Mathematics Subject Classification 03C20, 03H05, 26E35, 03B35, 68T15, 03B15

© 2000, Jacques D. Fleuriot and Lawrence C. Paulson

This paper consists of two main parts: Sections 2–8 are concerned with the construction of the hyperreals, while Sections 9–16 describe their application to mechanized analysis. We start by giving a description of Isabelle, and of the HOL object logic in which this work was carried out.

2. Isabelle/HOL

Isabelle [25] is a generic theorem-prover, written in ML, into which users can encode their own object-level logics. Examples of such object logics are higher-order logic (HOL), Zermelo–Fraenkel set theory (ZF), and first-order logic (FOL). Terms from the object logics are represented and manipulated in Isabelle’s intuitionistic higher-order meta-logic, which supports polymorphic typing.

2.1. Theories in Isabelle

Isabelle’s theories provide a hierarchical organization for the syntax, declarations and axioms of a mathematical development, and are developed using theory definition files [25]. A typical theory file will organize the definitions of types and functions. It may also contain the primitive axioms that are asserted (without proofs) by the user. A particular theory will usually collect (in a separate file) the proven named theorems, and make them available to all its children theories.

The meta-level connectives are implication (\implies), the universal quantifier and equality. In Figure 1, we give the description of some of the notations used in Isabelle/HOL. Throughout the presentation, we will mostly be using conventional mathematical notations when describing our development. However, there are cases where we might use the ASCII notations actually used to express terms and rules in Isabelle as explicit examples.

An inference rule with n premises or antecedents has the following form in Isabelle:

$$[[\phi_1; \dots; \phi_n]] \implies \psi.$$

This abbreviates the nested implication $\phi_1 \implies (\dots \phi_n \implies \psi \dots)$. Such a rule can also be viewed as the proof state with subgoals ϕ_1, \dots, ϕ_n and main goal ψ [25]. Alternatively, this can be viewed as meaning ‘if $\phi_1 \wedge \dots \wedge \phi_n$ then ψ ’.

2.2. Proof construction

Rules can be combined in various ways to derive new ones using higher-order resolution; this process is known as ‘proof construction’, and can proceed in both backward and forward directions.

- In backward fashion, the user supplies a goal and reduces it to simpler subgoals by applying existing rules until they are solved. A goal is solved when it becomes the instance of some previously proved theorem.
- In forward proofs, the antecedents or assumptions of a rule can be resolved with other rules to derive new assumptions. This process can be carried on until either the conclusion is the instance of some assumption, or the goal is an instance of a theorem.

2.3. Higher-order logic in Isabelle

One of Isabelle’s logics is HOL, a higher-order logic that supports polymorphism and type constructors. Isabelle/HOL is based on Gordon’s HOL90 theorem-prover [12], which

<i>syntax</i>	<i>description</i>
&	\wedge , and
~	\neg , not
==>	\implies , implication (meta level)
-->	\longrightarrow , implication (object level)
=	\equiv , if and only if
! or ALL	\forall , for all
? or EX	\exists , exists
@	ε , Hilbert choice
%	λ , lambda abstraction
- A	\overline{A} , set complement
Union c	$\bigcup c$, union over sets of sets

Figure 1: ASCII notation for HOL.

itself originates from Church’s paper [7]. Isabelle/HOL is well developed and widely used. It has a wide library of theories defined in it, including the natural numbers, set theory, well-founded recursion, inductive definitions and equivalence relations. Isabelle/HOL has been applied to reasoning in many fields, including the verification of security protocols [26] and verifying the type system of the Java programming language [24].

Though Isabelle is mainly used interactively as a proof assistant, it also provides substantial support for automation. It has a generic simplification package, which is set up for many of the logics, including HOL. Isabelle’s simplifier performs conditional and unconditional rewritings and makes use of context information [25]. The user is free to add new rules to the simplification set (the *simpset*), either permanently or temporarily. Isabelle also provides a number of generic automatic tactics that can execute proof procedures in the various logics. The automatic tactics provided by Isabelle’s *classical reasoner* include a fast tableau prover called `Blast_tac`, coded directly in ML, and `Auto_tac`, which attempts to prove all subgoals by a combination of simplification and classical reasoning. Other powerful theorem-proving tactics include those which, unlike `Blast_tac`, construct proofs directly in Isabelle: for example, `Fast_tac` implements a depth-first search automatic tactic.

2.3.1. The HOL methodology

Isabelle/HOL has been chosen as the logic in which to carry out our proofs. One of the main reasons is that it provides strong typing, and therefore ensures that only type correct terms are accepted. Moreover, the HOL methodology, an approach that originated in Gordon’s early work using HOL88, admits only conservative extensions to a theory. This means defining and deriving the required mathematical notions rather than postulating them. The definitional approach of HOL requires that assertions be proved about some model instead of being postulated. Such a rigorous definitional extension guarantees consistency, which cannot be ensured when axioms are introduced. As pointed out by Harrison [14], such an approach provides a simple logical basis that can be seen to be correct once and for all. With regard to the foundations of infinitesimals, the definitional approach is certainly advisable when one considers the numerous inconsistent axiomatizations that have been proposed in the past [9]. Of course, care still needs to be exercised, as a wrong definition will almost certainly yield the wrong properties.

3. Properties of an infinitesimal calculus

We first look at some of the requirements for a set of infinitesimals that could be useful for the development of an infinitesimal calculus. Keisler [20] and Vesley [33], for example, discuss the various properties that need to hold for developing a calculus for the infinitesimals. Let the set `Infinitesimal` denote the set of infinitesimals, where an infinitesimal can, for the time being, be viewed intuitively as a number smaller in magnitude than all positive reals.

We would like the following properties.

- 1) Zero is an `Infinitesimal`.
- 2) There is a nonzero infinitesimal.
- 3) `Infinitesimal` is a ring.

It might seem reasonable to want the following properties as well.

- 4) `Infinitesimal` is a subring of the real numbers \mathbb{R} .
- 5) `Infinitesimal` is an ideal in \mathbb{R} :

$$\forall r \in \mathbb{R} \forall x \in \text{Infinitesimal } rx \in \text{Infinitesimal}.$$

- 6) Also, we expect `Infinitesimal` to be non-Archimedean:


$$\exists x \in \text{Infinitesimal}. \forall n. nx < 1.$$

The above properties, (1)–(6), look sufficient for a simple theory of infinitesimals, but unfortunately such a theory would be inconsistent. Furthermore, as Vesley [33] notes, if \mathbb{R} is the set of classical reals, then *any nontrivial ideal in \mathbb{R} is equal to \mathbb{R}* . Thus, if `Infinitesimal` satisfies properties (2), (4) and (5) then `Infinitesimal` = \mathbb{R} . This problem is tackled in NSA by dispensing with property (4). Instead, using the axioms of classical set theory, a set \mathbb{R}^* of hyperreals is obtained with properties that include `Infinitesimal` \subseteq \mathbb{R}^* , $\mathbb{R} \subseteq \mathbb{R}^*$ and properties (1)–(3) and (6), but *not* `Infinitesimal` \subseteq \mathbb{R} , and therefore not property (4). As a result, property (5) now requires `Infinitesimal` to be an ideal in the set of finite members of \mathbb{R}^* . This set includes the reals and the infinitesimals, amongst other numbers.

Though an axiomatic approach seems the easiest way to get quickly to the infinitesimals, there is always the possibility that the set of axioms might lead to an inconsistency, as we saw above. We would rather have a development of infinitesimals that is guaranteed to be sound — especially, given the stormy history of infinitesimals.

4. Constructions leading to the reals

There are various classical methods in existence in the literature on the construction of the various number systems. The usual approach is to arrange them in a lattice respecting the inclusions between the sets. Let \mathbb{Z} , \mathbb{Q} , \mathbb{R} be the sets of integers, rationals, and reals respectively, and \mathbb{Z}^+ , \mathbb{Q}^+ , \mathbb{R}^+ be their positive elements. Note that \mathbb{Z}^+ is the set of elements of type `pnat`.



As can be seen in the figure, there are several ways to reach \mathbb{R} from \mathbb{Z}^+ . These various paths, however, differ greatly in the technical details of the constructions along them. Conway [8] suggests that there is a best way through the lattice to the reals that avoids, as much as possible, case splits. These are tedious and unnecessary complications that are often treated superficially in textbooks. Conway proposes the following general methods that we implement in Isabelle.

To add negative numbers, that is to proceed, for example, from \mathbb{R}^+ to \mathbb{R} , the signed number, $x \in \mathbb{R}^+$, is represented as an ordered pair of unsigned numbers (a, b) , meaning $a - b$, and the equivalence relation

$$(a, b) \sim (c, d) \iff a + d = b + c \tag{1}$$

is used. This is better than the obvious approach of the signed-magnitude representation, which leads to too much case-splitting.

Similarly, one can go from \mathbb{Z} to \mathbb{Q} or from \mathbb{Z}^+ to \mathbb{Q}^+ by taking ordered pairs (a, b) meaning a/b and the equivalence relation

$$(a, b) \sim (c, d) \iff a \cdot d = b \cdot c. \tag{2}$$

To proceed from \mathbb{Q} to \mathbb{R} or from \mathbb{Q}^+ to \mathbb{R}^+ , the method of *Dedekind cuts* is used. There are several other methods available, such as Cauchy sequences and positional expansions [14]. The best path, as suggested by Conway, is $\mathbb{Z}^+ \rightarrow \mathbb{Q}^+ \rightarrow \mathbb{R}^+ \rightarrow \mathbb{R}$.

4.1. *Equivalence relations in Isabelle/HOL*

We use Isabelle’s `Equiv` theory, which defines equivalence relations in higher-order set theory, to define the new type of positive rationals. First, we recall the definitions of equivalence relations, set quotients and equivalence classes:

Definition 4.1. A relation \sim is said to be an *equivalence relation* if and only if it is reflexive ($x \sim x$), symmetric ($x \sim y \implies y \sim x$), and transitive ($x \sim y \wedge y \sim z \implies x \sim z$).

Definition 4.2. Given an equivalence relation \sim on a set S , then the *quotient* of S with respect to \sim is the set of all equivalence classes, and is defined by $S/\sim \equiv \{[x] \mid x \in S\}$ where $[x] \equiv \{y \in S \mid x \sim y\}$.

The set of all equivalence classes S/\sim is called the *quotient set of S* by \sim , and a member of an equivalence class is often referred to as a *representative* of the class.

```

PRAT = PNAT + Equiv +

constdefs
  (* equivalence relation *)
  pratreel :: "(pnat * pnat) * (pnat * pnat) set"
  "pratreel ≡ {p. ∃ a b c d. p = ((a,b),(c,d)) ∧ a*d = b*c}"

typedef
  prat =
    "{x::(pnat*pnat). True}/pratreel" (Equiv.quotient_def)

instance
  prat :: {ord, plus, times}

constdefs
  prat_of_pnat :: pnat ⇒ prat
  "prat_of_pnat m ≡ Abs_prat(pratreel^{(m,Abs_pnat 1)})"

  qinv      :: prat ⇒ prat
  "qinv Q ≡ Abs_prat(⋃(x,y)∈Rep_prat(Q). pratreel^{(y,x)})"

defs
  prat_add_def
  "P + Q ≡ Abs_prat(⋃(a,b)∈Rep_prat(P). ⋃(c,d)∈Rep_prat(Q).
    pratreel^{(a*d + b*c, b*d)})"

  ...

  prat_less_def
  "P < (Q::prat) ≡ ∃T. P + T = Q"

end

```

Figure 2: Isabelle/HOL theory for rationals using equivalence classes.

4.2. Example: constructing \mathbb{Q}^+ from \mathbb{Z}^+

In this section, we illustrate, by means of an example, how a new type can be introduced in Isabelle as the quotient set of some equivalence relation. We also show how primitive functions are defined on the new type using *abstraction* and *representation* functions. Other operations derived from the primitive functions are also introduced.

The theory PRAT, shown in Figure 2 and developed on our way to the reals (and beyond), defines the type `prat` of positive rational numbers and its associated operations. The new type is defined on pairs of elements of `pnat`, which denotes the positive natural numbers, introduced as an explicit type in Isabelle.

Under the `constdefs` keyword, we declare and define the equivalence relation (2) specified at the beginning of Section 4 above, that enables us to proceed from \mathbb{Z}^+ to \mathbb{Q}^+ in the lattice:

$$\text{pratreel} \equiv \{p. \exists a b c d. p = ((a, b), (c, d)) \wedge a \cdot d = b \cdot c\}.$$

Using `typedef`, we declare the new type `prat`:

$$\text{prat} \equiv \{x. \text{True}\} / \text{pratreel} \quad (\text{Equiv.quotient_def}).$$

The representing set of elements is defined as the set of equivalence classes of fractions;

that is, the set of equivalence classes consisting of ordered pairs of positive natural numbers. The theorem `quotient_def` (from the theory `Equiv`) acts as a witness to prove the non-emptiness of the new type, and is given in brackets next to the new type. Non-emptiness needs to be proved to ensure that the quantifier rules of HOL are sound [27], otherwise the new type is rejected.

Once a new type has been successfully introduced, Isabelle provides coercion functions — the abstraction and representation functions — that enable us to define basic operations on the new type. Thus, in this particular example, the functions

$$\begin{aligned} \text{Abs_prat} &:: (\text{pnat} * \text{pnat}) \text{ set} \Rightarrow \text{prat} \\ \text{Rep_prat} &:: \text{prat} \Rightarrow (\text{pnat} * \text{pnat}) \text{ set} \end{aligned}$$

are added to the theory such that `prat` is isomorphic to

$$\{x. \text{True}\} / \text{pratrel}$$

by `Rep_hyprat` and its inverse `Abs_prat`. Using these functions and other operations from Isabelle’s `Set` and `Equiv` theories, we are now ready to define operations on the positive rationals. For example, the inverse function `qinv`, which swaps the elements of the ordered pairs (x, y) representing x/y around to give y/x , is constructed in Isabelle by:

$$\text{qinv } Q \equiv \text{Abs_prat } (\bigcup (x, y) \in \text{Rep_prat } (Q). \text{pratrel}^{\wedge}\{(y, x)\})$$

where

$$\begin{aligned} \bigcup x \in A. B[x] &\equiv \{y. \exists x \in A. y \in B\} && \text{(union of family of sets);} \\ r^{\wedge} s &\equiv \{y. \exists x \in s. (x, y) \in r\} && \text{(image of set } s \text{ under relation } r). \end{aligned}$$

Once the primitive operations such as addition and multiplication have been defined, we can use them to derive other operations such as the ordering relation:

$$P < Q \equiv \exists T. P + T = Q.$$

We then show that the operations on the new type respect the various field properties, and that we have indeed defined the densely ordered (but not Dedekind-complete) field of the positive rationals.

As a final remark, it is worth noting that the use of equivalence classes leads to simpler machine proofs than using notions of greatest common divisors (`gcd`) to choose unique representatives.

4.3. A few important theorems

In this section, some of the more important theorems that we proved during our constructions leading up to the reals are given. We are especially concerned with those that will be needed to establish properties of hyperreals and nonstandard real analysis later on.

Theorem 4.1 (Completeness of the reals). *The supremum property states that every non-empty set of reals X that has an upper bound has a least upper bound:*

$$\begin{aligned} \forall X. (\exists x. x \in X) \wedge (\exists U. \forall x \in X. x \leq U) \\ \implies \exists u. (\forall x \in X. x \leq u) \wedge \forall u'. (\forall x \in X. x \leq u') \implies u \leq u'. \end{aligned}$$

This simple result has far-reaching implications since it rules out the existence of infinitely small quantities or infinitesimals in \mathbb{R} .

Theorem 4.2 (The Archimedean property for the reals). *Any such infinitesimal in \mathbb{R} would mean that its reciprocal is an upper bound of \mathbb{N} in \mathbb{R} , thereby contradicting the Archimedean property:*

$$\forall x. \exists n. x < n.$$

Various mechanizations of standard analysis (see, for example, Harrison’s work using the HOL-Light system[13, 14]) have developed theories of limits, derivatives, continuity of functions and so on, taking as their foundations the real numbers. Our work, however, will now go one step further, and show how the reals can be used to build a richer number system.

5. Filters and ultrafilters

In this section, the preliminaries necessary to our construction are presented. The definitions and theorems that we need, and their formalization in the set theory of Isabelle/HOL, are reviewed. Our aim is to establish an equivalence relation on the set of all infinite sequences of reals, and use the system of equivalence classes as a model for \mathbb{R}^* . We start with the concept of a filter.

Definition 5.1. Let S be any non-empty set. A filter \mathcal{F} over S is a collection of subsets of S such that

- F1) $S \in \mathcal{F} \wedge \emptyset \notin \mathcal{F}$;
- F2) $X \in \mathcal{F} \wedge Y \in \mathcal{F} \implies X \cap Y \in \mathcal{F}$;
- F3) $X \in \mathcal{F} \wedge X \subseteq Y \subseteq S \implies Y \in \mathcal{F}$.

Every filter is a *nonempty* collection of subsets since $S \in \mathcal{F}$, and filters are closed under finite intersection and supersets. There are numerous examples of filters including the *trivial filter* $\{S\}$ and, if S is infinite, the *Fréchet* or *cofinite* filter $\{X. \text{finite } (S - X)\}$. In Isabelle, we develop a theory `Filter`, and formalize the notions described above as follows:

$$\begin{aligned} \text{Filters } S \equiv \{X. & F \subseteq \text{Pow } S \wedge S \in F \wedge \\ & (\forall X \in F. \forall Y \in F. X \cap Y \in F) \wedge \\ & (\forall X Y. X \in F \wedge Y \subseteq S \implies Y \in F)\}. \end{aligned}$$

We note in the above definition the occurrence of `Filters` S , which is defined to be the set of all filters over S . We adopt this general approach of defining sets of the various structures that are dealt with for clarity; this is possible since in Isabelle/HOL’s set theory the type α set is isomorphic to the type $\alpha \Rightarrow \text{bool}$ [25].

Let us mention some of the terminology often encountered when filters and related concepts are used. A set $X \subseteq S$ is sometimes said to be *large* [30] or *quasi-big* [15] if $X \in \mathcal{F}$. Other terms used include *residual* or *generic* when dealing with directed sets or Baire category theory. Moreover, and of relevance to our development, a condition P on points $x \in S$ is said to be satisfied *almost everywhere* (a.e.) or *almost always*, or is \mathcal{F} -true or *almost true*, if the set $\{x \in S. P \text{ is satisfied at } x\}$ is a member of \mathcal{F} .

A refinement of the concept of a filter is now introduced by defining the notion of an ultrafilter over the nonempty set S .

Definition 5.2. An ultrafilter \mathcal{U} over S is a filter over S such that

- U1) $\mathcal{U} \subseteq \mathcal{F} \wedge \mathcal{F} \in \text{Filters } S \implies \mathcal{U} = \mathcal{F}$.

An ultrafilter is thus a *maximal* filter; that is, a filter that cannot be enlarged. An ultrafilter (and hence a filter) is said to be *free* if and only if it does not contain any finite sets. A filter which is not free is said to be *fixed*. We are mainly interested in free ultrafilters. The definitions used in Isabelle’s `Filter` theory follow:

$$\begin{aligned} \text{Ultrafilters } S &\equiv \{X. X \in \text{Filters } S \wedge \\ &\quad \forall G \in \text{Filters } S. X \subseteq G \longrightarrow X = G\}; \end{aligned}$$

$$\begin{aligned} \text{FreeUltrafilters } S &\equiv \{X. X \in \text{Ultrafilters } S \wedge \\ &\quad \forall x \in X. \neg \text{finite } x\}. \end{aligned}$$

We proceed to prove various properties of filters, ultrafilters and so on from these definitions. These include a theorem about ultrafilters that reads as follows.

Theorem 5.1. *\mathcal{U} is an ultrafilter on S if and only if for any subset A of S , either A belongs to \mathcal{U} or else its complement $S - A$ belongs to \mathcal{U} , but not both:*

$$\mathcal{U} \in \text{Ultrafilters } S \iff (\mathcal{U} \in \text{Filters } (S) \wedge \forall A \in \text{Pow } S. A \in \mathcal{U} \vee S - A \in \mathcal{U}).$$

The content of this theorem is critically important to our development, and an outline of its proof in Isabelle is given below.

Proof. Suppose that \mathcal{U} is a filter such that for every $A \subseteq S$ either $A \in \mathcal{U}$ or $S - A \in \mathcal{U}$. Let G be a *superfilter* of \mathcal{U} ; that is, a filter such that $\mathcal{U} \subseteq G$, and suppose that $B \in G$ and $B \notin \mathcal{U}$. But then, from our initial assumption, it follows that $S - B \in \mathcal{U} \subseteq G$, and so $\emptyset = B \cap (S - B) \in G$ which contradicts property (F1) for a filter. Hence there is no proper filter including G , and so \mathcal{U} is an ultrafilter.

Conversely, suppose that \mathcal{U} is an ultrafilter and $A \notin \mathcal{U}$. Define a set $G \equiv \{X \subseteq S. \exists J \in \mathcal{U}. A \cap J \subseteq X\}$. Then $\mathcal{U} \subseteq G$ and $\mathcal{U} \neq G$ since $A \in G$, and so G cannot be a filter since by assumption \mathcal{U} is maximal. But G is not empty, and if $B, C \in G$ and $B \subseteq D$ then $B \cap C \in G$ and $D \in G$ (verifying conditions (F2) and (F3) for G to be a filter). Since $S \in G$, G can fail to be a filter only if $\emptyset \in G$. That is, we have $A \cap J = \emptyset$ for some $J \in \mathcal{U}$ for which we must then have $J \subseteq (S - A)$. It follows that $S - A \in \mathcal{U}$. \square

From this result, it can be seen, using the axiom of choice, that the Fréchet filter on an infinite set \mathcal{S} is not an ultrafilter, though it follows that it is free. What is needed to progress any further in the development is to show the existence of a free ultrafilter on any infinite set — this result is a corollary of the important *ultrafilter theorem* [16, 30]. Using the result above, we can see that for an ultrafilter \mathcal{U} to be free, every cofinite subset of \mathcal{S} , and hence the Fréchet filter, has to be contained in \mathcal{U} . This result will be useful to us in Section 5.2 but first, we give an overview of our proof of *Zorn’s lemma* and how we appeal to it to guarantee the existence of an ultrafilter. We then extend this result, and show that the ultrafilter can be free as well.

5.1. Zorn’s lemma

The existence of free ultrafilters is not obvious at first sight. To show that the ultrafilter theorem holds and to carry out our construction, we need Zorn’s lemma. This is an equivalent form of the axiom of choice (AC), and first needs to be proved in Isabelle/HOL.

Lemma 5.1 (Zorn’s lemma). *Let S be a nonempty set of sets such that each chain $c \subseteq S$ has an upper bound in S . Then S has a maximal element y ; that is, a set $y \in S$ such that no member of S properly contains y .*

The statement of Zorn’s lemma involves the idea of a partially ordered set and related concepts. We present briefly various mathematical concepts, and theorems about them, needed in Isabelle/HOL to express Zorn’s lemma.

Paulson has already proved Zorn’s lemma in Isabelle’s Zermelo–Fraenkel set theory (Isabelle/ZF) [28] by mechanizing a paper by Abrial and Laffitte [1]. Reporting on the mechanization, Paulson remarks that the formal language used by Abrial and Laffitte is close to higher-order logic, and thus should be useful to Isabelle/HOL amongst other proof assistants. In our current work, we have adapted the mechanization of Zorn’s lemma developed in Isabelle/ZF to Isabelle/HOL. Below, we briefly mention how our formalization in Isabelle/HOL compares with the one in Isabelle/ZF.

The definitions used by Abrial and Laffitte require the *choice* operator since, starting from AC, they prove Hausdorff’s maximal principle and then derive Zorn’s lemma. Unlike its ZF counterpart, Isabelle/HOL provides such an operator, the so-called *Hilbert epsilon operator*, ε . Thus, the formulation of the various theorems in Isabelle/HOL is somewhat simpler than that given by Paulson for ZF. The latter requires that the existence of the choice function be stated explicitly as a temporary additional assumption [28].

We also use Isabelle’s inductive package to define a set that is totally ordered by set inclusion. In general, the construction of the inductive set relies on defining a suitable successor function which, in our case, is defined using the choice operator:

$$\text{succ } S \ c \equiv \text{if } (c \notin \text{chain } S \vee c \in \text{maxchain } S) \\ \text{then } c \text{ else } (\varepsilon c'. c' \in \text{super } S \ c).$$

Our other definitions of set of chains, super chains and maximal chains are similar to those in Isabelle/ZF. Note that the definitions suppose that the set S has some *partial ordering* defined on it, which is denoted by \leq :

$$\text{chain } S \equiv \{F. F \subseteq S \wedge (\forall x \in F. \forall y \in F. x \leq y \vee y \leq x)\} \\ \text{super } S \ c \equiv \{d. d \in \text{chain } S \wedge c \subseteq d\} \\ \text{maxchain } S \equiv \{c. c \in \text{chain } S \wedge \text{super } S \ c = \emptyset\}.$$

We tried to simplify these definitions at first by removing references to the inductive set S , since it is actually used by Abrial and Laffitte to provide typing in their version of ZF. Thus, S as a parameter seems redundant when working in Isabelle’s typed higher-order logic. However, relying on the type made some of our proofs about ultrafilters unnecessarily complicated, and prompted us to refer explicitly to the underlying set in definitions, and hence in our proof of Zorn’s lemma. In outline, with these definitions, we prove the theorem of Hausdorff: every partially-ordered set contains a maximal chain. So, with the subset relation as the partial ordering on S , we have

$$\exists c. c \in \text{maxchain } S.$$

We then consider an upper bound u of such a maximal chain c — this is guaranteed to exist according to the premise of Zorn’s lemma. The last step in the proof simply involves showing that u is in fact a maximal element that we are looking for. Expressed formally in Isabelle, the following theorem is established:

$$\forall c \in \text{chain } S. \exists u \in S. \forall x \in c. x \subseteq u \\ \implies \exists y \in S. \forall x \in S. y \subseteq x \longrightarrow y = x.$$

5.2. The ultrafilter theorem

The ultrafilter theorem (UFT) is a complicated but important principle that lies midway between AC and the axiom of choice for finite sets [30]. Moreover, the ultrafilter theorem, like the axiom of choice, has many important equivalent forms. Schechter presents and discusses twenty-five of these, occurring in many areas of mathematics [30], and points to the many more equivalents occurring in the literature. The version that we are interested in is as follows.

Theorem 5.2 (Ultrafilter theorem (Cartan): UFT). *If \mathcal{F} is a filter on a set S then there is an ultrafilter \mathcal{U} on S with $\mathcal{F} \subseteq \mathcal{U}$.*

This result can be proved using Zorn’s lemma. In fact, we are really interested in proving a corollary of the ultrafilter theorem about the existence of free ultrafilters. (Some authors like Hoskins [16] and Keisler [20] state the corollary — or even one of its special cases — as the actual ultrafilter theorem.)

Corollary 5.1. *On every infinite set there exists a free ultrafilter. Expressed in Isabelle, we want to prove that*

$$\neg \text{finite } S \implies \exists u. u \in \text{FreeUltrafilters } S.$$

To do so, we define in the theory `Filter`, the set, `SuperFréchet S`, of all filters on S that contain the Fréchet filter (that is, the set of superfilters of the Fréchet filter):

$$\begin{aligned} \text{Fréchet } S &\equiv \{A. \text{finite } S - A\}; \\ \text{SuperFréchet } S &\equiv \{G. G \in \text{Filters } S \wedge \text{Fréchet } S \subseteq G\}. \end{aligned}$$

Our proof consists first in showing that `SuperFréchet S` contains a maximal element, that is, an ultrafilter on the (infinite) set S , and then in showing that this maximal element does not contain any finite sets. Stated formally in Isabelle, the following goal needs to be established:

$$\begin{aligned} \neg \text{finite } S \implies \exists U \in \text{SuperFréchet } S. \\ \forall G \in \text{SuperFréchet } S. U \subseteq G \longrightarrow U = G \wedge \\ \forall x \in U. \neg \text{finite } x. \end{aligned}$$

5.2.1. Existence of the ultrafilter

We split the main goal above into two parts, and outline in this section how the existence of the ultrafilter is proved. Formally, we need to prove that

$$\begin{aligned} \neg \text{finite } S \implies \exists U \in \text{SuperFréchet } S. \\ \forall G \in \text{SuperFréchet } S. U \subseteq G \longrightarrow U = G. \end{aligned}$$

Applying Zorn’s lemma (as an introduction rule in Isabelle) and with some simplification, this reduces the above to the following new subgoal:

$$\begin{aligned} [|\neg \text{finite } S; c \in \text{chain } (\text{SuperFréchet } S)|] \\ \implies \exists u \in \text{SuperFréchet } S. \forall x \in c. x \subseteq u. \end{aligned}$$

Thus, we now have to show that each chain of `SuperFréchet S` has an upper bound in `SuperFréchet S`. Since the empty set is also a chain, we need to consider the two possibilities for the chain c , as follows.

- 1) $c = \emptyset$. We simply use the fact that $\text{Frechet } S \in \text{Filters } S$ and hence that $\text{Frechet } S \in \text{SuperFrechet } S$ to prove the theorem for this case. (We have noticed that many proofs given in the literature neglect to consider the case where c is the empty chain.)
- 2) $c \neq \emptyset$. This case is trickier. The proof consists in choosing the union of the nonempty chain c , $\bigcup c$, as the upper bound we are looking for. It is trivially true that $x \subseteq \bigcup c$ for all $x \in c$. To bring the proof to conclusion, it just remains to show that $\text{SuperFrechet } S$ is closed under the union of nonempty chains:

$$[[c \neq \emptyset; \neg \text{finite } S; c \in \text{chain } (\text{SuperFrechet } S)]] \\ \implies \bigcup c \in \text{SuperFrechet } S.$$

The proof requires showing that $\bigcup c$ is a filter. Property (F1) for a filter is proved easily using Isabelle’s classical reasoner. In outline, to prove (F2), we choose $x_0 \in \bigcup c$, and $x_1 \in \bigcup c$. Then $x_0 \in G_0$ and $x_1 \in G_1$ for some filters G_1 and G_2 in the chain c . Since c is a chain we have that $G_1 \subseteq G_2$ or $G_2 \subseteq G_1$. If $G_1 \subseteq G_2$ then $x_0, x_1 \in G_2$ and so, by (F1), $x_0 \cap x_1 \in G_2 \subseteq \bigcup c$; the case $G_2 \subseteq G_1$ is proved in a similar way. Finally, we prove that Property (F3) also holds from the properties of chains and unions. We shall omit the details for this last step, since they are easily deduced.

5.2.2. Freeness property

The second part of the main goal consists in proving that the ultrafilter does not contain any finite set. Making use of the statement proved in the previous part, this reduces to solving the following subgoal (that is, deriving a contradiction) in Isabelle:

$$[[U \in \text{SuperFrechet } S; x \in U; \text{finite } x]] \implies \text{False}.$$

To prove this, we first deduce that $(S - x) \in U$ since $\text{finite } (S - (S - x))$ and $\text{Frechet } S \subseteq U$. Hence, since U is closed under set intersection, it follows that $\emptyset = x \cap (S - x) \in U$, which is a contradiction of Property (F1) of the filter. Thus U is free.

This concludes our proof of the existence of a free ultrafilter on any infinite set. This important theorem will be used in the next section to define the hyperreals by considering a special case known as the *weak ultrafilter theorem*.

We have described so far the mathematical foundations set up in Isabelle to enable the definition of the new types of numbers going beyond the traditional number systems. After carrying out constructions up to the reals, proving Zorn’s lemma in Isabelle and developing a theory of filters, we are now ready to apply the so-called *ultrapower* construction to get the *hyperreals*.

6. Ultrapower construction of the hyperreals

Our aim is to construct a linearly ordered field \mathbb{R}^* that contains an isomorphic copy of the reals \mathbb{R} extended with other elements. This new, strictly larger field is known as a *nonstandard* or hyperreal number system and obeys the same field laws as the reals.

As several authors have pointed out [17, 31], the construction of the hyperreals is reminiscent of the construction of the reals from the rationals using equivalence classes induced by Cauchy sequences. In this case, however, we use a free ultrafilter to partition the set of all sequences of real numbers into equivalence classes. The set of these equivalence classes, that is the quotient set, is used to define the new type `hyperreal`, denoting the hyperreal numbers.

```

HYPREAL = REAL + FILTER +

constdefs
  UN :: "nat set set"
  "UN ≡ (εU. u ∈ FreeUltrafilters (UNIV::nat set))"

  (* equivalence relation *)
  hyprel "( (nat ⇒ real) * (nat ⇒ real)) set"
  "hyprel ≡ {p. ∃ r s. p = (r,s) ∧ {n. r n = s n} ∈ UN}"

typedef
  hypreal ≡ "{x::(nat ⇒ real). True}/hyprel"      (Equiv.quotient_def)

instance
  hypreal :: {ord, plus, times}

defs
  hypreal_zero_def  "0hr ≡ Abs_hypreal(hyprel^^{λn::nat. 0r})"
  hypreal_one_def   "1hr ≡ Abs_hypreal(hyprel^^{λn::nat. 1r})"

constdefs
  hypreal_minus :: hypreal ⇒ hypreal
  "- P ≡ Abs_hypreal(⋃X∈Rep_hypreal(P). hyprel^^{λn::nat. - (X n)})"

  (* embedding for the reals *)
  hypreal_of_real :: real ⇒ hypreal
  "hypreal_of_real r ≡ Abs_hypreal(hyprel^^{λn::nat. r})"

  hrinv :: hypreal ⇒ hypreal
  "hrinv P ≡ Abs_hypreal(⋃X∈ Rep_hypreal(P).
    hyprel^^{λn. if X n = 0r then 0r else rinv (X n)})"

defs
  hypreal_add_def
  "P + Q ≡ Abs_hypreal(⋃X∈Rep_hypreal(P). ⋃Y∈Rep_hypreal(Q).
    hyprel^^{λn::nat. X n + Y n})"

  ...

  hypreal_less_def
  "P < (Q::hypreal) ≡ ∃X Y. X∈Rep_hypreal(P) ∧ Y∈Rep_hypreal(Q) ∧
    {n::nat. X n < Y n} ∈ UN"

```

Figure 3: Isabelle/HOL theory for hyperreals.

6.1. Choosing a free ultrafilter

To start the construction, a free ultrafilter $U_{\mathbb{N}}$ is chosen on the set of natural numbers \mathbb{N} . Such an ultrafilter exists, according to the weak ultrafilter theorem.

Theorem 6.1 (Weak ultrafilter theorem: WUF). *There exists a free ultrafilter on \mathbb{N} .*

As can be seen, this is a special case of the ultrafilter theorem's corollary given in Section 5.2. In fact, we have the implications $AC \Rightarrow UFT \Rightarrow WUF$, which are not reversible. Thus, the ultrafilter theorem is strictly weaker than the axiom of choice, and the weak ultrafilter theorem is weaker still. To prove the weak ultrafilter theorem, we show that the set of naturals is not finite by an inductive proof, and then discharge the premise of the ultrafilter theorem's corollary.

This ultrafilter need not be explicitly defined; it does not matter which ultrafilter on \mathbb{N} is used. The set of all free ultrafilters on \mathbb{N} determines a set of isomorphic fields from which we can choose any member to be the set of hyperreal numbers. Thus, in our formalization, we use Hilbert's ε -operator to define $U_{\mathbb{N}}$:

$$U_{\mathbb{N}} \equiv (\varepsilon U. U \in \text{FreeUltrafilters } (\text{UNIV} :: \text{nat set})).$$

In this definition, $(\text{UNIV} :: \text{nat set})$ denotes $\{n :: \text{nat}. \text{True}\}$, the set \mathbb{N} . Higher-order logic provides a typed set theory in which the universal set exists.

Once we have defined $U_{\mathbb{N}}$, its properties that will be used in the proofs involving the hyperreals are established. We give here a list of the theorems that we have proved, many of which follow from the definitions given in the previous sections.

Theorem 6.2. $(\text{UNIV} :: \text{nat set}) \in U_{\mathbb{N}}$.

Theorem 6.3. $\emptyset \notin U_{\mathbb{N}}$.

Theorem 6.4. $X \in U_{\mathbb{N}} \wedge Y \in U_{\mathbb{N}} \implies X \cap Y \in U_{\mathbb{N}}$.

Theorem 6.5. $X \in U_{\mathbb{N}} \wedge X \subseteq Y \implies Y \in U_{\mathbb{N}}$.

Theorem 6.6. $X \in U_{\mathbb{N}} \implies \neg \text{finite } X$.

Theorem 6.7. $X \in U_{\mathbb{N}} \iff -X \notin U_{\mathbb{N}}$.

Theorem 6.8. $\{n. P(n)\} \in U_{\mathbb{N}} \implies \exists n. P(n)$.

Theorem 6.9. $X \cup Y \in U_{\mathbb{N}} \implies X \in U_{\mathbb{N}} \vee Y \in U_{\mathbb{N}}$.

6.2. Equality

Using $U_{\mathbb{N}}$, the hyperreals are constructed by considering the set of all sequences of real numbers indexed by \mathbb{N} and defining the following equivalence relation on this set.

Definition 6.1. Given two sequences of real numbers $\langle r_n \rangle$ and $\langle s_n \rangle$,

$$\langle r_n \rangle \sim_{U_{\mathbb{N}}} \langle s_n \rangle \iff \{n \in \mathbb{N} \mid r_n = s_n\} \in U_{\mathbb{N}}.$$

The sequences $\langle r_n \rangle$ and $\langle s_n \rangle$ are sometimes said to be equal *almost everywhere* (a.e.). This terminology is used to mean that the entries of a sequence determine some set in the ultrafilter $U_{\mathbb{N}}$.

Figure 3 shows Isabelle’s theory HYPREAL, in which the new type `hypreal` is introduced using the definition above. The relation `hyprel` denotes $\sim_{U_{\mathbb{N}}}$ in the theory:

$$\text{hyprel} \equiv \{p. \exists r s. p = (r, s) \wedge \{n. r(n) = s(n)\} \in U_{\mathbb{N}}\}.$$

The first property that we prove is that `hyprel` is an equivalence relation.

Proposition 6.1. *The relation $\sim_{U_{\mathbb{N}}}$ is an equivalence relation.*

Proof. Let $\langle a_n \rangle, \langle b_n \rangle, \langle c_n \rangle$ be sequences of real numbers.

- *Reflexivity:* since $\mathbb{N} \in U_{\mathbb{N}}$, we have $\langle a_n \rangle \sim_{U_{\mathbb{N}}} \langle a_n \rangle$ and thus $\sim_{U_{\mathbb{N}}}$ is reflexive.
- *Symmetry:* if $\langle a_n \rangle \sim_{U_{\mathbb{N}}} \langle b_n \rangle$ then, by symmetry of equality, $\langle b_n \rangle \sim_{U_{\mathbb{N}}} \langle a_n \rangle$, implying that $\sim_{U_{\mathbb{N}}}$ is symmetric.
- *Transitivity:* now, given $\langle a_n \rangle \sim_{U_{\mathbb{N}}} \langle b_n \rangle$ and $\langle b_n \rangle \sim_{U_{\mathbb{N}}} \langle c_n \rangle$, let $A = \{n \in \mathbb{N} \mid a_n = b_n\}$ and $B = \{n \in \mathbb{N} \mid b_n = c_n\}$, and $C = \{n \in \mathbb{N} \mid a_n = c_n\}$; then $A \cap B \subseteq C$. Since $A, B \in U_{\mathbb{N}}$, it follows that $A \cap B \in U_{\mathbb{N}}$ since $U_{\mathbb{N}}$ is \cap -closed, and hence $C \in U_{\mathbb{N}}$ since $U_{\mathbb{N}}$ is also \subseteq -closed. Therefore, $\langle a_n \rangle \sim_{U_{\mathbb{N}}} \langle c_n \rangle$.

□

6.3. Defining operations on the hyperreals

Arithmetic operations on the new type, that is on the equivalence classes, are usually defined in terms of the pointwise operations on the sequences. Let $[\langle X_n \rangle]$ denote the equivalence class containing $\langle X_n \rangle$. Addition, for example, is defined by

$$[\langle X_n \rangle] + [\langle Y_n \rangle] \equiv [\langle X_n + Y_n \rangle]. \tag{3}$$

In Isabelle, however, using the abstraction and representation functions, we define addition on hyperreals P and Q as follows:

$$P + Q \equiv \text{Abs_hypreal} \left(\bigcup X \in \text{Rep_hypreal}(P). \bigcup Y \in \text{Rep_hypreal}(Q). \text{hyprel}^{\wedge\wedge} \{\lambda n. X n + Y n\} \right).$$

Then we prove equation (3) above as a theorem. It can then be supplied to the simplifier for use in many of the proofs. In Isabelle, equation (3) takes the following form:

$$\begin{aligned} &\text{Abs_hypreal} (\text{hyprel}^{\wedge\wedge} \{\lambda n. X n\}) + \text{Abs_hypreal} (\text{hyprel}^{\wedge\wedge} \{\lambda n. Y n\}) \\ &= \text{Abs_hypreal} (\text{hyprel}^{\wedge\wedge} \{\lambda n. X n + Y n\}). \end{aligned} \tag{4}$$

Properties such as commutativity and associativity follow straightforwardly from the corresponding properties of the reals. We can similarly prove $0_{\text{hr}} + P = P$ when 0_{hr} is defined as shown in Figure 3. Multiplication is defined in a similar way to addition. Associativity, commutativity, and distributivity of multiplication are all directly inherited from the reals, and are easily proved.

6.4. Ordering

The ordering relation on the hyperreals is defined as follows:

$$P < Q \equiv \exists X \in \text{Rep_hypreal } P. \\ \exists Y \in \text{Rep_hypreal } Q. \{n. X n < Y n\} \in U_{\mathbb{N}}.$$

We prove the following simplification theorem expressing the order relation in terms of equivalence classes of sequences of real numbers. A hyperreal $[\langle X_n \rangle]$ is less than a hyperreal $[\langle Y_n \rangle]$ if and only if X_n is less than Y_n *almost everywhere*:

$$\text{Abs_hypreal } (\text{hyprel}^{\wedge\wedge} \{X n\}) < \text{Abs_hypreal } (\text{hyprel}^{\wedge\wedge} \{Y n\}) \\ \iff \{n. X n < Y n\} \in U_{\mathbb{N}}.$$

Also, the system of hyperreal numbers generated by the free ultrafilter is a totally ordered field. To show this, we first prove that the ordering relation is total. This proof is relatively simple and follows from the fact that, given any two hyperreals $[\langle x_n \rangle]$ and $[\langle y_n \rangle]$, either they are equal, leading to

$$\{n \in \mathbb{N} \mid x_n = y_n\} \in U_{\mathbb{N}}$$

or else, by the complement property of the ultrafilter as given in Section 6.1, we find that

$$\{n \in \mathbb{N} \mid x_n \neq y_n\} \in U_{\mathbb{N}}.$$

In the second case, since the reals are totally ordered, we have to consider the sets $\{n \in \mathbb{N} \mid x_n < y_n\}$ and $\{n \in \mathbb{N} \mid y_n < x_n\}$. We know that only one of these can belong to the free ultrafilter $U_{\mathbb{N}}$ (since otherwise, closure of $U_{\mathbb{N}}$ under intersection would entail that $\emptyset \in U_{\mathbb{N}}$, which contradicts property (F1) of the filter).

6.5. Multiplicative inverse

To show that \mathbb{R}^* is a field, we need only prove that each non-zero element $[\langle X_n \rangle] \in \mathbb{R}^*$ has a multiplicative inverse. For any non-zero element, we have

$$\{n \in \mathbb{N} \mid X_n = 0\} \notin U_{\mathbb{N}}$$

and therefore, once more by the complement property of $U_{\mathbb{N}}$,

$$\{n \in \mathbb{N} \mid X_n \neq 0\} \in U_{\mathbb{N}}.$$

Therefore, define $Y_n = 1/X_n$ for each value of n for which $X_n \neq 0$, and set $Y_n = 0$ otherwise. Then the set $\{n \in \mathbb{N}. X_n \cdot Y_n = 1\} \in U_{\mathbb{N}}$, so that $[\langle X_n \rangle] \cdot [\langle Y_n \rangle] = [\langle 1 \rangle]$. This motivates the following definition, in Isabelle, for the inverse function `hrinv`:

$$\text{hrinv } P \equiv \text{Abs_hypreal } (\bigcup X \in \text{Rep_hypreal}(P). \\ \text{hyprel}^{\wedge\wedge} \{\lambda n. \text{if } X n = 0r \text{ then } 0r \text{ else } r \text{inv } (X n)\}).$$

It is easily proved that for all non-zero x , `hrinv` $x \cdot x = 1r$ as required. A few points worth mentioning are that `hrinv` x stands for the more conventional notation x^{-1} when x is a hyperreal; the inverse function for the reals is itself denoted by `rinv`, while `0r` and `1r` are defined as the zero and one respectively of the real field. Once again, for simplification purposes, we prove the useful theorem about inverse involving the equivalence classes of real sequences:

$$\text{hrinv } (\text{Abs_hypreal } (\text{hyprel}^{\wedge\wedge} \{X n\})) \iff \\ \text{Abs_hypreal } (\text{hyprel}^{\wedge\wedge} \{\text{if } X n = 0r \text{ then } 0r \text{ else } r \text{inv } (X n)\}).$$

We have shown in the discussion above that \mathbb{R}^* is a totally ordered field. The next important step is to show that \mathbb{R}^* contains a proper subfield that is isomorphic to the reals \mathbb{R} .

7. Structure of the hyperreal number line

In this section, we continue our investigation by introducing and defining the various elements that make up the new totally ordered field, which we show to be a proper extension of the reals. We also define a number of concepts that follow from this classification of the elements of \mathbb{R}^* .

7.1. Embedding the reals

Since our free ultrafilter has been fixed, we have effectively restricted our attention to one particular totally ordered field \mathbb{R}^* , though as we mentioned previously, there are infinitely many distinct but isomorphic number systems. We now embed the reals in our hyperreals by defining a map `hypreal_of_real :: real \Rightarrow hypreal` in Isabelle. This embedding is defined by

$$\text{hypreal_of_real } r = [\langle r, r, r, \dots \rangle]$$

and is expressed in Isabelle as

$$\text{hypreal_of_real } r \equiv \text{Abs_hypreal } (\text{hyprel}^{\wedge\wedge} \{ \lambda n :: \text{nat}. r \}).$$

In what follows, any embedded real r will be denoted by \tilde{r} unless the embedding function `hypreal_of_real` is used explicitly. Thus, the additive identity element `0hr` and the multiplicative identity element `1hr` of the hyperreals are the explicit images of the real numbers zero (`0r`) and one (`1r`) respectively under the embedding. To show that `hypreal_of_real` maps \mathbb{R} to a proper subfield of \mathbb{R}^* , we first define the following hyperreal number:

$$\omega \equiv \text{Abs_hypreal } (\text{hyprel}^{\wedge\wedge} \{ \lambda n :: \text{nat}. \text{real_of_nat } n \})$$

where `real_of_nat :: nat \Rightarrow real` maps its natural argument n to the real $n + 1$. For clarity, we omit the details of the various intermediate embeddings (`nat \Rightarrow pnat`, `pnat \Rightarrow prat`, `prat \Rightarrow preal`, and so on) required for defining `real_of_nat`, though we do need to prove their various properties (for example, that they are injective and order-preserving) explicitly in Isabelle. This sort of detail is not usually mentioned in textbooks, where it is assumed that one can define a map in one step.

We can now exhibit a member of \mathbb{R}^* that is not equal to any real number, since there is no r such that $\tilde{r} = \omega$. This is because the set on which $\langle r, r, r, \dots \rangle$ and $\langle 1, 2, 3, \dots \rangle$ coincide can consist of at most one element. Hence, by the definition of ultrafilter $U_{\mathbb{N}}$, the two sequences cannot be equivalent since no finite set can belong to $U_{\mathbb{N}}$. In fact, as we shall see shortly, $\tilde{r} < \omega$ for any real number r ; that is, ω is a so-called *infinite* number. Similarly, $\epsilon = \omega^{-1} = [\langle 1, \frac{1}{2}, \frac{1}{3}, \dots \rangle]$ is an *infinitesimal*.

We will call all members of \mathbb{R}^* that are images of the reals, the *standard* elements of \mathbb{R}^* . We then define the set of standard reals `SReal` in the theory NSA as follows,

$$\text{SReal} \equiv \text{range } (\text{hypreal_of_real})$$

where

$$\text{range } f = \{y. \exists x. y = f x\}.$$

We can now view `SReal` as the real numbers embedded in \mathbb{R}^* ; that is, as a sub-ordered field, if we agree to identify each real number r with the corresponding standard element \tilde{r} of \mathbb{R}^* . We then have that \mathbb{R}^* is an extension or enlargement of \mathbb{R} . We shall come across the general concept of set extensions later in the paper (see Section 10).

7.2. Properties of nonstandard numbers

We have exhibited in the previous section a hyperreal, ω , that does not belong to `SReal`. There are infinitely many of these so-called *nonstandard* hyperreal numbers. They can be classified into various sets that include, for example, infinitesimals and the infinite numbers. We start this section with a preamble, where the absolute value function for the hyperreals is introduced. This function is needed in order to define the various types of numbers found in our theory.

The definitions of infinitesimal, finite, and infinite numbers use the absolute value function. This function, which we also defined on the reals, needs to be extended to the hyperreal numbers. The definition that we use is analogous to that used for the reals. Using the if-then-else construct of Isabelle/HOL, we have

$$\text{hrabs } x \equiv \text{if } 0 \text{hr} \leq x \text{ then } x \text{ else } -x.$$

In fact, an alternative definition exists in which the (real) absolute value function is simply applied pointwise to an equivalence class representative in \mathbb{R}^* . In Isabelle, with `rabs` denoting the absolute value function for the reals, this takes the form of the following theorem:

$$\begin{aligned} \text{hrabs } (\text{Abs_hypreal } (\text{hyprel}^{\wedge}\{X\})) = \\ \text{Abs_hypreal } (\text{hyprel}^{\wedge}\{\lambda n. \text{rabs } (X n)\}). \end{aligned}$$

This result, taken in conjunction with the definitions of the operations such as addition, multiplication and reciprocal, hints at a general technique in which functions can be defined on the hyperreals through *extensions* of the analogous ones defined on the reals using our free ultrafilter $U_{\mathbb{N}}$. We examine this notion of extension later in this work.

The intuitive notion of an infinitesimal number can now be formally defined. Sets of finite and infinite numbers are also formally introduced.

Definition 7.1. An element x of \mathbb{R}^* is said to be an *infinitesimal* if and only if for every positive standard real number r we have $|x| < r$. It is *finite* if and only if for some standard real number r we have $|x| < r$, and *infinite* if and only if for every standard real number r we have $r < |x|$.

In the literature, the definition will often just say that an infinitesimal is less in magnitude than any positive (standard) *real* number. Here, since we have different types, it becomes explicit that such a definition is actually referring to the standard copy in \mathbb{R}^* . This leads to the following definition in Isabelle for the set of `Infinitesimal`:

$$\begin{aligned} \text{Infinitesimal} &:: \text{hypreal set} \\ \text{Infinitesimal} &\equiv \{x. \forall r \in \text{SReal}. 0 \text{hr} < r \longrightarrow \text{hrabs } x < r\}. \end{aligned}$$

This definition can be considered as a high-level one. Indeed, it is possible to define the set of infinitesimals by going down to the level of our free ultrafilter $U_{\mathbb{N}}$ itself. We thus prove the next theorem, which turns out to be useful when supplied to Isabelle’s simplifier

in cases where one wants to deal with real sequences rather than infinitesimals:

$$(x \in \text{Infinitesimal}) \iff (\exists X \in \text{Rep_hypreal } x. \forall u. 0r < u \longrightarrow \{n. \text{rabs } (X n) < u\} \in U_{\mathbb{N}}).$$

We give below the definitions for the sets `Finite` and `Infinite` of finite and infinite numbers respectively, as declared in Isabelle, and the equivalent theorems derived in terms of the free ultrafilter:

```
Finite :: hypreal set
Finite ≡ {x. ∃r ∈ SReal. hrabs x < r};

(x ∈ Finite) ≡ (∃X ∈ Rep_hypreal x.
    ∃u. {n. rabs (X n) < u} ∈ Uℕ);

Infinite :: hypreal set
Infinite ≡ {x. ∀r ∈ SReal. r < hrabs x};

(x ∈ Infinite) ≡ (∃X ∈ Rep_hypreal x.
    ∀u. {n. u < rabs (X n)} ∈ Uℕ).
```

We can view the low-level theorems as lemmas that enable us to translate properties involving the hyperreals into those depending on the ultrafilter. This is useful in our mechanization when we deal with real functions and their extensions.

An important point, highlighted through the definition of infinite and infinitesimal numbers, and already mentioned in Section 3, is that the set of hyperreal numbers is non-Archimedean. This is because not every bounded subset of \mathbb{R}^* has a least upper bound or greatest lower bound. For example, the set of infinite numbers is bounded below by any finite number, but has no greatest lower bound.

7.3. On infinitesimal, finite and infinite numbers

We have proved various properties of infinitesimal, finite and infinite numbers. A few of the theorems are listed below.

Theorem 7.1. *The set `Finite` of finite elements is a subring of \mathbb{R}^* ; that is, sums, differences, and products of finite elements are finite.*

Theorem 7.2. *The set `Infinitesimal` of infinitesimals is also a subring of \mathbb{R}^* .*

Theorem 7.3. *The set `Infinitesimal` is an ideal in `Finite`; that is, the product of an infinitesimal and a finite number is infinitesimal.*

Theorem 7.4. *Element x is infinite if and only if $\text{hrinv } x$ is infinitesimal for all non-zero x .*

The hyperreal number ω defined in Section 7.1 is a member of `Infinite`: for any given real number x , for all sufficiently large values of n , we have $x < n$. The infinitesimal number ϵ defined by the equivalence class containing the sequence $\langle 1/n \rangle$ is a member of

Infinitesimal since for any given x , for all sufficiently large value of n , we have $0 < 1/n < x$. We have also proved that ω is the multiplicative inverse of ϵ , since

$$\begin{aligned} \omega \cdot \epsilon &= [(1, 2, 3, \dots)] \cdot [(1, 1/2, 1/3, \dots)] \\ &= [(1 \cdot 1, 2 \cdot 1/2, 3 \cdot 1/3, \dots)] \\ &= [(1, 1, 1, \dots)] \\ &= 1\text{hr}. \end{aligned}$$

We next introduce an important equivalence relation that will be extremely useful to our mechanization.

Definition 7.2. Two hyperreal numbers x and y are said to be *infinitely close*, $x \approx y$, if and only if their difference $x - y$ is infinitesimal.

It is easily proved that x is an infinitesimal if and only if $x \approx 0$. To show that \approx is an equivalence relation is trivial. In addition, we prove the following theorems (amongst others).

Theorem 7.5. $a \approx b \wedge c \approx d \implies a + c \approx b + d$.

Theorem 7.6. $(a + b \approx a + c) \iff b \approx c$.

Theorem 7.7. $a \approx b \wedge c \in \text{Finite} \implies a \cdot c \approx b \cdot c$.

Theorem 7.8. $a \approx b \wedge c \approx d \wedge b \in \text{Finite} \wedge c \in \text{Finite} \implies a \cdot c \approx b \cdot d$.

Theorem 7.9. $a \in \text{Finite} \wedge a \approx b \implies b \in \text{Finite}$.

Theorem 7.10. $a \in \text{SReal} \wedge a \neq 0\text{hr} \implies (a \cdot x \approx a \cdot y) = (x \approx y)$.

Theorem 7.11. $x \in \text{SReal} \wedge y \in \text{SReal} \implies (x \approx y) = (x = y)$.

Theorem 7.12. $x \approx y \wedge y \in \text{Finite} - \text{Infinitesimal} \implies \text{hrinv } x \approx \text{hrinv } y$.

Theorem 7.13. $x \approx y \implies \text{hrabs } x \approx \text{hrabs } y$.

We continue in the next section with another basic fact about the structure of \mathbb{R}^* , which defines a function from the set of finite numbers onto the reals.

7.4. The standard part theorem

The *standard part* of a finite nonstandard number is defined to be the unique real infinitely close to it. We use Hilbert's choice operator, ϵ , to express this in Isabelle:

$$\text{st } x \equiv (\epsilon r. r \in \text{SReal} \wedge r \approx x).$$

We now prove the existence and uniqueness of the standard part. Existence needs to be demonstrated in any case whenever Hilbert's operator is used.

Proposition 7.1. *Let x be a finite hyperreal number. Then, there exists a unique standard real number r such that $r \approx x$.*

Proof. Let $A = \{y \in \mathbb{R} \mid y \leq x\}$. Since x is finite, A is nonempty and is bounded above. Let r be the least upper bound of A . For any real $\epsilon > 0$, $r - \epsilon \in A$ and $r + \epsilon \notin A$, and thus $r - \epsilon \leq x < r + \epsilon$. So $|r - x| \leq \epsilon$, from which it follows that $r \approx x$.

To show uniqueness, suppose that there exists a real number s such that $s \approx x$. Then, since \approx is transitive, $s \approx r$ and so $r - s \approx 0$. But $r - s$ is real, so $r - s = 0$ and $r = s$. \square

The proof just given glosses over many of the details that need to be satisfied for mechanization. The completeness of the reals, and hence of the embedded reals, is needed in the form of the *supremum property*, which ensures that any nonempty set of reals that is bounded above has a least upper bound. We first proved the property for the positive real numbers (`preal`) and then extended it to the reals (`real`). Now, since we are dealing with the hyperreals and identifying the reals with the proper subfield of \mathbb{R}^* which is isomorphic to \mathbb{R} , we have to transfer this theorem explicitly to the isomorphic copy of \mathbb{R} , namely `SReal`.

Once the existence of the standard part has been proved, we prove various of their properties: for any $x, y \in \text{Finite}$, we have

$$\begin{aligned} x \approx y &\iff \text{st } x = \text{st } y; \\ x \approx \text{st } x; \\ \text{st } (x + y) &= \text{st } x + \text{st } y; \\ \text{st } (x \cdot y) &= \text{st } x \cdot \text{st } y; \\ \text{if } \text{st } y \neq 0 &\text{hr then } \text{st } (x \cdot \text{hrinv } (y)) = \text{st } x \cdot \text{hrinv } (\text{st } y); \\ \text{st } (\text{st } x) &= \text{st } x; \\ \text{st } (\text{hrabs } x) &= \text{hrabs } (\text{st } x). \end{aligned}$$

From some of these theorems, we can see that the map preserves algebraic structure. The standard part function can be defined in other ways. For example, it corresponds to the order homomorphism of `Finite` with kernel `Infinitesimal` onto \mathbb{R} [32]. The standard part is an important concept that can be used when formulating the nonstandard definition for the limit of a sequence of reals, and also when defining the *slope* of a real function at a real point.

8. The hypernatural numbers

We can construct a set of numbers \mathbb{N}^* that contains both finite elements, identifiable with the ordinary natural numbers, and infinite numbers greater than all natural numbers. This discrete set is known as the *hypernaturals*. They will be needed in the nonstandard formalization of real sequences and series in the next part of this mechanization.

The construction of the hypernaturals in Isabelle is analogous to that of the hyperreals: we use the same free ultrafilter $U_{\mathbb{N}}$ but replace sequences of reals by sequences of natural numbers. Thus, \mathbb{N}^* is now characterized explicitly as the set of equivalence classes $[\langle m_n \rangle]$ determined by sequences m_n of natural numbers. The new equivalence relation on sequences is denoted by `hypnatrel` in Isabelle. In what follows, we make some observations on the construction and properties that apply to members of \mathbb{N}^* . These are interesting in their own right, but also in view of the applications to mechanization of analysis using nonstandard methods.

We define an embedding function that identifies each natural number m with the hypernatural number determined by the constant sequence $\langle m, m, \dots, m \rangle$. In Isabelle, we define

the function `hypnat_of_nat :: nat ⇒ hypnat`:

$$\text{hypnat_of_nat } m \equiv \text{Abs_hypnat } (\text{hypnatrel}^{\wedge\wedge} \{\lambda n::\text{nat}. m\}).$$

Using the map `hypnat_of_nat`, we easily define the set `SHNat` of *standard* natural numbers embedded in \mathbb{N}^* :

$$\text{SHNat} \equiv \text{range } (\text{hypnat_of_nat}).$$

In what follows, a natural number n embedded in the hypernaturals will also be denoted by \bar{n} in some cases.

8.1. Infinite hypernaturals

We define a hypernatural Ω denoting $[\langle n \rangle] = [\langle 0, 1, 2, \dots \rangle]$ by

$$\Omega \equiv \text{Abs_hypnat } (\text{hypnatrel}^{\wedge\wedge} \{\lambda n::\text{nat}. n\}).$$

We prove that for any embedded natural number $n \in \text{SHNat}$, $\Omega \neq n$ meaning that \mathbb{N}^* properly includes \mathbb{N} . This motivates the following definition for the set of non-standard hypernaturals:

$$\text{HNatInfinite} \equiv - \text{SHNat}$$

where ‘ $-$ ’ denotes set complement in Isabelle. To establish that the only nonstandard hypernaturals are the infinite ones, we prove the following equivalence theorem:

$$\text{HNatInfinite} \iff \{N. \forall n \in \text{SHNat}. n < N\}.$$

Thus, \mathbb{N}^* consists of the finite standard copies of the ordinary natural numbers and of the infinite hypernatural numbers only.

8.2. Properties of the hypernaturals

Some of the properties proved for the hypernatural numbers are as follows.

- 1) \mathbb{N}^* is a discrete subset of \mathbb{R}^* .
- 2) \mathbb{N}^* is closed under addition and multiplication.
- 3) Every infinite number has an immediate predecessor, which is also infinite.

The first property can be proved either by defining directly an embedding function from the hypernaturals to the hyperreals, or by taking the nonstandard extension of the set of natural numbers (embedded in the reals).

An important observation, following from the third property above, is that the non-empty set of infinite hypernatural numbers, `HNatInfinite`, does not have a least element. Thus, the well-ordering property of the natural numbers does not extend to the hypernaturals. This observation shows that, though most properties of the natural numbers are transferred to the hypernaturals, there are important exceptions. It will be seen in our subsequent exposition that properties such as the one above and the Archimedean property extend only to *special* subsets of the hypernaturals and hyperreals respectively. In what follows, we review the development of concepts from real analysis in Isabelle.

9. Mechanized infinitesimal calculus

Classical or standard analysis is mostly concerned with the study of the real numbers, and with the properties of functions defined on them. We shall now describe the use of

the hyperreals as tools for mathematical analysis. Through the existence of infinitesimals, finite, and infinite numbers, NSA provides us with a rich structure which we use to formalize alternative treatments of topics in classical analysis. Such treatments are valuable, not only for the additional light that they cast on analysis, but also for the simplification they bring. As will be seen, the mechanization of analysis can benefit directly from this simplification, since difficult instantiation steps in proofs are eliminated in many cases. We start by showing how functions defined over the reals and naturals can be systematically extended to the hyperreals and hypernaturals, respectively. These notions are crucial to nonstandard real analysis. We then proceed to develop some elementary analysis that will make use of the new classes of numbers, the infinitely close relation, and other notions induced on them.

10. Extending a relation to the hyperreals

There are systematic methods through which functions defined on the reals are extended to the hyperreals. This process of extending a relation from \mathbb{R} to \mathbb{R}^* is known as the **-transform* [17].

10.1. Internal sets and nonstandard extensions

Many properties of the reals, suitably reinterpreted, can be transferred to the hyperreal number system. For example, we have seen that \mathbb{R}^* , like \mathbb{R} , is a totally ordered field. Also, just as \mathbb{R} contains the natural numbers \mathbb{N} as a discrete subset with its own characteristic properties, \mathbb{R}^* contains the hypernaturals \mathbb{N}^* as a corresponding discrete subset with analogous properties. Moreover, subsets \mathbb{Z}^* (the hyperintegers) and \mathbb{Q}^* of \mathbb{R}^* exhibit relations to \mathbb{N}^* similar to those that \mathbb{Z} and \mathbb{Q} bear to \mathbb{N} in \mathbb{R} .

However, there are properties of \mathbb{R} that do not transfer to \mathbb{R}^* . This is the case for the fundamental supremum property of the reals. It is easy to see that this upper bound property does not necessarily hold by considering, for example, the set \mathbb{R} itself, which we regard as embedded into the hyperreals (that is, the set SReal from Section 7.1). This is a non-empty set which is bounded above (by any of the infinite numbers in \mathbb{R}^*) but does not have a least upper bound in \mathbb{R}^* .

Theorem 10.1. *The set $\mathbb{R} \subseteq \mathbb{R}^*$ does not have a least upper bound in \mathbb{R}^* .*

Proof. Suppose that r is the least upper bound of \mathbb{R} . Then it follows that r is infinite, since it is an upper bound. But as $r \in \text{Infinite}$, it follows that $r - 1 \in \text{Infinite}$, so $r - 1$ is a smaller upper bound, which is a contradiction. \square

We now introduce an important refinement that classifies subsets of \mathbb{R}^* into two types: *internal* and *external* subsets [17, 16]. With this done, we shall be able to prove the following statement, for example, about the supremum property for the hyperreals.

Every non-empty *internal* subset of \mathbb{R}^* which has an upper bound in \mathbb{R}^* has a least upper bound in \mathbb{R}^* .

Definition 10.1. Let A_n , for $n \in \mathbb{N}$, be any sequence of sets of real numbers. This sequence determines a certain set A of hyperreals according to the following rule. The hyperreal number $x = [(X_n)]$ is a member of set $A \subseteq \mathbb{R}^*$ if and only if the set $\{n \in \mathbb{N}. X_n \in A_n\}$ belongs to the ultrafilter $U_{\mathbb{N}}$.

This definition is analogous to the one we used to define hyperreals in terms of sequences of reals. The sequences of sets of real numbers can then be used to define the so-called *internal sets* of hyperreals. In Isabelle, we have the following declaration and definition for an internal set:

$$\begin{aligned} *sn* &:: (\text{nat} \Rightarrow \text{real set}) \Rightarrow \text{hypreal set} \\ *sn* A &\equiv \{x. \forall X \in \text{Rep_hypreal}(x). \{n. X(n) \in A(n)\} \in U_{\mathbb{N}}\}. \end{aligned}$$

We are particularly interested in the special case when the sequence is constant; that is, $A_n = A$ for all (or almost all) n . The internal set determined by such a sequence is called the *nonstandard extension* of A and, since this is the actual property that will be used more often in the course of our mechanization, it is defined explicitly:

$$\begin{aligned} *s* &:: \text{real set} \Rightarrow \text{hypreal set} \\ *s* A &\equiv \{x. \forall X \in \text{Rep_hypreal}(x). \{n. X(n) \in A\} \in U_{\mathbb{N}}\}. \end{aligned}$$

Thus, it follows that $*s* A = *sn* (\lambda n. A)$. In the literature, the nonstandard extension of a set A is usually denoted by A^* . We shall make use of this conventional mathematical notation as well. However, the actual Isabelle/HOL notation ($*s* A$) will also be used in many cases, especially to show how a particular concept is expressed in the theorem-prover.

It can be shown that any non-empty, internal subset of \mathbb{R}^* has the supremum property though the proof will not be given here [16]. In fact, for any subset of S of \mathbb{R}^* that fails to have a least upper bound, one can infer that it is not internal. Any subset of hyperreals that is not internal is called *external*.

The process of extending a set of real numbers to a set of hyperreals has shown an example of the $*$ -transform at work. In general, this transformation procedure can be applied to any n -ary relation on the reals, extending it to an n -ary relation on the hyperreals. This is done using the rule that P holds on an n -tuple in $(\mathbb{R}^*)^n$ if the index set where P holds on the representative real n -tuple sequence is in the chosen free ultrafilter. More instances of $*$ -transforms will be met when nonstandard extensions of functions are introduced.

10.2. Properties of extended sets

Various properties of nonstandard extensions of sets of real numbers can now be derived. The first result to be proved (in one step, using Isabelle’s automatic tactic) is that \mathbb{R}^* is the nonstandard extension of \mathbb{R} . The nonstandard extensions of sets of reals will, in general, be different from the original set. The exception occurs for finite sets, since then the extension function simply degenerates to the embedding function. This is confirmed by the following theorem, where the symbol “ denotes the image operator for relations:

$$\text{finite } A \implies *s* A = \text{hypreal_of_real} \text{ “} A.$$

If the set A is infinite, however, then we prove that A^* contains elements that are not standard copies of the members of A . This leads to the following theorem relating the embedding of a set of real numbers A to its nonstandard extension A^* :

$$\text{hypreal_of_real} \text{ “} A \subseteq *s* A.$$

The nonstandard extension provides us with a new set that is an enlargement of A . Thus, the enlargement of \mathbb{R} yields a new set that contains infinitesimals and infinite elements that have no counterparts in the real number system. A number of useful results involving boolean

operations on nonstandard extensions of sets are proved:

$$(*s* \emptyset) = \emptyset; \tag{5}$$

$$*s* (-A) = -(*s* A); \tag{6}$$

$$A \subseteq B \implies (*s* A) \subseteq (*s* B); \tag{7}$$

$$*s* (A \cup B) = (*s* A) \cup (*s* B); \tag{8}$$

$$*s* (A \cap B) = (*s* A) \cap (*s* B); \tag{9}$$

$$\forall n. X(n) \notin M \implies \text{Abs_hypreal}(\text{hypreal}^{\wedge}\{X\}) \notin (*s* M). \tag{10}$$

The proofs of these theorems all follow from the properties of the free ultrafilter (see Section 6). For example, property (5) holds because no filter contains the empty set. Property (9) holds because filters are closed under the \cap and \subseteq operations. Proving properties (8) and (6) needs the fact that for any subset A of \mathbb{N} , either A or \bar{A} belongs to the ultrafilter. The proofs are all straightforwardly carried through with the help of Isabelle’s automatic tactic.

10.3. Internal functions and nonstandard extensions

Given a *standard* function that takes real arguments, we want to be able to define an analogous one that will also take *nonstandard* arguments. This leads to the notions of internal functions, and to nonstandard extensions. These concepts are crucial, as they will enable the formulation of familiar constructs in analysis using nonstandard definitions. Also, they give a systematic way of extending any function over the reals to one over the hyperreals. We give the definition for the case dealing with internal functions of one real variable [16].

Definition 10.2. Let $\langle F_n \rangle$ be any sequence of standard functions from \mathbb{R} to \mathbb{R} . This sequence determines an *internal function* $f \equiv [\langle F_n \rangle]$ from \mathbb{R}^* to \mathbb{R}^* according to the rule $x = [\langle X_n \rangle] \in \mathbb{R}^*$ maps into $y = [\langle Y_n \rangle] = f(x) \in \mathbb{R}^*$ if and only if $\{n \in \mathbb{N}. Y_n = F_n(X_n)\} \in U_{\mathbb{N}}$.

Expressed in Isabelle, we have this rather more concise definition for the internal function:

$$\begin{aligned} *fn* &:: (\text{nat} \Rightarrow (\text{real} \Rightarrow \text{real})) \Rightarrow \text{hypreal} \Rightarrow \text{hypreal} \\ *fn* F x &\equiv \text{Abs_hypreal} (\bigcup X \in \text{Rep_hypreal}(x). \text{hypreal}^{\wedge}\{\lambda n. (F n)(X n)\}). \end{aligned}$$

Thus, according to this definition, with F and x defined as above, the value of the internal function $(*fn* F)$ at x is given by

$$(*fn* F) x = [\langle F_1(X_1), F_2(X_2), \dots, F_n(X_n), \dots \rangle].$$

Of interest here, as well, is the special type of internal function known as the *nonstandard* extension of a standard function F . The nonstandard extension is obtained by having a constant sequence of functions; that is, one for which $F_n = F$ for (almost) all n . We define the special case in Isabelle as follows:

$$\begin{aligned} *f* &:: (\text{real} \Rightarrow \text{real}) \Rightarrow \text{hypreal} \Rightarrow \text{hypreal} \\ *f* F x &\equiv \text{Abs_hypreal} (\bigcup X \in \text{Rep_hypreal}(x). \text{hypreal}^{\wedge}\{\lambda n. F(X n)\}). \end{aligned}$$

We will denote the nonstandard extension of a given function either by f^* or by the equivalent Isabelle notation $(*f* f)$. Referring back to the construction of the hyperreals in Isabelle as described in Section 6.3, the definitions given for the field operations on them can all be viewed as nonstandard extensions of the analogous operations on the reals (for

example, addition on the hyperreals is actually $+$). We also note that our definition of nonstandard extension corresponds to Keisler’s *function axiom*, which states that ‘for each real function f of n variables there is a corresponding function f^* of n variables, called the natural extension of f ’ [20].

10.4. Properties of extended functions

We prove, as we did for set extensions, a number of useful properties about nonstandard extensions of functions. One of the first and most useful simplification theorems shows that the nonstandard extension of a function f^* is equivalent to applying f elementwise to an equivalence class representative in \mathbb{R}^* :

$$(*f* f) (\text{Abs_hypreal } (\text{hyprel } ^\wedge \{\lambda n. Xn\})) = (\text{Abs_hypreal } (\text{hyprel } ^\wedge \{\lambda n. f(Xn)\})).$$

This enables us to prove theorems about nonstandard functions by using the properties of the corresponding standard real function, the reals, and the free ultrafilter. We then prove theorems about the preservation of rules across the $*$ -transformation and other properties. Some of these Isabelle theorems are listed next. Most of the proofs are mechanized in two steps or fewer with the help of Isabelle’s automatic tactic *auto_tac*; the tactic is supplied with simplification rules such as the theorem above, and others about addition, multiplication and other operations. (We recall that \tilde{r} stands for the image of real number r in \mathbb{R}^* , as described in Section 7.1.)

$$(*f* (\lambda y. f y + g y)) x = (*f* f) x + (*f* g) x \tag{11}$$

$$(*f* (\lambda y. f y \cdot g y)) x = (*f* f) x \cdot (*f* g) x \tag{12}$$

$$(*f* (f \circ g)) = (*f* f) \circ (*f* g) \tag{13}$$

$$(*f* \lambda y. k) x = \tilde{k} \tag{14}$$

$$(*f* (\lambda y. - f y)) x = - (*f* f) x \tag{15}$$

$$(*f* (\lambda y. y)) x = x \tag{16}$$

$$(*f* f) (\tilde{a}) = \widetilde{f(a)} \tag{17}$$

$$(*f* (\lambda h. f(y + h))) x = (*f* f) (\tilde{y} + x) \tag{18}$$

$$(*f* (\lambda h. f(g(y + h)))) x = (*f* (f \circ g)) (\tilde{y} + x) \tag{19}$$

$$*f* \text{ rabs} = \text{hrabs} \tag{20}$$

$$x \neq 0_{\text{hr}} \implies (*f* \text{ rinv}) x = \text{hrinv } x \tag{21}$$

$$(*f* f) x \in *s* A \implies x \in *s* \{y. f y \in A\} \tag{22}$$

Theorem (17) is important, as it tells us that the extended function has the same solutions as its standard counterpart for all (embedded) real arguments. Theorems (18) and (19) are proved because of their importance in the nonstandard definition of derivatives. Theorems (20) and (21) confirm that the hyperreal absolute and inverse functions are nonstandard extensions of their real counterparts. Theorem (22) is a general lemma, needed for proofs in elementary real topology. One might try to picture these various theorems mentally, to get a better, more intuitive feel for the properties. If we combine $*$ -transforms of both sets and functions, we can derive further theorems, such as

$$*s* (f `` A) = (*f* f) `` (*s* A);$$

$$*s* \{x. \text{rabs } (f x - y) < r\} = \{x. \text{hrabs } ((*f* f) x - \tilde{y}) < \tilde{r}\}.$$

We note that any real constant is mapped to its embedded counterpart in the transform, as expected, while the functions are replaced by their nonstandard extensions.

The importance of internal sets and functions cannot be overstated. Lindstrøm calls them the ‘nice’ subsets and functions of nonstandard analysis [23], and draws an analogy to topology where, for example, the nice sets and functions are the open sets and continuous functions. Nice concepts are those that we are interested in whenever a new mathematical structure is introduced. In NSA, they are important because they enable hyperreal sets and functions to inherit properties from their standard counterparts in a natural way. They also enable us to express familiar concepts for our new mathematical structure that may be only partially inherited (such as the supremum property, which applies only to internal subsets of \mathbb{R}^*). The strict typing of Isabelle/HOL makes the new concepts clearer, and definitions ensure that their use is rigorous. We will later introduce some further extensions that enable us to deal with functions from \mathbb{N} to \mathbb{R} , for example.

11. Towards an intuitive calculus

Consider the real function $f(x) = x^2$. This extends naturally to a function f^* over \mathbb{R}^* . Now, if a is finite and ϵ is infinitesimal, then $f^*(a + \epsilon) = (a + \epsilon)^2 = a^2 + \epsilon(2a + \epsilon) \approx a^2 = f^*(a)$ since the set *Infinitesimal* is an ideal in *Finite*. Thus, an infinitesimal change in the argument x produces only an infinitesimal change in f . This is, intuitively, the behaviour expected from a continuous function such as $f(x)$ above; broadly speaking, one does not expect any sudden gap or jumps in the graph that represents the behaviour of the function. As pointed out by Keisler [20] and others [31], students who are just beginning to study calculus often find it difficult to cope with formulas involving quantifiers. The traditional epsilon and delta approach is a sudden leap from the intuitive calculus of school to the rigour of real analysis. One of the advantages of introducing the hyperreals is the simplification that this brings to the statement of many properties such as limits and continuity. For example, the ϵ and δ condition for a function f to be continuous at a ,

$$\forall \epsilon. (0 < \epsilon \longrightarrow \exists \delta. (0 < \delta \wedge \forall x. (0 < |x - a| < \delta \longrightarrow |f(x) - f(a)| < \epsilon)))$$

can be simplified to

$$\forall x. x \approx \tilde{a} \longrightarrow f^*(x) \approx \tilde{f}(a).$$

The approach, through the formal use of infinitesimals and relations such as \approx , retains much of the intuition that was present in school mathematics. The nonstandard treatment has been expounded in textbooks by Keisler [20], by Henle and Kleinberg [15] and more recently by Hoskins [16], for example. Keisler’s text has even been used successfully as an introductory textbook in calculus courses. There is much to be gained from carrying out proofs using a nonstandard formulation, and as this work shows next, even the mechanization of analysis becomes simpler and shorter due to the more algebraic nature of nonstandard analysis.

In applying nonstandard analysis to the formalization, we first introduce the standard and nonstandard formulations for the basic definitions in the theory. In the next step we prove that the standard and nonstandard definitions are equivalent. The nonstandard equivalents are then applied, whenever appropriate, to produce often shorter mechanical proofs of standard results. Thus, the use of NSA can effectively ease the task of mechanization. In the next sections, we will illustrate these points by mechanizing basic notions from the theories of limits for real sequences and series, elementary topology on the reals, limits and continuity of functions, and differentiability. We introduce and prove in Isabelle propositions stating

that the standard and nonstandard definitions for the various concepts are equivalent.

12. Real sequences and series

A real sequence $\langle a_n \rangle$ is viewed as a standard function, a , mapping the natural numbers into the reals; that is, $a : \mathbb{N} \rightarrow \mathbb{R}$. The notation $a(n)$ is also used to denote a typical term a_n of the sequence.

The function a has a nonstandard extension a^* which maps the hypernaturals into the hyperreals. The $*$ -transform of a is thus the function $a^* : \mathbb{N}^* \rightarrow \mathbb{R}^*$ where $a^*([\langle X_n \rangle]) = [a(X_n)]$ for any $[\langle X_n \rangle] \in \mathbb{N}^*$. We therefore define this in a similar fashion to the extension $*f*$ for real functions. In Isabelle, the nonstandard extension of a is given by $(*fNat* a)$ and defined as

```
*fNat* :: (nat => real) => hypnat => hypreal
*fNat* a N ≡ Abs_hypreal (⋃ X ∈ Rep_hypnat(N). hypreal ^ ^ {λn. a(Xn)}).
```

As can be seen, the nonstandard extension results in a sequence of hyperreals indexed not by the natural numbers, but by the hypernaturals. For this reason, the extended sequence is also known as a *hypersequence*.

Similar theorems to those presented in Section 10.4 about $*f*$ are proved, together with some new ones such as

$$(*fNat* (\lambda n. a(\text{Suc } n))) N = (*fNat* a) (N + 1).$$

Of particular importance is the theorem $(*fNat* a)(\bar{n}) = \widetilde{a(\bar{n})}$, which shows that the hypersequence agrees with the original sequence on \mathbb{N} ; that is, for any $n \in \mathbb{N}$, a_n^* is simply the image of a_n in the hyperreals. (From now on, we shall assume, for clarity, that 0 and 1 are overloaded over all the various types of numbers, and will refrain from using $0r$, $0hr$, $1hr$, and so on, which were defined in Section 6.)

12.1. On limits

The hyperreals are now used to define the concept of *limit*. A few observations about the notation need to be made first. The symbol ∞ is usually used in the real number system to denote that which is potentially arbitrarily large. The expression $\lim_{n \rightarrow \infty} a_n$ thus denotes the limiting value of a as n becomes an arbitrarily large natural number. In \mathbb{R}^* , the symbol ∞ can be viewed as having a similar meaning, but this time ‘arbitrarily large’ means a number larger than any *finite* number in \mathbb{R}^* . So the expression $\lim_{n \rightarrow \infty} a_n^*$ denotes the value infinitely close to a_n^* for any *infinite* hypernatural number n . This motivates the nonstandard definition for sequential limit that is given below.

With regard to the formalization in Isabelle, we decided to follow an approach similar to that used by Harrison in the HOL-Light system [14] and formulate both a relational and functional form for sequential limits. We declare and define an infix ‘tends to’ relation ‘ \longrightarrow ’ and use it to express statements such as a_n tends to l by $a_n \longrightarrow l$. The standard definition used in Isabelle is

$$X \longrightarrow l \equiv \forall r. (0 < r \longrightarrow (\exists N. \forall n. N \leq n \longrightarrow \text{rabs } (Xn - L) < r)).$$

Our formalization, however, also has a second version of the predicate, denoted by $\overset{NS}{\longrightarrow}$; this second notion of convergence is defined using nonstandard concepts and expressed by

the following simpler statement, not involving any existential quantifiers:

$$X \xrightarrow[NS]{} l \equiv (\forall N \in \mathbb{H}\text{NatInfinite}. (*f\text{Nat}* X)N \approx \tilde{l}).$$

The first task is to prove the equivalence of the two definitions. Before coming to this, we briefly make some remarks about the functional form of a sequential limit. We declare a constant `lim` and use it to denote the statement $\lim_{n \rightarrow \infty} a_n$ by `lim a` (equivalent by η -expansion to `lim ($\lambda n. a_n$)`). A nonstandard version of the function is also introduced, and that is denoted by `nslim`. The following definitions are made, using Hilbert’s ε -operator to denote the unique limit (if it exists):

$$\begin{aligned} \text{lim } a &\equiv \varepsilon l. a \longrightarrow l \\ \text{nslim } a &\equiv \varepsilon l. a \xrightarrow[NS]{} l. \end{aligned}$$

The relational form is preferred to prove properties about limits since the functional form is less powerful [14]. We do not have $a \longrightarrow \text{lim } a$ because all functions in HOL and, of course, Isabelle/HOL are total. The interested reader should consult Harrison’s PhD thesis for an extended discussion on binders, relational versus functional forms of mathematical statements, and other related issues arising from HOL’s lack of partial functions [14]. These points are equally relevant to the aspects of analysis that we have formalized in Isabelle. One last point is that, for a convergent sequence, the following theorem suggests an alternative definition for `nslim` (and hence `lim`):

$$\text{nslim } a = \text{st } ((*f\text{Nat}* a) \Omega)$$

where Ω denotes the infinite hypernatural $[\langle n \rangle]$. This is an interesting characterization of limit that arises due to the nonstandard framework.

We will now outline the steps needed to prove the equivalence of the standard and nonstandard definitions for limits.

12.2. Equivalence of standard and nonstandard definitions

Proving the equivalence of the standard and nonstandard formulations of a property is important, as it justifies using nonstandard methods to prove standard theorems. The proof that the nonstandard definition implies the standard definition is usually the trickier part. We need to go down to the level of the ultrafilter and use the theorems that recast properties such as those belonging to the set `Infinitesimal` in terms of membership of $U_{\mathbb{N}}$.

Theorem 12.1. *A sequence $a : \mathbb{N} \rightarrow \mathbb{R}$ converges to the real number l as its limit if and only if for each infinitely large hypernatural number $\eta = [\langle m_n \rangle] \in \mathbb{N}^* - \mathbb{N}$ we have that a_{η}^* is infinitely close to l . In symbols, $a \longrightarrow l \iff a \xrightarrow[NS]{} l$.*

Proof. We mechanize the proof given by Hurd [17] as follows.

$$1) \ a \longrightarrow l \implies a \xrightarrow[NS]{} l.$$

Assume that the sequence $\langle a_n \rangle$ converges to l . Let $0 < r$ be given and let $\eta = [\langle m_n \rangle]$ be any given infinite hypernatural number. Since $a \longrightarrow l$, there exists a natural number N such that $|a_n - l| < r$ for all $N \leq n$. Now, since η is an infinite hypernatural with representative sequence $\langle m_n \rangle$, we know, from the properties of infinite hypernaturals (see Section 8), that $N \leq m_n$ for almost all the m_n ; that is, $\{n. N \leq m_n\} \in U_{\mathbb{N}}$. But we can also prove that

$$\{n. N \leq m_n\} \subseteq \{n. |a_{m_n} - l| < r\}$$

from which it immediately follows that $\{n. |a_{m_n} - l| < r\} \in U_{\mathbb{N}}$. Thus, given any positive real number r , we have that $|a_{m_n} - l| < r$ for almost all the a_{m_n} . From this it follows that $a_{\eta}^* - \tilde{l}$ is infinitesimal; that is, $a_{\eta}^* \approx \tilde{l}$.

$$2) a \xrightarrow{NS} l \implies a \longrightarrow l.$$

Suppose that $\langle a_n \rangle$ does not converge to l . Then, there is some standard real $r > 0$ and a function $f : \mathbb{N} \rightarrow \mathbb{N}$ satisfying $n \leq f(n)$ and $r \leq |a_{f(n)} - l|$ for all $n \in \mathbb{N}$. Now, writing $f(n) \equiv f_n$, the sequence $\langle f_n \rangle$ defines a hypernatural number η , which we prove to be infinite. We have $\{n. r \leq |a_{f_n} - l|\} \in U_{\mathbb{N}}$ since it coincides with \mathbb{N} . Thus, it follows that $a_{\eta}^* - \tilde{l}$ is not infinitesimal. □

12.3. Remarks on the proof

There are several points that need to be made about the mechanical proof of the theorem above. As we mentioned already, the first part of the proof was relatively easy to mechanize, given that we had already proved various theorems expressing each class of hyperreal numbers in terms of the free ultrafilter. The second part needed several lemmas since it is more complicated. It involves, for example, the use of the axiom of choice (AC), which textbooks often fail to mention explicitly. In our mechanization, we use Hilbert’s description operator to prove the next theorem, which enables the existential quantifier to be pulled across the universal quantifier:

$$\forall x. \exists y. Q x y \implies \exists f. \forall x. Q x (f x). \tag{23}$$

This theorem allows us to introduce a function from \mathbb{N} to \mathbb{N} — effectively a sequence of natural numbers — that can be used to define an infinite hypernatural number. A brief remark is needed about theorem (23) above: in this particular proof where we only want to obtain a function $f : \mathbb{N} \rightarrow \mathbb{N}$, the use of AC is not strictly needed. Indeed, we can prove a special case of theorem (23) simply by letting $f(x)$ be the least y such that $Q x y$. This translates as a direct proof in Isabelle: we use the defined `LEAST` operator and let f be $\lambda x. \text{LEAST } y. Q x y$. (We thank the anonymous referee for pointing this out, and motivating the alternative mechanization in Isabelle.)

The following lemma is then proved on the way to the main result:

$$\forall n. n \leq f n \implies \text{Abs_hypnat} (\text{hypnatrel} \hat{\wedge} \{f\}) \in \text{HNatInfinite}.$$

Another important observation is that the structure of the proof follows a general pattern that will occur again when we mechanize the equivalence proofs for other properties. We typically need to use AC when proving that a particular nonstandard definition implies the standard one. Mechanical theorem-proving benefits from the re-use of code and lemmas.

The general pattern in the proofs is not a coincidence, and can be related to one of the central features of nonstandard analysis, known as the *transfer principle*. This provides a context in which true statements about \mathbb{R} are transformed into statements about \mathbb{R}^* . Within a typed logic, this procedure would involve lifting results from the type `real` to the type `hypreal`, from `nat` to `hypnat`, or (viewed more generally) from any type to its extended counterpart.

In the subsequent survey of the development of NSA in Isabelle, we shall state the standard and nonstandard formulations of various concepts, but often omit explicit details of the equivalence proof unless they differ considerably from the proof just given. We

shall, however, mention any interesting lemmas that were needed, as well as any particular difficulties encountered.

12.4. Properties of sequential limits

With the nonstandard formulation, the proofs of basic properties of sequences all become trivial. Indeed, their mechanization mostly involves simple algebraic manipulations that can be handled automatically by Isabelle’s simplifier. We prove the following theorems.

$$\frac{X \xrightarrow{NS} a \quad Y \xrightarrow{NS} b}{(\lambda n. X n + Y n) \xrightarrow{NS} a + b} \tag{24}$$

$$\frac{X \xrightarrow{NS} a \quad Y \xrightarrow{NS} b}{(\lambda n. X n \cdot Y n) \xrightarrow{NS} a \cdot b} \tag{25}$$

$$\frac{X \xrightarrow{NS} a}{\lambda n. - X n \xrightarrow{NS} -a} \tag{26}$$

$$\frac{X \xrightarrow{NS} a \quad a \neq 0}{(\lambda n. \text{rinv } X n) \xrightarrow{NS} \text{rinv } a} \tag{27}$$

$$\frac{X \xrightarrow{NS} a \quad X \xrightarrow{NS} b}{a = b} \tag{28}$$

For the proof of theorem (24) above, for example, we have that $X_n^* \approx \tilde{a}$ and $Y_n^* \approx \tilde{b}$, and hence that $X_n^* + Y_n^* \approx \tilde{a} + \tilde{b}$ for any infinite hypernatural n , since the infinitely close relation is closed under addition (see Section 7.3). The proof is done in one step using Isabelle’s automatic tactic. The other theorems are all proved as simply, the only exception being theorem (27). This requires a bit more work, and the following lemma:

$$X^*N \neq 0 \implies (\lambda m. \text{rinv } (X m))^*N = \text{hrinv } (X^*N).$$

This result effectively performs the $*$ -transform over both the inverse function and the sequence function, since $\text{hrinv} = \text{rinv}^*$. Once these basic properties have been proved, we can deal with the important concept of Cauchy sequences and their associated theorems.

12.5. Sequences

In this section, we examine some of the important properties of sequences formalized in Isabelle. We first examine the concept of a bounded sequence.

12.5.1. Boundedness and monotonicity

We define the standard and nonstandard notions of a bounded sequence as follows:

$$\begin{aligned} \text{Bseq } X &\equiv \exists K. (0 < K \wedge \forall n. \text{rabs } (X n) \leq K) \\ \text{NSBseq } X &\equiv \forall N \in \text{HNatInfinite}. (*\text{fNat}^* X) N \in \text{Finite}. \end{aligned}$$

The equivalence of the standard and nonstandard definitions for boundedness is first proved, thereby making two characterizations of the concept available for use in our proofs. The nonstandard definition, NSBseq , makes it immediately obvious that boundedness is a necessary condition for convergence. We have the following theorem.

Theorem 12.2. $\text{NSconvergent } X \implies \text{NSBseq } X$ where

$$\text{NSconvergent } X \equiv (\exists l. X \xrightarrow[\text{NS}]{} l).$$

This reduces, in Isabelle, to proving the following goal:

$$\begin{aligned} & \exists l. \forall N \in \text{HNatInfinite}. (*fNat* X) N \approx \tilde{l} \\ \implies & \forall N \in \text{HNatInfinite}. (*fNat* X) N \in \text{Finite}. \end{aligned}$$

Proof. Suppose that $\langle X_n \rangle$ converges to some $\alpha \in \mathbb{R}$. Then $X_n^* \approx \tilde{\alpha}$ for every infinite hypernatural n , and must therefore be finite by the following lemma:

$$x \in \text{Finite} \wedge x \approx y \implies y \in \text{Finite}.$$

□

The theorem is proved in one step by Isabelle’s `blast_tac`. We also prove that boundedness is a sufficient condition for convergence, provided that a given sequence is *monotone*:

$$\text{Bseq } X \wedge \text{monoseq } X \implies \text{convergent } X$$

where the monotonicity of a sequence X is defined by

$$\text{monoseq } X \equiv ((\forall m n. m \leq n \longrightarrow X m \leq X n) \vee (\forall m n. m \leq n \longrightarrow X n \leq X m)).$$

The proof of the theorem above proceeds through a mixture of standard and nonstandard arguments: for some of the lemmas, it is easier to prove a standard version rather than a nonstandard one. This is the case for the following result, for example:

$$\forall n. m \leq n \longrightarrow X n = X m \implies \exists l. X \longrightarrow l.$$

The standard proof is trivial since the variables are easy to instantiate by a routine examination of the goal. Isabelle’s automatic tactic then proves the theorem without difficulty. A nonstandard proof, however, would require proving a more demanding theorem:

$$\forall n. m \leq n \longrightarrow X n = X m \implies \exists l. \forall N \in \text{HNatInfinite}. (*fNat* X) N \approx \tilde{l}.$$

This is one of the few cases where we have noticed that a nonstandard proof seems to be more complicated than its standard counterpart. The difficulty here lies in finding the right instantiation for the existential variable.

12.5.2. Cauchy sequences

The following statements are equivalent.

Theorem 12.3 (Convergence). *The sequence $\langle a_n \rangle$ converges; that is, $\exists l. a_n \longrightarrow l$.*

Theorem 12.4 (Hyperreal Cauchy condition). *For all infinite hypernatural numbers N and M , $a_N^* \approx a_M^*$.*

Theorem 12.5 (Real Cauchy condition). *For all $0 < \epsilon$ there is an integer M such that for all $m, n \geq M$, $|a_m - a_n| < \epsilon$.*

The standard proof that a sequence is Cauchy if and only if it is convergent can be obtained from most traditional textbooks on analysis. Harrison [14], for example, uses the proof from Burkill and Burkill [5]. Although, the mechanization is reported as being a direct formalization in the HOL-Light system, Harrison’s proof is rather complicated. This is partly due to difficulties inherent in finding the right instantiations for variables in ϵ and δ proofs, especially since HOL-Light does not allow unknown variables whose instantiations can be delayed. Owing to this problem, Harrison suggests that Isabelle might provide a more natural environment for ϵ and δ proofs. Although this in itself seems a reasonable argument, we actually go one step further by using nonstandard arguments: our formalization avoids the need for ϵ and δ arguments altogether.

To prove the Cauchy criterion for convergence, Burkill and Burkill, and hence Harrison, define the extra notion of a subsequence. They then prove that every sequence has a monotonic subsequence. Although the main theorem is not difficult to reach once this result and a few other lemmas have been set up (Harrison also needs to define a ‘reindexing’ function in his formalization, for example), one might feel that the need for various auxiliary notions diverts attention from what is actually being proved. The need to introduce and use the properties of subsequences is not immediately obvious to anyone trying to prove the theorem (without the help of a textbook, for example).

Our formalization avoids the notion of a subsequence and goes for a direct and more intuitive proof. First we prove the equivalence of the real (standard) and hyperreal (non-standard) Cauchy conditions. This resembles that of Theorem 12.1. With this equivalence set up, the proof of the main result is simple and direct within the nonstandard framework.

Theorem 12.6. *The sequence $\langle X_n \rangle$ converges if and only if it is Cauchy.*

Proof. If $\langle X_n \rangle$ converges to l then $X_n^* \approx \tilde{l} \approx X_m^*$ for all infinite n and m by the nonstandard definition of convergence; so $\langle X_n \rangle$ is a Cauchy sequence by the nonstandard definition of the Cauchy criterion.

Conversely, if $\langle X_n \rangle$ is a Cauchy sequence then $\langle X_n \rangle$ is bounded and so X_n^* is finite for all infinite n . Therefore, using the standard part theorem, there exists a standard (embedded) real number l infinitely close to X_Ω^* where Ω is our usual infinite hypernatural number (see Section 12.1, for example). Thus, we have that $X_n^* \approx X_\Omega^* \approx l$ for all infinite n (nonstandard Cauchy criterion), and so $\langle X_n \rangle$ converges to l (nonstandard formulation for convergence). □

One lemma, also needed by Harrison, requires proving that every Cauchy sequence is bounded. We use the nonstandard version of this theorem involving the hyperreal formulations of both the Cauchy and boundedness properties.

A historical note: though infinitesimals do not appear in the standard definition of Cauchy convergence, Cauchy used them as a tool in his *Cours d’analyse* (1821) [22]. Indeed, Cauchy explicitly states the following as an alternative version of convergence: ‘in other words, it is necessary and sufficient that, for infinitely large values of the number n , the sums $s_n, s_{n+1}, s_{n+2}, \dots$ differ from the limit s , and consequently among themselves, by infinitely small quantities.’ Reinterpreted, within the context of nonstandard real analysis, this corresponds exactly to the hyperreal Cauchy condition. Laugwitz (further) mentions that Euler was the first, much earlier, in 1735, to state that $s_n - s_m$ being infinitesimal for

infinitely large m, n was a necessary and sufficient condition for convergence [22]. Such use of infinitesimals, especially by the rigorous Cauchy, gives yet another indication of their power as a tool in analysis throughout the centuries.

12.5.3. Sequences and hyperreals

There is, as expected, a close relationship between sequences and hyperreal numbers. Indeed since the development of the hyperreals has been based on the use of sequences of real numbers, we can prove the following theorems.

Theorem 12.7. *If $\langle a_n \rangle$ is bounded then $[\langle a_n \rangle]$ is finite; expressed as a theorem of Isabelle, we have*

$$\text{NSBseq } X \implies \text{Abs_hypreal } (\text{hyprel}^{\wedge\wedge}\{X\}) \in \text{Finite.}$$

Theorem 12.8. *If $\langle a_n \rangle$ converges to zero then $[\langle a_n \rangle]$ is an infinitesimal.*

Theorem 12.9. *If $\langle a_n \rangle$ is an unbounded sequence then $[\langle a_n \rangle]$ is an infinite hyperreal.*

12.6. Series

In standard analysis an infinite series is the limit of a sequence of finite sums. Despite the notation

$$\sum_{i=0}^{\infty} a_i$$

one does not try in classical analysis to interpret it literally as an infinite number of additions. Instead, one considers the sums of finitely many of the terms of the series, and examines the behaviour of such sums as an increasingly large, but still finite, number of terms is allowed. Using our framework, however, it is possible to use the nonstandard criterion for sequential convergence to define *literally* infinite sums.

12.6.1. Infinite sums and infinite series

Given a real sequence (f_n) , we define the standard notion of a finite sum $(\sum_{i=m}^{n-1} f_i)$ using Isabelle's *primrec* package, which implements primitive recursion:

```
consts sumr :: [nat, nat, (nat  $\Rightarrow$  real)]  $\Rightarrow$  real
primrec
sumr m 0 f = 0
sumr m (Suc n) f = if n < m then 0
                    else sumr m n f + f(n).
```

The first line declares `sumr` to be a constant. The `primrec` declaration provides a safe way of defining primitive recursion on datatypes [27]. Isabelle checks whether the reduction rules given for `sumr` satisfy a primitive recursive definition, thereby ensuring consistency, and supplies the reduction rules to the simplifier.

The expected theorems about finite sum are easily proved, mostly through induction followed by simplification. We shall not list them here, but will instead describe how the canonical nonstandard extension of `sumr` is defined.

Consider a sequence of finite sums: this constitutes a mapping from \mathbb{N} to \mathbb{R} which has a unique nonstandard extension defined, for any infinite hypernatural numbers $M = [\langle X_n \rangle]$ and $N = [\langle Y_n \rangle]$, as

$$\sum_{i=M}^N * a_i \equiv \left[\left\langle \sum_{i=X_1}^{Y_1} a_i, \sum_{i=X_2}^{Y_2} a_i, \sum_{i=X_3}^{Y_3} a_i, \dots \right\rangle \right]. \tag{29}$$

This enables one to talk of the sum being taken to N terms (M can be set to 0), where N is any hypernatural number. The value of such an *infinite* sum is a hyperreal number which depends on the number of terms taken. The formalization of the nonstandard extension in definition (29) is given in Isabelle by

```
sumhr :: (hypnat * hypnat * (nat => real)) => hyperreal
sumhr p ≡ ( λ(M, N, f).
  Abs_hyperreal( ( ∪ X ∈ Rep_hypnat M.
    ∪ Y ∈ Rep_hypnat N.
    hyperreal ^^ {λn. sumr ((Xn), (Yn), f)})) ) p.
```

As is usual in such cases, the corresponding simplification theorem is proved; it can be added to Isabelle’s simplifier when needed:

```
sumhr (Abs_hypnat (hypnatrel ^^ {λn. X n}),
  Abs_hypnat (hypnatrel ^^ {λn. Y n}), f)
= Abs_hyperreal (hyperreal ^^ {λn. sumr (X n, Y n, f)}).
```

Using this definition, theorems similar to the two reduction rules in the recursive definition of `sumr` are proved:

```
sumhr (m, 0, f) = 0
sumhr (m, n + 1, f) = if n < m then 0
  else sumhr (m, n, f) + (*fNat* f) n.
```

The nonstandard extension, with its possibly infinite hypernatural limits, preserves the formal behaviour of finite summation. In fact, with the help of the theorems just introduced, the properties of the finite sum are directly transferred from `sumr` to `sumhr`. A few of the theorems proved in Isabelle are as shown below.

$$\text{sumhr } (m, n, f) + \text{sumhr } (m, n, g) = \text{sumhr } (m, n, \lambda i. f i + g i) \tag{30}$$

$$\text{sumhr } (0, \Omega, \lambda i. 1) = \omega - 1 \tag{31}$$

$$\text{sumhr } (0, 2\Omega, \lambda i. (-1)^{\text{Suc } i}) = 0 \tag{32}$$

$$\text{sumhr } (0, 2\Omega - 1, \lambda i. (-1)^{\text{Suc } i}) = 1 \tag{33}$$

$$\text{sumhr } (m, n, \lambda i. r \cdot (f i)) = \tilde{r} \cdot \text{sumhr } (m, n, f) \tag{34}$$

$$\text{hrabs } (\text{sumhr } (m, n, f)) \leq \text{sumhr } (m, n, \lambda i. \text{rabs } (f i)) \tag{35}$$

$$n < p \implies \text{sumhr } (0, n, f) + \text{sumhr } (n, p, f) = \text{sumhr } (0, p, f) \tag{36}$$

$$(\forall r. m \leq r \wedge r < n \implies f r = g r) \implies \text{sumhr } (\bar{m}, \bar{n}, f) = \text{sumhr } (\bar{m}, \bar{n}, g) \tag{37}$$

$$\text{sumhr } (0, N, f) = (*fNat* (\lambda n. \text{sumr } (0, n, f)))N \tag{38}$$

In theorem (31), Ω once more refers to the infinite hypernatural $[\langle 0, 1, 2, \dots \rangle]$, while ω refers to the infinite hyperreal $[\langle 1, 2, \dots \rangle]$ (see Sections 7.1 and 8.1). The sum involved in this theorem can thus be literally taken as infinite. It is proved by observing that, according

to the definitions formalized in Isabelle,

$$\begin{aligned} \sum_{i=0}^{\Omega^*} 1 &= \left[\left\langle \sum_{i=0}^0 1, \sum_{i=0}^1 1, \sum_{i=0}^2 1, \dots \right\rangle \right] \\ &= \{0, 1, 2, \dots\} \\ &= \{1, 2, 3, \dots\} - \{1, 1, 1, \dots\} \\ &= \omega - 1. \end{aligned}$$

Of the other theorems shown, theorem (38) is perhaps the best illustration that `sumhr` is the nonstandard extension of `sumr`. It shows how the framework naturally extends any standard function (of a single variable), enabling it to take a nonstandard argument. This theorem is important to the derivation of results about convergence of series. Theorems (32) and (33) illustrate the comment made above that the value of the infinite sum depends on the number of terms taken.

Following Harrison [14], a relation `sums` is defined to denote that an infinite series converges to some limit a as its sum. An infinite series $\sum_{i=0}^{\infty} f_i$ ‘sums to’ some real number a if and only if the sequence of *partial sums* $\sum_{i=0}^n f_i$ converges to a as its limit. This provides the following definition in Isabelle:

$$f \text{ sums } a \equiv (\lambda n. \text{sumr } (0, n, f)) \longrightarrow a.$$

Hence, it also follows that the infinite series is convergent if and only if the sequence $(\lambda n. \text{sumr } (0, n, f))$ is a Cauchy sequence.

In nonstandard terms, the definition of a convergent series is given as follows.

Definition 12.1. The infinite series defined by the sequence $\langle f_n \rangle$ is said to *converge* if there exists some real number a such that for every infinite hypernatural number N ,

$$\sum_{i=0}^N f_i \approx a.$$

In Isabelle, this definition becomes

$$f \text{ NSsums } a \equiv (\forall N \in \text{HNatInfinite}. \text{sumhr } (0, N, f) \approx \tilde{a}).$$

From this definition, the following theorems are proved.

Theorem 12.10. A necessary and sufficient condition for an infinite series to converge is that for any two infinite hypernatural numbers M and N , we have

$$\sum_{i=0}^M f_i \approx \sum_{i=0}^N f_i$$

or, equivalently in Isabelle,

$$\exists a. f \text{ NSsums } a \iff \forall M \in \text{HNatInfinite}. \forall N \in \text{HNatInfinite}. \text{sumhr } (0, M, f) \approx \text{sumhr } (0, N, f).$$

Theorem 12.11. The theorem above is also expressed in an alternative form using result (36) from the list of theorems given about `sumhr`:

$$\exists a. f \text{ NSsums } a \iff \forall M \in \text{HNatInfinite}. \forall N \in \text{HNatInfinite}. M < N \longrightarrow \text{sumhr } (M, N, f) \approx 0.$$

As we have seen, NSA does indeed simplify the treatment of real sequences and infinite series. As a further benefit, the nonstandard extension of sums enables us to treat finite and infinite series in a homogeneous fashion. There is no need to use ∞ as a purely notational device in defining infinite series; it is now possible to take the sum to N terms, where N can be a natural number or an infinite hypernatural. The ∞ symbol now stands for any member of HNatInfinite .

13. Some elementary topology of the reals

We now survey the development of some basic topology on the reals in Isabelle. The aim of this formalization is to see the benefits that might be gained using nonstandard analysis when dealing with elementary topological notions such as open sets and neighbourhoods.

13.1. Neighbourhoods

We begin by giving the standard and nonstandard definitions of the *neighbourhood* of a point. For the standard definition, the concept of a *ball* is first defined. If a is any point in \mathbb{R} and r is any real number, then the set of all real points x whose distance from a is less than r is defined as

$$\text{rBall } a \ r \equiv \{x. \text{rabs } (a - x) < r\}.$$

Definition 13.1. A set $M \subseteq \mathbb{R}$ is said to be a (*standard*) *neighbourhood* of point $a \in \mathbb{R}$ if and only if there exists some $r > 0$ such that

$$\text{rBall } a \ r \subseteq M.$$

Expressing this in Isabelle, we have

$$\text{isnbhd } a \ M \equiv \exists r. 0 < r \wedge \text{rBall } a \ r \subseteq M.$$

The nonstandard formulation, on the other hand, is given by the following definition.

Definition 13.2. A set $M \subseteq \mathbb{R}$ is said to be a (*nonstandard*) *neighbourhood* of point a if and only if every hyperreal x infinitely close to a belongs to the nonstandard extension M^* of M .

In Isabelle, this is formalized as

$$\text{isNSnbhd } a \ M \equiv \text{monad } (\tilde{a}) \subseteq \text{*s* } M.$$

As can be seen, the concept of a *monad* (named as a tribute to Leibniz) enables the definition to be expressed concisely. The monad is a set of hyperreals, formally defined by

$$\text{monad } x \equiv \{y. x \approx y\}.$$

The next step, as usual, is to prove the equivalence of the two definitions as a theorem in Isabelle. The proof is mechanized without much difficulty with the help of result (10) from Section 10.2. This lemma is necessary to prove that the nonstandard definition implies the standard one. The formulations are next used to deal with the notion of open sets.

13.2. Open sets

A subset G of \mathbb{R} is said to be *open* if and only if G is a neighbourhood of each of its points. This leads to the following direct formalization of the standard and nonstandard characterizations:

$$\begin{aligned} \text{isOpen } G &\equiv \forall a \in G. \text{isnbhd } a \ G; \\ \text{isNSOpen } G &\equiv \forall a \in G. \text{isNSnbhd } a \ G. \end{aligned}$$

The equivalence proof follows trivially from that of neighbourhood. The theorems given next are all proved automatically. They are direct consequences of the results obtained about boolean operations on nonstandard extensions of sets (see Section 10.2).

$$\begin{array}{c} \overline{\text{isOpen } \emptyset} \quad \overline{\text{isOpen (UNIV :: real set)}} \\ \frac{\overline{\text{isOpen } A} \quad \overline{\text{isOpen } B}}{\overline{\text{isOpen } A \cap B}} \quad \frac{\overline{\text{isOpen } A} \quad \overline{\text{isOpen } B}}{\overline{\text{isOpen } A \cup B}} \\ \frac{\text{isOpen } A}{\text{isOpen } (\bigcup A)} \end{array}$$

By contrast, and as an example, a *standard* proof in Isabelle that open sets are closed under finite intersections requires several steps including an explicit instantiation of variables, a case split, and the use of the following lemma:

$$r_1 < r_2 \wedge x \in \text{rBall } a \ r_1 \implies x \in \text{rBall } a \ r_2.$$

The gain from using nonstandard analysis seems obvious once again. In this development of elementary real topology, several other concepts (such as closed sets, limit points, and derived sets) are also introduced. Their various properties are formalized, and in most cases the proofs are automatic. One of the main results to be formalized in this theory using a nonstandard approach is the Bolzano–Weierstrass theorem. Its proof, as given by Hurd [17], is extremely short and simple compared to the standard proof.

14. Limits and continuity

There are several notions of limits that share a number of common theorems (such as uniqueness). It is clear that an efficient mechanization of standard analysis should seek to minimize proof replication by developing a generic treatment of limits. Harrison uses the well-known theory of convergence nets to prove a number of general theorems that can then be specialized to fit each notion of ‘limit’ [14].

Since our development involves standard theorems about limits using a nonstandard approach, we did not feel a need for such a generic treatment of limits. Moreover, this is only an initial investigation into the benefits to be gained from working in the hyperreals. So there is scope for further improvement. An interesting idea would be to seek a generalization for the nonstandard theory of limits as well. However, since we are already working with much simpler and more algebraic formulations than in the standard case, the gains might not be worth the trouble. After all, as we noticed in our development, having independent notions of sequential and pointwise limits does not represent a lot of extra work since the proofs of similar properties are all done automatically. Having said this, it is probably wise in any mechanization to favour the approach that cuts down on work. This would prevent

us from having two similar-looking theorems like the ones below, which were used for sequential and pointwise limits respectively:

$$\begin{aligned}
 (*fNat* f) N \approx l \wedge (*fNat* g) N \approx m & \\
 \implies (*fNat* (\lambda y. f y + g y)) N \approx l + m & \\
 (*f* f) x \approx l \wedge (*f* g) x \approx m & \\
 \implies (*f* (\lambda y. f y + g y)) x \approx l + m. &
 \end{aligned}$$

All this falls under the more general concept of preservation of properties across nonstandard extensions. We would like to prove general properties that hold for all nonstandard extensions of functions, rather than deal with specific cases like those above. Textbooks usually state the properties that we presented in Section 10.4 as general results that apply to all extensions. In our case, since we extend each type of function explicitly, we need to prove similar properties each time.

Let us now return to the standard and nonstandard characterizations of the notions of pointwise limits. A function f is said to have a limit l as x approaches a point a if and only if for any given $\epsilon > 0$, there exists a $\delta > 0$ such that for every value of x satisfying the inequality $0 < |x - a| < \delta$, we have $|f(x) - l| < \epsilon$. This is the standard ϵ and δ definition for the limit of a function at a point. The conventional notation is $\lim_{x \rightarrow a} f(x) = l$. We will, however, use a relational approach in this case as well, and denote the condition by $f \xrightarrow{a} l$ for the standard case, and by $f \xrightarrow[NS]{a} l$ for the nonstandard one. In Isabelle, we have:

$$\begin{aligned}
 f \xrightarrow{a} l \equiv \forall \epsilon. 0 < \epsilon \longrightarrow (\exists \delta. 0 < \delta \wedge \\
 (\forall x. 0 < \text{rabs } (x - a) \wedge \text{rabs } (x - a) < \delta \\
 \longrightarrow \text{rabs } (f x - l) < \epsilon)).
 \end{aligned}$$

The nonstandard definition, once again, is more concise and captures the intuition behind the notion:

$$f \xrightarrow[NS]{a} l \equiv \forall x. x \neq \tilde{a} \wedge x \approx \tilde{a} \longrightarrow (*f* f) x \approx \tilde{l}.$$

The equivalence of the two definitions is not too difficult to prove and has the same structure as the other, similar proofs. We make use of a few lemmas such as

$$\begin{aligned}
 \forall n. \text{rabs } (X n - x) < \text{rinv } (n) \implies \\
 (\text{Abs_hypreal } (\text{hyprel } \hat{\wedge} \{X\}) - \tilde{x}) \in \text{Infinitesimal},
 \end{aligned}$$

which enables us to define a hyperreal infinitely close to a real number x , given a real sequence converging towards that number.

We prove properties analogous to those presented in Section 12.4. In this part of the formal investigation, however, we decided to prove some of the properties twice: first using only the standard approach and then using the nonstandard approach. The aim was to examine more closely the gains from using nonstandard analysis in terms of the number of steps required to complete each proof, instantiations of variables, and theorems used. If we consider, for example, the formalization of the addition property,

$$\frac{f \xrightarrow[NS]{a} l \quad g \xrightarrow[NS]{a} m}{(\lambda x. f(x) + g(x)) \xrightarrow[NS]{a} l + m}$$

a few interesting remarks can be made.

- The nonstandard proof expands the definitions and is completed automatically in one step (0.08 seconds):

```
Goalw [NSLIM_def]
" [| f -- a --NS> l; g -- a --NS> m | ]
==> (%x. f(x) + g(x)) -- a --NS> (l + m) ";
by (auto_tac (claset() addSIs [starfun_add_inf_close],
simpset() addsimps [hypreal_real_add]));
```

while the standard proof, with our direct formalization, takes some 15 steps.

- We need to give instantiations of variables in several steps for the standard proof. The level of automation is thus fairly low, and requires the user to pay attention to a lot of details. Moreover, there is the added difficulty of deciding what the instantiation should be, and dealing with a three-way case split arising from the linear ordering of the reals.
- The standard proof requires theorems about the transitivity of the ordering relation, the absolute function (triangle inequality theorem), the monotonicity of the ordering relation under addition, and so on, while the nonstandard proof needs only a theorem about the monotonicity of the \approx relation under addition, and one about the preservation of the addition operation by the embedding function for the reals. Both of these are supplied to Isabelle’s automatic tactic as shown.

Therefore, we notice that the nonstandard proof offers a clear gain in automation. The user is freed from some of the more tedious steps through the use of the simpler formalization.

In addition, theorems such as

$$f \xrightarrow{a} l \iff (\lambda h. f(a + h)) \xrightarrow{0} l$$

are simple to prove using the nonstandard formulation. This is a useful lemma that can be used to simplify theorems about continuity and differentiability, for example. Next, the standard notion of continuity is examined. A standard real function f is continuous at a point a when $f(x)$ tends to $f(a)$ as x tends to a . In Isabelle,

$$\text{isCont } f \ a \equiv (f \xrightarrow{a} f \ a).$$

We give again the nonstandard definition of continuity that we mentioned in Section 11. A standard real function f is *continuous* at the point a if and only if $f^*(x)$ is infinitely close to $f(a)$ for every hyperreal x infinitely close to a . Expressed formally in Isabelle,

$$\text{isNSCont } f \ a \equiv (\forall x. x \approx \tilde{a} \longrightarrow (*f*) x \approx \widetilde{f(a)}).$$

Once again, the formalization makes it explicit that the definition is referring to the embedded copies of a and $f(a)$ in the hyperreals. The equivalence of the two definitions follows immediately from that of standard and nonstandard limits. A number of useful theorems are proved immediately. Examples are:

$$\begin{aligned} \text{isNSCont } f \ a &\iff (f \xrightarrow[NS]{a} f \ a) \\ \text{isCont } f \ a &\iff (\lambda h. f(a + h)) \xrightarrow{0} f \ a. \end{aligned}$$

We also have two distinct ways of proving the usual theorems about continuous functions.

- 1) The theorems can be proved as results of the corresponding theorems for pointwise limits. This is a conventional approach and, although (some of) the limit theorems themselves might have been proved using NSA, the process is wholly standard.

- 2) They can be proved as simple algebraic consequences of the nonstandard formulation of continuity. This approach bypasses the limit results, and provides alternative simple proofs. Moreover, it has an added power. It can prove at least one elementary result — the composition of continuous functions — that does not follow from limit theorems. This is examined next.

We prove that the sum, product, and division of continuous functions are also continuous. These are results that can be proved in either of the two ways mentioned above. We also prove the following theorem.

Theorem 14.1. *The composition of continuous functions is continuous:*

$$\text{isCont } f \ a \wedge \text{isCont } g \ (f \ a) \implies \text{isCont } (g \ o \ f) \ a.$$

Proof. If $x \approx \tilde{a}$ then $f^*(x) \approx \widetilde{f(a)}$, and so it follows that $g^*(f^*(x)) \approx g(\widetilde{f(a)})$. □

This result is proved automatically by Isabelle’s `auto_tac`. Contrast this with Harrison’s corresponding proof, which is longer, and requires the instantiation of ϵ and δ properties. In a sense, this also hints at another powerful aspect of nonstandard techniques in mechanical theorem-proving: their simple algebra enables them to deal uniformly with a wide range of theorems. The standard approach, on the other hand, required Harrison to go back to a direct formalization in the HOL-Light system because the theorem does not follow from any of the results about limits. An analogous difficulty occurs if the standard treatment is used to formalize the chain rule of differentiation.

Using the nonstandard framework, it is an interesting exercise to prove more involved theorems such as the following topological characterization of continuity. A function f is continuous on \mathbb{R} if and only if the inverse image $\{x \in \mathbb{R}. f(x) \in A\}$ of any open set A is itself always an open set. In Isabelle, the following theorem is proved without any difficulties:

$$(\forall x. \text{isCont } f \ x) \iff (\forall A. \text{isOpen } A \implies \text{isOpen } \{x. f(x) \in A\}).$$

Proofs of the theorems about limits, topological notions and so on only refer to the free ultrafilter when we are proving the equivalence of the standard and nonstandard definitions. All the other theorems are proved at the more intuitive algebraic level. The equivalence theorems are essential because the standard formulations are the ones that are in widespread use. With the success and widening acceptance of NSA, it might be that in a few decades the so-called ‘nonstandard’ definitions will become the established ones.

Using the various continuity theorems, we have mechanized nonstandard proofs by Hurd [17] of some important results of real analysis.

Theorem 14.2 (Intermediate value theorem). *If f is continuous on the closed interval $[a, b]$ and $f(a) < d < f(b)$ for some d , then there exists a term c between a and b with $f(c) = d$. The proof considers the points $x_k = a + k(a - b)/n, 0 \leq k \leq n$ and the values of f at x_k . The proof then proceeds through a $*$ -transform.*

Theorem 14.3 (Extreme value theorem). *If f is continuous on the closed and bounded interval $[a, b]$, then there exists a term c between a and b so that $f(x) \leq f(c)$ for all x between a and b . The proof proceeds succinctly using arguments similar to the ones above. The points $x_{n,k} = a + k(b - a)/n, 0 \leq k \leq n$ are considered this time.*

15. Differentiation

The development of the theory of differentiation builds upon the results of the previous section. The standard formulation states that a function f has a derivative d at a point x if $(f(x + h) - f(x))/h \rightarrow d$ as $h \rightarrow 0$. In Isabelle, we formalize the relational definition $\text{DERIV}(x) f := d$ meaning ‘the derivative of f at x is d ’ as

$$\text{DERIV}(x) f := d \equiv (\lambda h. (f(x + h) - f(x)) \cdot \text{rinv } h) \xrightarrow{0} d.$$

The notation $\text{DERIV}(x)$ can be regarded as a variation of the Leibniz notation, and as standing for d/dx . We prove this equivalent form of the standard definition, which is useful for some of our proofs:

$$\text{DERIV}(x) f := d \iff (\lambda z. (f(z) - f(x)) \cdot \text{rinv } (z - x)) \xrightarrow{x} d. \quad (39)$$

The nonstandard definition is stated as

$$\begin{aligned} \text{NSDERIV}(x) f := d \equiv & \forall h \in \text{Infinitesimal} - \{0\}. \\ & ((\text{*f* } f)(\widetilde{x} + h) - \widetilde{f}(x)) \cdot \text{hrinv } h \approx \widetilde{d}. \end{aligned}$$

We first prove that this nonstandard definition can also be given in terms of limits, exactly as the standard definition. The proof does not cause much difficulty and, from it, we see immediately that the two definitions of derivative are equivalent. In addition, using Theorem (39), we provide a second useful nonstandard characterization for the differentiability of a function f at a point x :

$$\begin{aligned} \text{NSDERIV}(x) f := d \iff & \forall y. y \approx x \wedge y \neq x \longrightarrow \\ & ((\text{*f* } f)(y) - \widetilde{f}(x)) \cdot \text{hrinv } (y - \widetilde{x}) \approx \widetilde{d}. \end{aligned}$$

We then proceed to prove the standard results in an extremely simple fashion. For example, we prove that a function f , differentiable at a point x , is continuous at that point:

$$\text{NSDERIV}(x) f := d \implies \text{isNSCont } f \ x.$$

This is a simple algebraic theorem using the nonstandard formulation, since $f^*(\widetilde{x} + h) - \widetilde{f}(x) \approx \widetilde{d} \cdot h$ for all $h \approx 0$, and so $f^*(\widetilde{x} + h) \approx \widetilde{f}(x)$; that is, f is continuous at x .

A functional form is also defined for the derivative using the standard part function and the non-zero infinitesimal ϵ defined previously:

$$\text{nsderiv}(x) f \equiv \text{st } (((\text{*f* } f)(\widetilde{x} + \epsilon) - \widetilde{f}(x)) \cdot \text{hrinv } \epsilon).$$

15.1. *Standard properties of derivatives*

We prove the familiar rules about the differentiation of simple functions and their combination as follows.

$$\frac{}{\text{NSDERIV}(x) (\lambda x. k) \text{ :> } 0}$$

$$\frac{\text{NSDERIV}(x) f \text{ :> } d}{\text{NSDERIV}(x) (\lambda y. c \cdot f(y)) \text{ :> } c \cdot d}$$

$$\frac{\text{NSDERIV}(x) f \text{ :> } d}{\text{NSDERIV}(x) (\lambda y. - f(y)) \text{ :> } - d}$$

$$\frac{\text{NSDERIV}(x) f \text{ :> } d \quad \text{NSDERIV}(x) g \text{ :> } e}{\text{NSDERIV}(x) (\lambda y. f(y) + g(y)) \text{ :> } d + e}$$

$$\frac{\text{NSDERIV}(x) f \text{ :> } d \quad \text{NSDERIV}(x) g \text{ :> } e}{\text{NSDERIV}(x) (\lambda y. f(y) \cdot g(y)) \text{ :> } d \cdot g(x) + e \cdot f(x)}$$

The absence of any explicit notions of limits makes many of the standard results about derivatives straightforward to derive. The properties follow from simple algebraic manipulations of infinitesimals. As a result, the simplifier of Isabelle plays an important part in these proofs, in doing the tedious term manipulation and cancellation. To achieve this, we might need to add rules for associative-commutative rewriting, for example. However, there are cases when we need to prove lemmas explicitly to help the simplifier to re-arrange terms. For example, to prove the theorem about the derivative of product, we need the following lemma:

$$(a \cdot b) - (c \cdot d) = b \cdot (a - c) + c \cdot (b - d).$$

15.2. *Chain rule*

One of the important theorems about differentiation is the *chain rule*. In his formalization of differentiation, Harrison reports on the problems that arise when proving this theorem directly. The main difficulty is that, when using the standard definition, the theorem does not follow directly from any limit results. Indeed, unlike continuity, limits are not compositional. To deal with this problem, Harrison had to formalize an alternative, rather different characterization of differentiability, the so-called ‘Carathéodory derivative’. In our case, however, due to the nonstandard formulation, the chain rule admits an entirely straightforward derivation. The Isabelle theorem is given as follows.

Theorem 15.1.

$$\frac{\text{NSDERIV}(a) g \text{ :> } d \quad \text{NSDERIV}((g a)) f \text{ :> } e}{\text{NSDERIV}(a) (f \circ g) \text{ :> } d \cdot e}$$

Proof. This follows immediately from

$$\frac{f^*(g^*(x)) - f^*(\widetilde{g(a)})}{x - \widetilde{a}} = \frac{f^*(g^*(x)) - f^*(\widetilde{g(a)})}{g^*(x) - \widetilde{g(a)}} \frac{g^*(x) - \widetilde{g(a)}}{x - \widetilde{a}} \approx d \cdot e.$$

□

This nonstandard proof, unlike its standard counterpart, reflects nicely and directly the intuition behind the Leibnizian notation for the rule:

$$\frac{df}{dx} = \frac{df}{dg} \frac{dg}{dx}.$$

It should be noted that Ballantyne and Bledsoe’s NSA prover [2] could not prove the chain rule automatically. In our case, we use a simple lemma to help set up the required product of fractions:

$$y \neq 0 \implies x \cdot z = (x \cdot \text{hrinv } y) \cdot (y \cdot z).$$

The main proof is directly formalized, though we have to do some manipulations explicitly — for example, we need to use one of Isabelle’s instantiation tactics with the lemma above to set the variable y in it to the correct binding. The level of automation could be made higher by building stronger routines in the simplifier to deal with division. For example, the recent addition of generic simplification procedures for subtraction have been helpful to many algebraic proofs. This is a case where the development of new theories can call for more support from the prover. This ultimately benefits many other theories.

Coming back to our development, we prove the theorems about the inverses and quotients of functions using the chain rule and the fact that, for non-zero x , the derivative of $f(x) = 1/x$ is $-1/x^2$. The proofs remain simple and algebraic. Stated as theorems of Isabelle, these various extra results (shown in terms of the equivalent standard notation) are formalized as follows.

$$\frac{x \neq 0}{\text{DERIV}(x) (\lambda x. \text{rinv } x) \text{ :> } - \text{rinv } (x^2)}$$

$$\frac{\text{DERIV}(x) f \text{ :> } d \quad f(x) \neq 0}{\text{DERIV}(x) (\lambda x. \text{rinv } (f x)) \text{ :> } -d \cdot \text{rinv } (f(x)^2)}$$

$$\frac{\text{DERIV}(x) f \text{ :> } d \quad \text{DERIV}(x) g \text{ :> } e \quad g(x) \neq 0}{\text{DERIV}(x) (\lambda z. f(z) \cdot \text{rinv } (g z)) \text{ :> } (d \cdot g(x) - e \cdot f(x)) \cdot \text{rinv } (g(x)^2)}$$

15.3. Rolle’s theorem

Rolle’s theorem involves notions from both continuity and differentiability.

Theorem 15.2 (Rolle’s theorem). *If f is defined and continuous on the finite closed interval $[a, b]$, $f(a) = f(b)$, and differentiable at least on the open interval (a, b) , then there exists x_0 between a and b such that $f'(x_0) = 0$.*

The formalized proof is taken from Hoskins [16], and proceeds through a case analysis on the values that f can take in the interval between a and b . The argument is once again nonstandard, and yields a direct formalization. In Isabelle, the theorem is given by

$$\begin{aligned} & a < b \wedge \\ & f(a) = f(b) \wedge \\ & \forall x. a \leq x \wedge x \leq b \implies \text{isNSCont } f \ x \wedge \\ & \forall x. a < x \wedge x < b \implies f \ \text{NSdifferentiable } x \wedge \\ & \implies \exists x_0. a < x_0 \wedge x_0 < b \wedge \text{NSDERIV}(x_0) f \text{ :> } 0 \end{aligned}$$

where the nonstandard infix predicate `NSdifferentiable` stands for ‘the real function f is differentiable at x ’ and is defined by

$$f \text{ NSdifferentiable } x \equiv \exists d. \text{NSDERIV}(x) f \text{ :> } d.$$

In the previous sections, we have presented an initial investigation of analysis using a nonstandard treatment. There are several important aspects of elementary analysis that still need to be formalized, including Taylor and power series and the theory of Integration. A nonstandard approach promises to be useful for these as well.

16. On the transfer principle

We now expand on the *transfer principle*, on which we remarked briefly in Section 12.3. Consider the statement, true in \mathbb{R} , stating that the set of natural numbers \mathbb{N} is unbounded as a subset of \mathbb{R} : the Archimedean property holds for the reals. Formalized in Isabelle, this is expressed by

$$\forall x::\text{real}. \exists n::\text{nat}. x < \text{real_of_nat } n.$$

Using the definitions of hyperreals and hypernaturals, and the properties of the free ultrafilter, we can then deduce the theorem that the set of hypernaturals \mathbb{N}^* is unbounded as a subset of the hyperreals \mathbb{R}^* . Stated in Isabelle/HOL, with explicit typing information shown, we have

$$\forall x::\text{hypreal}. \exists n::\text{hypnat}. x < \text{hypreal_of_hypnat } n.$$

This second statement about the hyperreals thus appears to be, in some sense, a transform of the original statement about the reals. One can go from one to the other, as this example illustrates, by making certain specific changes about the types of the terms (and the embedding functions) appearing in each. The crux of nonstandard analysis is that the transformation of statements along these lines can be carried out generally. It is this general idea that is captured by the transfer principle [16].

Theorem 16.1 (The transfer principle for real analysis). *There exists a set \mathbb{R}^* such that*

- 1) \mathbb{R} is a proper subset of \mathbb{R}^* ;
- 2) to each function $f : \mathbb{R} \rightarrow \mathbb{R}$ there corresponds a function $f^* : \mathbb{R}^* \rightarrow \mathbb{R}^*$ which agrees with f on \mathbb{R} ;
- 3) to each n -place relation P on \mathbb{R} there corresponds a n -place relation P^* on \mathbb{R}^* which agrees with P on \mathbb{R} .

Further, every well-formed statement φ formulated in terms of

- particular real numbers r_1, r_2, \dots, r_m ,
- particular functions f_1, f_2, \dots, f_m ,
- particular relations P_1, P_2, \dots, P_m ,
- logical connectives and quantifiers, with variables ranging over \mathbb{R}

is true with respect to \mathbb{R} if and only if the statement φ^* obtained from φ by replacing each f_k by f_k^* and each P_k by P_k^* , and by allowing variables to range over \mathbb{R}^* , is true with respect to \mathbb{R}^* .

In the current work, proving the equivalence of the standard and nonstandard formulations has involved working with sequences and checking whether certain sets belong to the

ultrafilter or not, each time a new property is introduced. By implementing some form of the transfer principle, one should be able to capture much of the power that NSA derives from the use of such metatheorems. This has not been investigated thoroughly — we have formalized part (1), and particular cases of parts (2) and (3) above, though — and so producing an effective form of the principle provides scope for further research. Our work has shown, however, that a powerful theory is still possible if one is willing to transfer properties by separate proofs. In fact, if we consider the example given at the beginning of this section, and its proof in Isabelle, we get an idea of what would be needed in most cases to enable transfer. So, for

$$\exists n::\text{hypnat}. x < \text{hypreal_of_hypnat } n,$$

the proof simply boils down to showing that the following theorem involving our free ultrafilter $U_{\mathbb{N}}$ holds:

$$\begin{aligned} &\forall m. X(m) < \text{real_of_nat } f\ m \\ \implies &\{n.X(n) < \text{real_of_nat } (f\ n)\} \in U_{\mathbb{N}}. \end{aligned}$$

This final subgoal, which is trivial to prove, is reached through three simple steps which involve

- 1) recasting the hyperreal x in terms of its underlying equivalence class;
- 2) expressing the Archimedean property of the reals in terms of the real sequence X introduced above to give

$$\forall m. \exists n. X(m) < \text{real_of_nat } n$$

which by the axiom of choice yields:

$$\forall m. X(m) < \text{real_of_nat } (f\ m);$$

- 3) instantiating the existential variable in the goal to the hypernatural

$$\text{Abs_hypnat } (\text{hypnatrel}^{\wedge}\{f\})$$

which, as can be seen, is defined using the sequence f above.

The proof of the theorem is four lines long, and can be routinely done. As expected, this compares favourably with the (mechanized) proof of the Archimedean property for the reals.

Moreover, to help our formalization, general automatic tactics to check whether supersets, intersections, or complements of sets belong to the free ultrafilter have been coded. These enable many of the goals to be greatly simplified, and in quite a few cases to be proved automatically. The idea behind the main tactic (`ultra_tac`) exploits the facts that the ultrafilter $U_{\mathbb{N}}$ is proper (that is, it does not contain the empty set), that for any subset A of the naturals either $A \in U_{\mathbb{N}}$ or its complement $-A \in U_{\mathbb{N}}$, and that $U_{\mathbb{N}}$ is closed under finite intersection and supersets. As an example, if the tactic is used on the following goal,

$$A \in U_{\mathbb{N}} \wedge \dots \wedge B \notin U_{\mathbb{N}} \wedge \dots \wedge X \in U_{\mathbb{N}} \dots \implies Z \in U_{\mathbb{N}}$$

it tries to solve it by looking for a proof that

$$A \cap -B \cap \dots \cap X \cap -Z \subseteq \emptyset.$$

If it succeeds, this means, by the superset property, that the empty set is a member of $U_{\mathbb{N}}$, which immediately leads to a contradiction. The tactic is wrapped around Isabelle's

`auto_tac`, which is used to perform simplification. This means that extra theorems can be added to the simplifier if they are needed to show that the intersection is empty.

17. Related work

The reals were first constructed in Automath in 1977 by Jutting [19], who translated Landau's famous monograph on the foundations of analysis [21]. More recently, Harrison has constructed the reals and formalized a substantial amount of analysis in the HOL-Light system [14]. The work of Harrison has influenced some of our decisions during mechanization, especially when formalizing analysis, where we have benefited from the observations made by him on notations, for example. As far as our own constructions up to the reals are concerned, we have mostly followed the presentation given by Gleason [11], since it matches the sequence of constructions that Conway advocates [8].

The automated theorem-proving community does not seem to have shown much interest in NSA, even though its importance has grown in many fields, such as physics, analysis and economics, where it has successfully been applied. Ballantyne and Bledsoe [2] implemented a prover using nonstandard techniques in the late seventies. Their work basically involved substituting any theorem in the reals \mathbb{R} by its analogue in the extended reals \mathbb{R}^* and proving it in this new setting. Even though the prover had many limitations, and the work was just a preliminary investigation, the authors argued that through the use of nonstandard analysis, they had brought some new and powerful mathematical techniques to bear on the problem.

Despite this rather promising work, there does not seem to have been much done over the last two decades. Chuaqui and Suppes [6] have proposed an axiomatic framework for doing proofs in NSA, and Bedrax has implemented a prototype for a simplified version of the Suppes–Chuaqui system called Infmal [3]. Infmal is implemented in Common Lisp and contains the various axioms (logical, algebraic and infinitesimal) required by the deduction system and extensions to the usual arithmetic operations. Unfortunately, Infmal is a simple experiment and, though interactive, is rather limited in the proofs it can carry out. There has also been some work carried out by Beeson [4] who developed a restricted axiomatic version of NSA using the logic of partial terms. The properties of the infinitely close relation, standard parts, infinitesimals and so on, are asserted as axioms leading to a theory similar in spirit to the one that could be developed starting from the axioms we give at the beginning of this paper in Section 3. Beeson uses NSA to ensure the correctness of applications of calculus in a system called *Mathpert* which combines computer algebra with theorem-proving.

In our development we have verified the various basic axioms asserted by Beeson in his approach. Moreover, we have also verified, through our strictly definitional approach, the axioms about properties of the hyperreals that were built into Ballantyne and Bledsoe's prover.

18. Concluding remarks

As far as we are aware, there has not been any previously published construction of the hyperreals using a mechanical theorem-prover. This paper has described the construction process resulting in a proper field extension of the reals. Various classes of numbers, including the notorious infinitesimals, have been formally defined, and their properties formalized. The \approx (infinitely close) relation has been introduced, which is crucial to the formalization of nonstandard real analysis and to our own work on the formalization of Newton's *Principia*

[10]. The framework has been shown to be flexible by allowing the hypernaturals, and their associated properties, to be formalized with minimal effort.

To reach the hyperreals has involved all the constructions up to the reals (which we have not described in much detail) and proving the various properties of each number system introduced; it also involved working in Isabelle/HOL set theory to formalize Zorn's lemma and the theory of filters and ultrafilters. As might be expected, a number of interesting remarks emerge from this development. We outline some of these next.

The formalization of filters is an important contribution. They have numerous applications in set theory, logic, algebra, and so forth. They can also be used to study the various notions of convergence; they yield essentially the same results as convergence nets [30]. Nets provide a natural generalization of sequences and are commonly used in analysis. In fact, nets are also useful to the mechanization of analysis, as was shown by Harrison [14]. Thus, Isabelle's theory of filters could be used for a general theory of convergence.

Since this work formalizes the ultrafilter theorem, the ultrapower construction becomes available for the development of other nonstandard number systems. For instance, the hyperintegers or hypercomplex numbers could be introduced. In particular, it becomes possible to construct the *hyperhyperreal* numbers from the hyperreals. These numbers were introduced by Henle and Kleinberg [15], for example, and are shown to contain, in addition to the field \mathbb{R}^* , numbers even smaller than the infinitesimals. The new hyperhyperreal field can be used, with benefits, for analysis over the hyperreals. On the other hand, ultrapowers also have other independent uses: they are important concepts in the study of Banach spaces, for instance.

The other main part of this work has dealt with the foundational development of real analysis in Isabelle, using nonstandard techniques. The approach used for mechanization of the calculus has proceeded strictly through definitions. We have introduced standard and nonstandard definitions of all the concepts formalized, and have proved their equivalence in each case.

We have also compared various aspects of our mechanization with corresponding ones from the formalization of real analysis by Harrison using standard techniques. We have highlighted the advantages that the more algebraic and often more intuitive nonstandard formulation of familiar concepts has over the standard approach. There is much scope for extending this development of real analysis in Isabelle. Some recent work, not covered in this paper, has involved developing Isabelle theories for power series and transcendental functions such as \exp , \sin , and \cos , using a combination of standard and nonstandard analysis techniques. This latest development points to another strength of our framework: it provides the option of both standard and nonstandard proof development. One may choose either to work with purely standard concepts, or to use mainly nonstandard ones, or a combination of both.

In summary, this work describes a rigorous investigation of the mechanization of analysis using nonstandard techniques. Our main aim has been to show that there are advantages to be gained by using nonstandard analysis as the framework for mechanized real analysis. The simplicity of the formulations and the ease with which many different results are mechanized justify the promises held by the approach.

Acknowledgements This work was carried out while the first author was a member of the Computer Laboratory. Support from the ORS and Cambridge Commonwealth Trusts is gratefully acknowledged. We thank James Margetson for several useful suggestions made during the development of this work, and for his comments on this paper.

Appendix A. *Isabelle theory files*

This appendix contains some of the theory files for the development of nonstandard analysis described in this paper.

The material is to be found at

<http://www.lms.ac.uk/jcm/3/lms1999-027/appendix-a/>.

The files should be used with Isabelle99, the current release of the theorem-prover. As this work evolves, up-to-date versions of the theory files will be available in the online Isabelle distribution [18].

References

1. J. R. ABRIAL and G. LAFFITTE, ‘Towards the mechanization of the proofs of some classical theorems of set theory’, Preprint, February 1993. 149
2. A. M. BALLANTYNE and W. W. BLEDSOE, ‘Automatic proofs of theorems in analysis using nonstandard analysis’, *J. ACM* 24 (1977) 353–374. 183, 186
3. T. BEDRAX, ‘Infmal: prototype of an interactive theorem prover based on infinitesimal analysis’, Master’s thesis, Pontifica Universidad Catolica de Chile, 1993. Liciendo en Mathematica con Mencion en Computation Thesis. 186
4. M. BEESON, ‘Using nonstandard analysis to ensure the correctness of symbolic computations’, *Internat. J. Found. Comput. Sci.* 6 (1995) 299–338. 186
5. J. C. BURKILL and H. BURKILL, *A second course in mathematical analysis* (Cambridge University Press, 1970). 172
6. R. CHUAQUI and P. SUPPES, ‘Free-variable axiomatic foundations of infinitesimal analysis: a fragment with finitary consistency proof’, *J. Symbolic Logic* 60 (1995) 122–159. 186
7. A. CHURCH, ‘A formulation of the simple theory of type’, *J. Symbolic Logic* 5 (1940) 56–68. 142
8. J. H. CONWAY, *On numbers and games* (Academic Press Inc. (London) Ltd, 1976). 144, 186
9. P. J. DAVIS and R. HERSH, *The mathematical experience* (Harmondsworth, Penguin, 1983). 142
10. J. D. FLEURIOT and L. C. PAULSON, ‘A combination of geometry theorem proving and nonstandard analysis, with application to Newton’s *Principia*’, *Automated deduction – CADE-15* (ed. C. Kirchner and H. Kirchner), Lecture Notes in Artificial Intelligence 1421 (Springer-Verlag, 1998) 3–16. 187
11. A. M. GLEASON, *Fundamentals of abstract analysis*, Series in Mathematics (Addison-Wesley, 1966). 186
12. M. GORDON and T. MELHAM, *Introduction to HOL: a theorem proving environment for higher order logic* (Cambridge University Press, 1993). 141
13. JOHN HARRISON, ‘Constructing the real numbers in HOL’, *Proceedings of the IFIP TC10/WG10.2 International Workshop on Higher Order Logic Theorem Proving and its Applications* (ed. L. J. M. Claesen and M. J. C. Gordon), IFIP Trans. A: Comput. Sci. Tech., vol. A-20 (North-Holland, IMEC, Leuven, Belgium, 1992) 145–164. 147

14. JOHN HARRISON, *Theorem proving with the real numbers* (Springer-Verlag, 1998). Also published as Technical Report 408 of the Computer Laboratory, University of Cambridge, 1996. 142, 144, 147, 167, 168, 168, 172, 175, 177, 186, 187
15. J. M. HENLE and E. M. KLEINBERG, *Infinitesimal calculus* (The MIT Press, 1979). 147, 166, 187
16. R. F. HOSKINS, *Standard and nonstandard analysis*, Math. Appl. (Ellis Horwood Limited, 1990). 148, 150, 162, 163, 164, 166, 183, 184
17. A. E. HURD and P. A. LOEB, *An introduction to nonstandard real analysis*, Pure Appl. Math. 118 (Academic Press Inc., Boston, MA, 1985). 151, 162, 162, 168, 177, 180
18. ‘Isabelle-online’,
<http://www.cl.cam.ac.uk/Research/HVG/Isabelle/library/HOL/HOL-Real/index.html>
188
19. L. S. JUTTING, ‘Checking Landau’s “Grundlagen” in the Automath system’, Ph.D. thesis, Eindhoven University of Technology, 1977. 186
20. H. J. KEISLER, *Foundations of infinitesimal calculus* (Prindle, Weber & Schmidt, 1976). 143, 150, 165, 166, 166
21. E. LANDAU, *Foundations of analysis* (Chelsea, 1951). 186
22. D. LAUGWITZ, ‘Infinitely small quantities in Cauchy’s textbooks’, *Historia Math.* 14 (1987) 258–274. 172, 173
23. T. LINDSTRØM, ‘An invitation to nonstandard analysis’, *Nonstandard analysis and its applications* (ed. N. Cutland), London Math. Soc. Student Texts 10 (Cambridge University Press, 1988) 1–105. 166
24. T. NIPKOW and D. VON OHEIMB, ‘Java_{light} is type-safe — definitely’, *Proc. 25th ACM Symp. Principles of Programming Languages* (ACM Press, New York, 1998) 161–170. 142
25. L. C. PAULSON, *Isabelle: a generic theorem prover*, Lecture Notes in Comput. Sci. 828 (Springer, 1994). 141, 141, 141, 142, 147
26. L. C. PAULSON, ‘The inductive approach to verifying cryptographic protocols’, *J. Computer Security* (1998) 85–128. 142
27. L. C. PAULSON, ‘Isabelle’s object-logics’, Tech. Rep. 286, Computer Laboratory, University of Cambridge (February 1998). 146, 173
28. L. C. PAULSON and K. GRĄBCZEWSKI, ‘Mechanizing set theory: cardinal arithmetic and the axiom of choice’, *Journal of Automat. Reason.* 17 (1996) 291–323. 149, 149
29. A. ROBINSON, *Non-standard analysis* (North-Holland, 1980). 140
30. E. SCHECHTER, *Handbook of analysis and its foundations* (Academic Press, 1997). 147, 148, 150, 150, 187
31. A. P. SIMPSON, ‘The Infidel is innocent’, *Math. Intelligencer* 12 (1990) 43–51. 151, 166
32. K. D. STROYAN and W. A. J. LUXEMBURG, *Introduction to the theory of infinitesimals* (Academic Press, 1976). 160
33. R. VESLEY, ‘An intuitionistic infinitesimal calculus’, *Constructive Mathematics* (ed. F. Richman), Lecture Notes in Math. 873 (Springer, 1983). 143, 143

Mechanizing NSA in Isabelle

Jacques D. Fleuriot jdf@dai.ed.ac.uk

Institute for Representation and Reasoning
Division of Informatics
University of Edinburgh

Lawrence C. Paulson lcp@cl.cam.ac.uk

Computer Laboratory
University of Cambridge