Constructing the real numbers in HOL

John Harrison

University of Cambridge Computer Laboratory, New Museums Site, Pembroke Street, Cambridge, CB2 3QG, England.

Abstract

This paper describes a construction of the real numbers in the HOL theorem-prover by strictly definitional means using a version of Dedekind's method. It also outlines the theory of mathematical analysis that has been built on top of it and discusses current and potential applications in verification and computer algebra.

Keywords: Mathematical Logic; Deduction and Theorem Proving

1 The real numbers

For some mathematical tasks, the natural numbers $\mathbb{N} = \{0, 1, 2, \ldots\}$ are sufficient. However for many purposes it is convenient to use a more extensive system, such as the integers ($\mathbb{Z}$) or the rational ($\mathbb{Q}$), real ($\mathbb{R}$) or complex ($\mathbb{C}$) numbers. In particular the real numbers are normally used for the measurement of physical quantities which (at least in abstract models) are continuously variable, and are therefore ubiquitous in scientific applications.

1.1 Properties of the real numbers

We can characterize the reals as the unique 'complete ordered field'. More precisely, the reals are a set $\mathbb{R}$ together with two distinguished constants $0 \in \mathbb{R}$ and $1 \in \mathbb{R}$ and the operations

$$\begin{array}{rl} + & : \mathbb{R} \times \mathbb{R} \to \mathbb{R} \\ . & : \mathbb{R} \times \mathbb{R} \to \mathbb{R} \\ - & : \mathbb{R} \to \mathbb{R} \\ inv & : \mathbb{R} - \{0\} \to \mathbb{R} \end{array}$$

having all the properties in the list given below. In what follows we use the more conventional notation xy for $x.y$ and x^{-1} for $inv(x)$. The use of such symbolism, including 0 and 1, is not intended to carry any connotations about what the symbols actually denote.

Firstly, the structure is nontrivial, i.e. has more than one element (note that the other properties below do not exclude this possibility):

- $1 \neq 0$

The reals form an abelian group under addition; i.e. addition is commutative and associative, there is an additive identity, and every element has an additive inverse:

- $\forall x\ y.\ x + y = y + x$
- $\forall x\ y\ z.\ x + (y + z) = (x + y) + z$
- $\forall x.\ 0 + x = x$
- $\forall x.\ (-x) + x = 0$

The *nonzero* reals form an abelian group under multiplication (of course 0 does not have a multiplicative inverse):

- $\forall x\ y.\ xy = yx$
- $\forall x\ y\ z.\ x(yz) = (xy)z$
- $\forall x.\ 1x = x$
- $\forall x.\ (x \neq 0) \Longrightarrow (x^{-1}x = 1)$

Addition and multiplication are related by the distributive law. Together with the above properties, this shows that the real numbers form a *field*.

- $\forall x\ y\ z.\ x(y + z) = xy + xz$

The reals are totally ordered by $<$ (i.e. the order is connected, transitive and irreflexive):

- $\forall x\ y.\ (x = y) \vee x < y \vee y < x$
- $\forall x\ y\ z.\ x < y \wedge y < z \Longrightarrow x < z$
- $\forall x.\ x \not< x$

The ordering relation interacts with the arithmetic operations in the following manner:

- $\forall y\ z.\ y < z \Longrightarrow \forall x.\ x + y < x + z$
- $\forall x\ y.\ 0 < x \wedge 0 < y \Longrightarrow 0 < xy$

All the above properties are also true of the rationals. The property which sets the reals apart is generally referred to as *completeness*, and can be stated in many equivalent forms. Perhaps the simplest is the *supremum property* which states that any nonempty set of reals which is bounded above has a least upper bound (supremum).

- $\forall S.\ (\exists x.\ x \in S)\ \wedge\ \ (\exists M.\ \forall x \in S.\ x \leq M)$
  $\Longrightarrow \exists m.\ (\forall x \in S.\ x \leq m)\ \wedge (\forall m' < m.\ \exists x \in S.\ x > m'))$

(Here we are using $\leq$ and $>$ as abbreviations for the obvious equivalents in terms of $<$ and $=$.) For example, the two sets $\{x \in \mathbb{R} \mid x^2 \leq 2\}$ and $\{x \in \mathbb{R} \mid x^2 < 2\}$ *both* have a supremum of $\sqrt{2}$, although one of the sets contains $\sqrt{2}$, as a maximum element, and the other does not.

1.2 Uniqueness of the real numbers

In higher-order logic, which allows the second-order quantification necessary to express completeness, the above properties determine the reals uniquely up to isomorphism, *provided* the inverse is defined only on the subset of nonzero reals. (For a proof, see [7], [9] or [21].) The role of partial functions in mathematics is rather more obscure and difficult than is often appreciated; we shall have more to say about this later. HOL functions are total, so making the inverse a partial function is highly inconvenient in our formalization. We therefore have an extra 'degree of freedom' that does not exist traditionally: we can define 0^{-1} to be whatever we like. To avoid surprises we keep it 'undefined' (effectively $\varepsilon x.\ F$), but it is important to realize that this is not the same as true undefinedness. For example we can prove $\forall x.\ 0x = 0$ and so in particular, assuming we define $x/y = xy^{-1}$ (as we do),

$$0/0 = 0$$

Conversely, a theorem true classically which is *not* true in our framework is:

$$\forall x \in \mathbb{R}.\ (tan(x) = 0) \Longrightarrow \exists n \in \mathbb{Z}.\ x = n\pi$$

because we cannot exclude the possibility that $cos(x)^{-1}$ is zero at odd multiples of $\pi/2$, in other words that $0^{-1} = 0$. There is even something to be said for *defining* $0^{-1} = 0$ since this makes the inverse into a bijection, so things like the following are true universally:

$$\forall x.\ (x^{-1})^{-1} = x$$

$$\forall x.\ 0 < x \equiv 0 < x^{-1}$$

We feel these issues are unlikely to present problems in practice, because division by zero is normally treated as a special case anyway, but one should be aware of them.

2 Constructing the real numbers

One approach to defining the reals in HOL is simply to introduce a new type `real` and the various operators on that type, then assert the above properties as axioms. However it is traditional to extend HOL only definitionally, to guarantee that consistency is preserved without employing any metalogical reasoning (except for ML programming, which at least enforces *recursive* metalogical reasoning).

Therefore we seek a way of creating a structure with the above properties out of previously defined objects. This problem has been solved by mathematicians in various different ways. Some leap in a single step from the natural numbers, others involve various intermediate stages such as the rational numbers. The following diagram shows a lattice of number systems under inclusion; the superscripted $\mathbb{Q}^+$ and $\mathbb{R}^+$ denote the positive or non-negative elements of $\mathbb{Q}$ and $\mathbb{R}$ respectively.

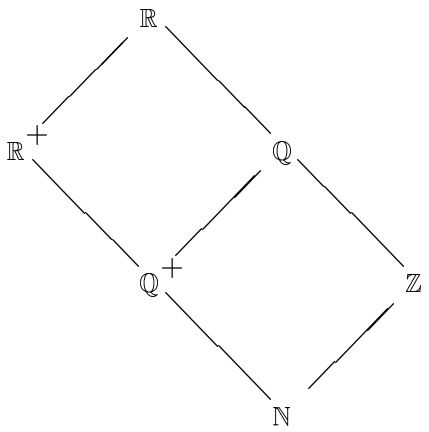

2.1 Straight from naturals (ℕ) to reals (ℝ)

Perhaps the most obvious approach is to model the real numbers by infinite positional (e.g. binary or decimal) sequences. It is necessary to take into account the fact that the representation is not unique (for example $0.99999\ldots$ and $1.00000\ldots$ both represent the same real number).

The definition of the ordering relation is reasonably straightforward, but addition is harder because it involves 'carries'. Nevertheless, a workable definition is carried through in [4]. An alternative, which is explored in [11], is to extend to sequences of arbitrary integers (not just those less than some base). The argument is that because of the non-uniqueness noted above, some form of equivalence relation would probably be used anyway, so we may as well avoid the problem of carries.

In either case, multiplication is harder, since it has a much more complicated relationship with the position of digits. One solution [7] is to define multiplication of terminating sequences (i.e. those which are zero beyond a certain point) and extend it to all sequences by a limiting process. But this is rather ugly and complicated. A very elegant alternative is proposed by Behrend [1]. He proves that a set $\mathbb{R}^+$ with operation $+ : \mathbb{R}^+ \times \mathbb{R}^+ \to \mathbb{R}^+$ and relation $<: \mathbb{R}^+ \times \mathbb{R}^+ \to \mathit{bool}$ obeying a few basic algebraic laws will turn out to be isomorphic with the strictly positive reals. In particular, for each $x \in \mathbb{R}^+$ there is a unique automorphism (i.e. bijection from $\mathbb{R}^+$ onto itself which respects the existing algebraic structure) x^* which maps $x^* : 1 \mapsto x$. It is now possible to define multiplication by

$$xy = y^*(x^*(1))$$

and prove all its required properties.

A more radical way of avoiding intermediate steps is to construct an extremely general number system using games, as explained by Conway [10]. In many ways this is a simple approach, but the recursive definitions of the operations seem hard to formalize in HOL.

2.2 From naturals ($\mathbb{N}$) to integers ($\mathbb{Z}$)

There are various possible representations of the integers in terms of the natural numbers, such as:

- A pair consisting of a boolean 'sign bit' and a natural number. For example $(true, 1)$ might represent $+1$ and $(false, 2)$ represent -2.

- A pair of natural numbers, where one imagines (m, n) standing for $m - n$ in the integers. Thus $(1, 0)$ represents $+1$ and $(5, 7)$ represents -2.

The main problem with both the above is non-uniqueness. Manifestly, every number has an infinity of representatives in the second case; $+1$ could equally well be represented by $(2, 1)$, $(3, 2)$, $(4, 3)$ and so on. Less egregious is the first case, but there are still two representations of zero, $(false, 0)$ and $(true, 0)$. There are two natural ways round this problem:

- Consider only a minimal set of representatives, which are in some sense canonical. For example one might in the first case exclude $(false, 0)$, and in the second insist that one or both numbers of the pair be zero.

- Define an equivalence relation expressing the effective identity of sets of terms, and use the equivalence classes under this relation, rather than the representatives themselves, to construct the new type.

2.3 From integers ($\mathbb{Z}$) to rationals ($\mathbb{Q}$)

This stage is a particular case of a well-known construction in abstract algebra, constructing the *field of fractions of an integral domain*, an integral domain being a nontrivial commutative ring with the property that

$$\forall x\ y.\ (xy = 0) \Longrightarrow (x = 0)\ \vee\ (y = 0)$$

The procedure consists of considering pairs of integers, which one thinks of as the numerator and denominator of a fraction; it is necessary to exclude 0 from the possible denominators. Then one uses equivalence classes of this subset of pairs under the obvious 'cross multiplication' equivalence relation

$$(x, y) \sim (x', y') \equiv (xy' = x'y)$$

As with the path from $\mathbb{N}$ to $\mathbb{Z}$, we have the option of eschewing equivalence classes in favour of choosing canonical elements. The natural choice of canonical form would be to insist that the pair of elements be coprime, i.e. represent a cancelled fraction (though this is not available in a general integral domain).

2.4 From the rationals ($\mathbb{Q}$) to the reals ($\mathbb{R}$)

There are two well-established classical methods for constructing the reals from the rationals, which were published independently by Cantor and Dedekind, both in 1872. (Cantor's method was largely anticipated by Méray in 1869, and one can find precursors of Dedekind's method as far back as Eudoxus with his theory of proportion.)

2.4.1 Cantor's Method

This method identifies a real number with the set of all rational sequences that converge to it. To say that a sequence (s_n) converges to s, written $s_n \to s$ means:

$$\forall \epsilon > 0.\ \exists N.\ \forall n \geq N.\ |s_n - s| < \epsilon$$

This is no good as a definition, because it contains the limit itself, which may not be rational. However the following variant is equivalent (as can be shown after completing the construction):

$$\forall \epsilon > 0.\ \exists N.\ \forall m \geq N, n \geq N.\ |s_m - s_n| < \epsilon$$

(It does not matter that we will restrict ϵ to rational values, since $\mathbb{Q}$ is dense in $\mathbb{R}$, i.e. between any two distinct reals there is a rational.) A sequence which satisfies this property is called a *Cauchy sequence*.

The fact that two series (s_n) and (t_n) converge to the same limit can also be expressed without using the limit itself:

$$\forall \epsilon > 0.\ \exists N.\ \forall n \geq N.\ |s_n - t_n| < \epsilon$$

It is easy to see that this defines an equivalence relation on Cauchy sequences, and the real numbers can be defined as its equivalence classes. The arithmetic operations can be inherited from those of the rationals in a fairly natural way, although the supremum presents slightly more difficulty. A complete treatment is given by Thurston [22].

2.4.2 Completion of metric and uniform spaces

Cantor's method admits of abstraction to more general structures. Given any metric space, that is, a set equipped with a 'distance function' on pairs of points (see later for formal definition), the process can be carried through in essentially the same way. This gives an isometric (distance-preserving) embedding into a complete metric space, i.e. one where every Cauchy sequence has a limit.

Since generality and abstraction are to be striven for in mathematics, it seems desirable to regard the construction of the reals as a special case of this procedure. Taken literally, however, this is circular, since the distance returned by a metric is supposed to be real-valued! On the other hand if we move to the more general structure of a topological space, the procedure seems to have no natural counterpart, since the property of being a Cauchy sequence is not preserved by homeomorphisms. Consider the action of the function from the set of strictly positive reals onto itself which maps $x \mapsto 1/x$. Clearly this is a homeomorphism (under the induced topology given by the usual topology on

$\mathbb{R}$) but it maps the sequence of positive integers, which is not a Cauchy sequence, to a Cauchy sequence.

Nevertheless there is a suitable structure lying between a metric and a topological space in generality. This is a uniform space, which while not equipped with an actual notion of distance, has nevertheless a system of *entourages* which intuitively indicate that certain pairs of points are the *same* distance apart. The completion procedure can be extended in a natural way to show that any uniform space can be embedded in a complete one by a uniformly continuous mapping which has an appropriate universal property. (From a categorical perspective, the 'morphisms' natural to topological, uniform and metric spaces are respectively continuous, uniformly continuous and isometric.)

A topological group is a structure which is both a group and a (Hausdorff) topological space, such that the group operations are continuous. It is not hard to see that a topological group has enough structure to make it a uniform space, where addition amounts to a 'rigid spatial translation'. Bourbaki [3] constructs the reals by first giving the rational numbers a topology, regarding this topological group as a uniform space and taking its completion. Although elegant in the context of general work in various mathematical structures, this is too complicated *per se* for us to emulate.

2.4.3 Dedekind's method

Dedekind's method identifies a real number with the set of all rational numbers less than it. Once again this is not immediately satisfactory as a definition, but it is possible to give an equivalent definition without referring to the bounding real number. We shall call such a set a *cut*. The four properties required of a set C for it to be a cut are as follows:

1. $\exists x.\ x \in C$
2. $\exists x.\ x \notin C$
3. $\forall x \in C.\ \forall y < x.\ y \in C$
4. $\forall x \in C.\ \exists y > x.\ y \in C$

These state respectively that a cut is not empty, is not $\mathbb{Q}$ in its entirety, is a 'down set', and has no greatest element. Again the arithmetic operations can be inherited from $\mathbb{Q}$ in a natural way, and the supremum of a set of cuts is simply its union.

3 The choice

It is harder than it might appear at first sight to make the above rigorous. Defining the integers as sign/magnitude pairs excluding $(false, 0)$ means that even proving the associative law of addition for integers represents a considerable amount of work, because there are eight different cases according to the sign of each of the three integers. (Ostensibly at any rate; avoiding the case split requires careful exploitation of symmetry and the distributivity of negation over addition.)

Defining the rationals as pairs of integers is even worse. Using cancelled fractions seems a bad idea because the proofs of the elementary theorems require some nontrivial lemmas about coprimality and divisibility. It seems that the use of equivalence classes is better; essentially we would have to prove similar theorems anyway in the case of cancelled fractions, with the added complexity of canonicalization. However the use of equivalence classes is slightly harder than intuition would suggest, particularly since the equivalence relation is *not* an equivalence relation over the whole type, but only on the subset with nonzero denominators. To see this, observe that

$$(1, 2) \sim (0, 0) \ \wedge \ (0, 0) \sim (1, 3)$$

But clearly $(1, 2) \not\sim (1, 3)$, so transitivity fails globally. Therefore many of the proofs are hedged with conditions about membership of the admissible subset, which renders them more complicated. (An alternative which is slightly cleaner, but still tedious, is to define the 'subtype' of nonzero integers and prove the closure of the arithmetic operators under certain conditions.) Finally, when it comes to constructing the reals from the rationals, neither of the classical methods is very appealing.

- Cauchy's method requires the generation of a significant body of 'analytical' lemmas about the rational numbers. Moreover, the problem of defining equivalence relations over subsets of types would have to be faced, which as noted above is tiresome.

- When proving the axioms for the structure, Dedekind's method requires several case-splits according to the sign of variables. This is mainly because the product of two negative rationals is positive, so the natural definition of multiplication on cuts

  $$XY = \{xy \mid x \in X \ \wedge \ y \in Y\}$$

  does not work. The two cuts X and Y extend to $-\infty$, so there will exist products of these large and negative numbers which are arbitrarily large and positive. Therefore the set is not a cut.

  This difficulty is usually noted in sketch proofs given in books, but to carry through in detail the complicated case splits they gloss over would be extremely tedious.

A rather more sophisticated approach is suggested by Conway. Although the novel method for constructing the reals explained in [10] did not seem to be possible to formalize in HOL, his incidental discussion of the classical methods is perceptive. Firstly, he emphasizes the difficulty of constructing $\mathbb{R}$ from $\mathbb{Q}$ by Dedekind cuts:

> Nobody can seriously pretend that he has ever discussed even eight cases in such a theorem — yet I have seen a presentation in which one theorem actually had 64 cases ... Of course an elegant treatment will manage to discuss several cases at once, but one has to work very hard to find such a treatment.

He advocates instead following the path on the lattice diagram through $\mathbb{Q}^+$ and $\mathbb{R}^+$, at least if Dedekind's method is to be used. This avoids the case splits (otherwise it is

essentially the same as the signed case presented above), and as we shall see, has other advantages as well.

One apparent drawback of using this path is that we lose the potentially useful intermediate types $\mathbb{Z}$ and $\mathbb{Q}$. However this is not really so, because the method used to construct $\mathbb{R}$ from $\mathbb{R}^+$ can be used almost unchanged (and this is where a computer theorem prover scores over a human) to construct $\mathbb{Z}$ and $\mathbb{Q}$ from their positive-only counterparts.

Landau's book [19] is one of the few to present Dedekind's construction in more or less full detail. This also, significantly, uses the path through $\mathbb{Q}^+$ and $\mathbb{R}^+$. Landau takes the step from $\mathbb{R}^+$ to $\mathbb{R}$ by introducing zero and negative numbers, which is similar to the sign/magnitude representation, and with the same problems of case splitting. Conway anticipates this and prefers instead the use of equivalence classes of pairs of numbers. In fact there is then quite a close analogy with the construction of the rationals, with addition taking the role of multiplication.

4 The theory of semirings

To journey along the path through $\mathbb{Q}^+$ and $\mathbb{R}^+$, we must know which algebraic laws for these structures we are going to need to get $\mathbb{R}$ out at the other end, and whether we are going to include 0. We will refer to structures like $\mathbb{N}$, $\mathbb{Q}^+$ and $\mathbb{R}^+$ as *semirings*, although to some authors [12] this implies that they contain a zero.

Semirings can be characterized by relatively few axioms. Ordering does not have to be primitive; we can *define* an ordering as follows, whether or not the structure contains a zero:

$$x < y \equiv (x \neq y) \wedge (\exists d.\ y = x + d)$$

Investigation reveals that the following set of axioms is sufficient to allow the derivation of a nontrivial ordered ring by the method of equivalence classes of pairs of numbers under $(x, y) \sim (x', y') \equiv (x + y' = x' + y)$. They are chosen to be true whether or not the semiring contains a zero, which explains why the penultimate one looks a bit peculiar.

1. $\forall x\ y.\ x + y = y + x$
2. $\forall x\ y\ z.\ x + (y + z) = (x + y) + z$
3. $\forall x\ y.\ xy = yx$
4. $\forall x\ y\ z.\ x(yz) = (xy)z$
5. $\forall x\ y\ z.\ x(y + z) = xy + xz$
6. $\forall x.\ 1x = x$
7. $\forall x\ y.\ (x = y)\ \vee (\exists d.\ x = y + d)\ \vee (\exists d.\ y = x + d)$
8. $1 + 1 \neq 1$
9. $\forall x\ y\ z.\ (x + (y + z) = x) \Longrightarrow (x + y = x)$

10. $\forall x\ y\ z.\ (x + y = x + z) \Longrightarrow (y = z)$

It is not hard to see that the last three axioms can all be derived from the single axiom

$$\forall x\ y.\ x + y \neq x$$

This is a strong argument for not including a zero in the structure: we will have to prove fewer axioms as primitive for the semirings we construct. Further, it allows certain theorems such as the field axiom and the Archimedean property (see below) to be written in a simpler form.

There is unfortunately a problem: the standard HOL theory of natural numbers does contain a zero. We could have defined a new type of nonzero natural numbers, but that seemed rather wasteful, so instead we wrote a procedure which works whether or not the semiring contains a zero. The procedure requires the full list of axioms above. However where it is possible, it is easier to prove $\forall x\ y.\ x + y \neq x$ and derive the others from that in a general way.

There are a few extra axioms we need for particular semirings. To get from $\mathbb{R}^+$ to $\mathbb{R}$, we require a form of the supremum property for $\mathbb{R}^+$. And to get from $\mathbb{Q}^+$ to $\mathbb{R}^+$ we need to prove for $\mathbb{Q}^+$ both the field axiom

$$\forall x.\ x^{-1}x = 1$$

and also a form of the *Archimedean property*. This states that if we define a function

$$addn\ n\ x = x + \ldots + x$$

where there are n terms in the sum (of course if we are sloppy with types, regarding $\mathbb{N}$ as a 'subset' of the semiring, $addn\ n\ x$ is just nx), then the following is true

$$\forall x\ y.\ \exists n.\ addn\ n\ x > y$$

Note that neither of the above would be true if $\mathbb{Q}^+$ contained a zero; they would both be consequent on $x \neq 0$.

5 Equivalence relations

Every step we are to take in the lattice, with the sole exception of the line from $\mathbb{Q}^+$ to $\mathbb{R}^+$, involves constructing a set of equivalence classes. To make this easier, we wrote a procedure to automate it, given:

- A name for the new type
- A theorem asserting that a (2-place curried) relation (say R) is an equivalence relation, in the following simple form:

$$\forall x\ y.\ xRy \equiv (Rx = Ry)$$

(Here we are using infix notation for the first instance of R; we might write instead Rxy.)

Supposing we are constructing the integers from pairs of natural numbers as explained above, the relation would be defined by

$$(x_1, y_1)R(x_2, y_2) = (x_1 + y_2 = x_2 + y_1)$$

- A list of operations on the representatives together with the desired name of the corresponding operators over the equivalence classes. For example, we might give it an addition operation defined on pairs of numbers as follows:

$$(x_1, y_1) + (x_2, y_2) = (x_1 + x_2, y_1 + y_2)$$

- A list of theorems asserting that the operations on representatives are all well-defined, in the sense that, taking addition as our example again:

$$(x\ R\ x')\ \wedge\ (y\ R\ y') \Longrightarrow (x + y)\ R\ (x' + y')$$

  When the relevant argument or result is not of the representing type, equality takes the place of R.

- A list of theorems about the operations on representatives, e.g. the associative law:

$$x + (y + z) = (x + y) + z$$

The procedure first constructs a type of equivalence classes. The characteristic predicate required to select the set $\{R\ x\}$ is formally:

$$\lambda C.\ \exists x.\ C = R\ x$$

Next, the appropriate operations on the new type are defined. For example + gives rise to a new operator $+^*$ (we use the star consistently, but in fact the user specifies the name of the operator) on the equivalence relation as follows:

$$X +^* Y = R((\varepsilon x.\ x \in X) + (\varepsilon y.\ y \in Y))$$

In other words, pick using the ε operator arbitrary representatives of each equivalence class, operate on them and then take the equivalence class of the result. If arguments or result are not of the representing type, then we avoid picking representatives or applying R, respectively. For example, an 'is positive' predicate would be elevated as follows:

$$ispos^* X = ispos(\varepsilon x.\ x \in X)$$

and the $addn$ function mentioned above as

$$addn^*\ n\ X = R(addn\ n\ (\varepsilon x.\ x \in X))$$

Finally, the proof procedure tries to convert the theorems on representatives into theorems about the new type. This is the only part which requires R to be an equivalence relation and the operations to be well defined. In the case of the associative law, we get:

$$X +^* (Y +^* Z) = (X +^* Y) +^* Z$$

Our procedure does *not* work if we have to deal with a subset of the basic type. We could have generalized the procedure: one would also have to supply it with theorems expressing the closure of the operations with respect to this subset, and the other theorems may become conditional on the variables belonging to the subset. The extra complexity was not necessary for the task in hand. Furthermore, there are a few problems to be resolved. Consider a definition of multiplicative inverse for a field of fractions in an integral domain:

$$(x, y)^{-1} = (y, x)$$

Then one needs a condition $x \neq 0$ on the closure theorem. Dealing with things like this in a regular way seems quite awkward.

6 Details of the HOL construction

We now look in a little more technical details at how the various parts were implemented inside HOL.

6.1 From $\mathbb{N}$ to $\mathbb{Q}^+$

This is reasonably straightforward, but we do have to deal with the problem that the natural numbers contain 0 when we would rather they did not. The solution chosen was to use (x, y) to represent $(x + 1)/(y + 1)$.

The use of $y + 1$ avoids zero denominators without using subsets, which would defeat the equivalence class procedure. Using $x+1$ avoids including zero in the rational semiring, which is what we want, and also makes the proofs more symmetrical and regular. The only drawback is that the definitions of the operations are somewhat more complicated. For example the addition of (x_1, y_1) and (x_2, y_2) is defined as

$$((x_1 + 1)(y_2 + 1) + (x_2 + 1)(y_1 + 1) - 1, (y_1 + 1)(y_2 + 1) - 1)$$

The apparent extra difficulty of the proofs can be overcome by a simple tactic which manages to eliminate a lot of the above complexity.

6.2 From $\mathbb{Q}^+$ to $\mathbb{R}^+$

This is the most difficult part of the whole procedure. Although we are dealing with semirings, the Dedekind cuts procedure is essentially identical to the full case, but includes none of the complicated case splits. We define the operations on cuts as follows (taking some liberties with the HOL notation):

- $sup\ S = \bigcup S$

- $X + Y = \{x + y \mid x \in X \ \wedge \ y \in Y\}$
- $XY = \{xy \mid x \in X \ \wedge \ y \in Y\}$
- $X^{-1} = \{w \mid \exists d < 1.\ \forall x \in X.\ wx < d\}$

Only the last of these is unobvious; the more natural definition is

$$X^{-1} = \{w \mid \forall x \in X.\ wx < 1\}$$

However this would mean that unless a cut denoted an irrational real, its inverse would not be a cut, e.g:

$$\{x \in \mathbb{Q}^+ \mid x < 1\ \}^{-1} = \{x \in \mathbb{Q}^+ \mid x \leq 1\ \}$$

The construction then consists of proving closure, i.e. showing that when applied to cuts, all the operations yield a cut (including the trivial instance of proving that the set representing real 1 is a cut) and that all the required axioms hold. These proofs, while mostly routine, are sometimes quite long, so it is not possible to discuss them all here. We sketch only the proof of the axiom

$$\forall X.\ X^{-1}X = 1$$

which is a fairly representative example. In the above and what follows, we assume X denotes a cut, without stating it explicitly; in fact in the HOL proofs X becomes $cut\ X$, where cut is the type bijection from the real number type to the set of rational cuts. This means that no explicit set constraint is necessary since $cut\ X$ is *always* a cut. Firstly we need the following sequence of lemmas.

Lemma 1

$$\forall X\ x\ y.\ x \in X \wedge y \notin X \Longrightarrow x < y$$

This follows easily from cut property 3.

Lemma 2

$$\forall X\ x\ y.\ x \notin X \wedge x < y \Longrightarrow y \notin X$$

Also a straightforward consequence of cut property 3.

Lemma 3 This states that we can get arbitrarily close to the "top" of a cut.

$$\forall X\ e.\ \exists x.\ x \in X \wedge x + e \notin X$$

To prove this, choose any $x_0 \in X$, and $x_1 \notin X$ (this is possible by virtue of cut properties 1 and 2). Then consider the sequence of rationals

$$\{x_0 + addn\ n\ e \mid n \in \mathbb{N}\}$$

It is a simple consequence of the Archimedean property that there is an n with

$$x_0 + addn\ n\ e > x_1$$

Then from Lemma 2 and the wellfoundedness of the naturals, we know there is a *least* n such that $x_0 + addn\ n\ e \notin X$, say k. Further, $k \neq 0$ because we know $x_0 \in X$. Consequently

$$x = x_0 + addn\ (k-1)\ e$$

has the required property. (Strictly, *addn* is not defined at 0 because $\mathbb{Q}^+$ has no zero, but the reasoning is essentially the same as the above.)

Lemma 4 This is just a multiplicative rather than additive version of Lemma 3.

$$\forall X\ u.\ u > 1 \Longrightarrow \exists x.\ x \in X \wedge ux \notin X$$

To prove this, choose any $x_0 \in X$, possible by cut property 1. If $ux_0 \notin X$, then we are finished. Otherwise let

$$e = x_0(u-1)$$

so by Lemma 3, we can pick an x with

$$x \in X \wedge x + e \notin X$$

But now we have

$$x_0 + x_0(u-1) \in X \wedge x + x_0(u-1) \notin X$$

Consequently, using Lemmas 1 and 2, together with simple properties of $<$, we deduce that $x_0 < x$, and therefore $x + x(u-1) \notin X$, as required.

Main theorem

$$\forall X.\ X^{-1}X = 1$$

Translated into cuts, we want to show

$$\{wx \mid x \in X \wedge \exists d < 1.\ \forall y \in X.\ wy < d\} = \{z \mid z < 1\}$$

Expressed formally in HOL, this means establishing the following logical equivalence

$$\forall z.\ (\exists w\ x.(z = wx) \wedge x \in X \wedge (\exists d < 1.\ \forall y \in X.\ wy < d)) = z < 1$$

This reduces to two implications. A little thought will show that the left to right implication is straightforward. For the other, suppose $z < 1$. Then it is a simple property of the rationals that we can find a d such that $z < d$ and $d < 1$ (for example $d = (z+1)/2$). Therefore $z^{-1}d > 1$. By Lemma 4, we can choose an x such that

$$x \in X \wedge z^{-1}dx \notin X$$

Now let $w = zx^{-1}$. Then w, x and d are going to be witnesses for the correspondingly named existentially quantified variables. It remains only to prove

$$\forall y \in X.\ wy < d$$

But if $y \in X$, then by Lemma 1, $y < z^{-1}dx$; and $z = wx$, so we have $y < w^{-1}x^{-1}dx$, i.e. $wy < d$, as required.

6.3 From $\mathbb{R}^+$ to $\mathbb{R}$

Most of this is the same as the construction of $\mathbb{Z}$ and $\mathbb{Q}$ from their half-counterparts. The extra theorem that has to be transferred across is the supremum property. The supremum property for the positive reals states that

> Every nonempty set of positive reals which is bounded above has a supremum.

The first step is to transfer this result to the type of reals. Although not vacuous (formally, the positive reals are a completely different type), this is straightforward because the type bijections define an isomorphism between the *type* of positive reals and the positive elements of the type of reals. The theorem now becomes

> Every nonempty set of real numbers which is bounded above, and all of whose elements are strictly positive, has a supremum.

We generalize this in two stages. Firstly it is simple to prove the following strengthening:

> Every nonempty set of real numbers which is bounded above, and which contains at least one strictly positive element, has a supremum.

(The property 'nonempty' is actually superfluous here, but we keep it in for regularity.) This follows because l is a supremum of the whole set if and only if it is a supremum of the strictly positive elements of it, since any positive number is greater than any negative number.

Finally we prove the lemma that for any d, positive or negative, l is a supremum of S if and only if $l + d$ is of $\{x + d \mid x \in S\}$. Now this can be used to reduce the case of any nonempty bounded set to the above, by choosing a d to 'translate' it such that it has at least one strictly positive element. We now have the full result:

> Every nonempty set of real numbers which is bounded above has a supremum.

7 Interface issues

It is desirable to use the normal arithmetic symbols like + both for the real numbers and the natural numbers (not to mention the integers, the rationals, and other uses such as binary summation in process calculi). At present, HOL does not have a mechanism to support overloading of constant names, but it is hoped to add something in the future.

At present the interface map feature is used to allow the user to swap easily between using + etc. for different purposes, but it is not possible to use the same symbols for different operations in the same term. Furthermore, a different notation is required for real number constants. At present the ampersand is used as an interface map for the 'inclusion' function $\iota : \mathbb{N} \to \mathbb{R}$, so the real constants can be written &0, &1 etc. This is reminiscent of programming languages like C and Standard ML where the floating-point constants are distinguished syntactically from integer ones by being written 1.0 or $1e1$ etc. One could in fact fix the HOL parser to allow just such notation; this is largely a matter of personal taste.

8 Building on the real number axioms

To make the reals useful as a library, we have built a fair amount of theory on top of the basic construction. This includes a large number of algebraic lemmas, and enough mathematical analysis to define and prove the main properties of the transcendental functions like *exp* and *sin*. We will outline the main results.

8.1 Topology and metric spaces

A *metric* on a set S is an abstraction of the notion of distance between a pair of points in ordinary Euclidean space $\mathbb{R}^n$. A function $\rho : S \times S \to \mathbb{R}$ is a metric iff it has the following properties:

1. $\forall x\ y.\ \rho(x, y) = 0 \equiv x = y$

2. $\forall x\ y\ z.\ \rho(y, z) \leq \rho(x, y) + \rho(x, z)$

The latter is usually called the *triangle law*; in the usual metric on $\mathbb{R}^2$ it states that one side of a triangle is no greater than the sum of the other two sides. From the above two properties, it follows quite easily that a metric is both nonnegative ($\forall x\ y.\ \rho(x, y) \geq 0$) and symmetric ($\forall x\ y.\ \rho(x, y) = \rho(y, x)$).

In subsequent developments we invariably use the usual metric on $\mathbb{R}$, namely:

$$\rho(x, y) = |x - y|$$

but many of the theorems we prove are true in a more general framework. (In fact some are proved in the even more general structure of a topological space.) Since on *any* set we can define the *discrete* metric:

$$\rho(x, y) = (x = y) \to 0 \mid 1$$

we can define a type operator $(\alpha)metric$. (HOL insists that all types be inhabited.)

8.2 Convergence nets

Nets generalize the notions of sequences and pointwise limits. They are simply functions out of a set with a directed partial order, i.e. one where for any x and y there is a z with $z \geq x$ and $z \geq y$. Together with their now more popular relatives, filters, they are important in their own right since many results about metric spaces remain true in arbitrary topological spaces if we consider nets, rather than merely sequences. For example, a generalization of the Bolzano-Weierstrass theorem is as follows: a set is compact iff every net has a limit point.

Our use of nets is quite prosaic: they avoid proving twice various theorems about combining limits both for the limits of sequences of reals, and for pointwise limits of functions $\mathbb{R} \to \mathbb{R}$. An example is the theorem which asserts 'the limit of a sum is the sum of the limits', i.e:

$$x \to l \wedge y \to m \implies (\lambda n.\ x_n + y_n) \to (l + m)$$

From the general net theorems, the two cases we are interested in can be derived by instantiating the partial order as follows:

- The usual order $\geq$ on the natural numbers.
- Closeness under a metric to the limit, i.e. $x \geq x'$ if $\rho(x, x_0) < \rho(x', x_0)$. In fact reverse inclusion of neighbourhoods works in an arbitrary topological space, but we do not need such generality.

More details of the theory of nets are given in the classic book by Kelley [17] and many more modern books on general topology.

8.3 Sequences and series

Firstly some net theorems are specialized to sequences, yielding various combining theorems for sequences. Further theorems specifically about sequences are then proved, including the following:

- A bounded and monotonic sequence converges. Suppose the sequence is increasing, the other case being analogous. Consider the set $\{x_n | n \in \mathbb{N}\}$. This must have a supremum l such that for any $\epsilon > 0$, there exists an N with $|x_N - l| < \epsilon$. But because the sequence is increasing, this means that $\forall n \geq N.\ l - \epsilon < x_n \leq l$, so the sequence in fact converges to l.
- Every sequence has a monotonic subsequence. Call n a *terrace point* if we have $\forall m > n.\ x_m \leq x_n$. If there are infinitely many such terrace points, we can just form a decreasing sequence by successively picking them. If on the other hand there are only finitely many terrace points, then suppose N is the last one (or $N = 0$ if there are none). Now for *any* $n > N$, there is an m with $x_m > x_n$ (otherwise n would be a terrace point). Hence we can choose a (strictly) increasing subsequence by repeatedly making such choices.

- Every Cauchy sequence converges. A Cauchy sequence is bounded because for any $\epsilon > 0$, say $\epsilon = 1$, we can find an N such that $\forall n \geq N,\ m \geq N.\ |x_m - x_n| < \epsilon$, so

$$max(x_0, \ldots, x_{N-1}, x_N + \epsilon)$$

  is an upper bound. Hence we can find a subsequence which is both bounded and monotonic, and hence convergent. But now because of the Cauchy criterion, the limit of the subsequence is in fact a limit for the sequence itself.

We also prove a very useful general principle, codifying Bolzano's notion of proof by bisection. Suppose we want to establish that a property holds for an interval $[a, b]$ (usual notation for $\{x | a \leq x \wedge x \leq b\}$). Suppose that this property is such that when it is true of two adjacent intervals, it is true for the combined interval. Then the principle of bisection states that it is sufficient to prove that it is true for any sufficiently small interval containing any given point of the original interval.

For the proof, suppose the property is false for some interval. Then we can divide the interval in half, and the property must fail for one of the halves (otherwise by the composition property it would be true of the whole interval). Picking the interval (or one of the intervals) where it is false, we can repeat the process of bisection. In this way we get a decreasing nest of intervals. Since both sets of successive endpoints form monotone sequences, it is not hard to show that there is precisely one point common to all of them. But we know the property is true of all sufficiently small intervals surrounding this point, which gives a contradiction. The formal statement of the principle is as follows:

$$\begin{array}{ll} \forall P. & (\forall a\ b\ c.\ a \leq b \wedge b \leq c \wedge P(a,b) \wedge P(b,c) \Longrightarrow P(a,c)) \wedge \\ & (\forall x.\ \exists \delta.\ 0 < \delta \wedge (\forall a\ b.\ a \leq x \wedge x \leq b \wedge (b-a) < \delta \Longrightarrow P(a,b)) \\ & \Longrightarrow \forall a\ b.\ a \leq b \Longrightarrow P(a,b) \end{array}$$

Next, we move on to infinite series, defined as limits of a finite sums. Additional theorems proved are mainly tests for convergence such as the comparison test and ratio test.

8.4 Limits, continuity and differentiability

Once again we specialize the net theorems to give various combining theorems for pointwise limits. Next we define the notion of continuity; a function f is continuous at a point x when, as $h \to 0$,

$$(\lambda h.\ f(x+h)) \longrightarrow f(x)$$

and proceed to prove some of the classic theorems of elementary real analysis:

- A function continuous on a closed interval is bounded. This can be proved by bisection, since boundedness obviously has the required composition property, and the boundedness for sufficiently small regions follows immediately from continuity.

- A function continuous on a closed interval attains its supremum and infimum. The following slick proof is taken from [5]. Suppose f does not attain its supremum M. Then the function defined by $\lambda x.\ (M - f(x))^{-1}$ is continuous on the interval (a previous theorem about continuity assures us of this because the denominator is never zero), and therefore *it* is bounded, by K say, which must be strictly positive. But this means that we have $M - f(x) \geq K^{-1}$, which is a contradiction because M is a *least* upper bound.

- Rolle's theorem: if f is continuous for $a \leq x \leq b$ and differentiable for $a < x < b$, and in addition $f(a) = f(b)$, then there is some $a < x < b$ with $f'(x) = 0$. We know that f attains its bounds, and in fact its derivative must be zero there, otherwise it would exceed its bounds on one side or the other.

- The Mean Value Theorem states that if f is continuous for $a \leq x \leq b$ and differentiable for $a < x < b$, then there is some $a < x < b$ with $f(b) - f(a) = (b - a)f'(x)$. A proof is easy by applying Rolle's theorem to the function:

$$\lambda x.\ f(x) - (f(b) - f(a))x/(b - a)$$

- A function whose derivative is zero on an interval is constant on that interval. This is an immediate corollary of the Mean Value Theorem. Note that it can also be proved directly by bisection, using the property:

$$P(x, y) \equiv f(y) - f(x) \leq C(y - x)$$

- Bolzano's Intermediate Value Theorem. This states that for any continuous function f and interval $[a, b]$, if $f(a) < f(b)$, then for any y between $f(a)$ and $f(b)$ there is an x between a and b such that $f(x) = y$. Intuitively this says that if a continuous function starts below a horizontal line and ends above it, then it must cross the line (this corresponds well with the intuitive idea of continuity). This, or to be precise its contrapositive, is also proved by bisection. Suppose f is continuous on $[a, b]$ but never attains the value y. Then it is easy to see by bisection that y cannot lie between $f(a)$ and $f(b)$.

We then move on to defining differentiation in the usual manner. The various combining theorems are mostly straightforward; one exception is the chain rule. In Leibnizian notation the theorem is very suggestive:

$$\frac{dy}{dx} = \frac{dy}{du}\frac{du}{dx}$$

It would seem that to prove it we need simply observe that the above is true for finite differences, and consider the limit. However this does not work easily, because we have to consider the possibility that du may be zero even when dx is not; crudely speaking, the problem is that limits are not compositional. However continuity *is* compositional, and the theorem follows quite easily from the following alternative characterization of

differentiability, due to Carathéodory, namely that f is differentiable at x with derivative $f'(x)$ if there is a function g_x, continuous at x and with value $f'(x)$ there, such that

$$f(x') - f(x) = g_x(x)(x' - x)$$

for all x'. The equivalence with the usual definition is easy to establish.

We also prove results about the continuity and differentiability of inverse functions; the latter is also eased somewhat by using the Carathéodory definition [18].

We define next a useful piece of ML code, a function `DIFF_CONV`. This is a conversion which is given an expression denoting a function $\mathbb{R} \to \mathbb{R}$ and a list of known derivatives. It will return a theorem about the derivative of the given function, applying the chain rule, product rule etc. recursively, and automatically generating necessary conditions, such as nonzero denominators. This is very useful for reasoning about the transcendental functions, where we will want to differentiate some quite complicated expressions which would be tedious to do by hand.

8.5 Power series

Here we bring together the theories of infinite series and differentiability, proving a few results about power series, in particular that they are characterized by a 'circle of convergence' within which they can be differentiated term-by-term. This latter result was in fact the most difficult proof in the whole undertaking. Had we been developing analysis for its own sake, we would have proved some general results about uniform convergence. As it is, we prove the result by direct manipulation of the definition of derivative, following Theorem 10.2 in [6]. The theorem requires both the first and second formal derivative series to converge within the radius of convergence. This does in fact follow in general, but we did not need to prove it because the power series we are concerned with differentiate to 'each other', so we already have convergence theorems.

We also prove Taylor's theorem in its full infinite series form. In fact it is no longer used in the subsequent development, but could be useful for giving error bounds when truncating infinite series.

8.6 The transcendental functions

The functions *exp*, *sin* and *cos* are defined by their power series expansions (we do not need Taylor's theorem to do this):

$$\begin{aligned}
exp(x) &= 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \ldots \\
sin(x) &= x - \frac{x^3}{3!} + \frac{x^5}{5!} - \frac{x^7}{7!} + \ldots \\
cos(x) &= 1 - \frac{x^2}{2!} + \frac{x^4}{4!} - \frac{x^6}{6!} + \ldots
\end{aligned}$$

We show using the ratio test that the series for *exp* converges, and by the comparison test that the other two do. Now by our theorem about differentiating infinite series

term by term, we can show that the derivative of sin is cos, and so on. Furthermore, a few properties like $cos(0) = 1$ are more or less immediate from the series. Now we are in a position to prove more interesting properties. As an example, to show that $exp(x + y) = exp(x)exp(y)$, consider the function:

$$\lambda x.\ exp(x + y)exp(x)^{-1}$$

Our automatic conversion, with a little manual simplification, shows that this has a derivative which is 0 everywhere. Consequently, by a previous theorem, it is constant everywhere. But at $x = 0$ it is just $exp(y)$, so the result follows.

The addition formulas for sin and cos can also be proved in this way, but we need the slight trick of combining the two as follows to make the procedure work:

$$\begin{array}{ll}\lambda x. & (sin(x + y) - (sin(x)cos(y) + cos(x)sin(y)))^2 + \\ & (cos(x + y) - (cos(x)cos(y) - sin(x)sin(y)))^2\end{array}$$

By the same method as above we can prove that this function is always zero, and the addition formulas follow easily. Periodicity of the trigonometric functions follows from the addition formulas and the fact that there is a least $x > 0$ with $cos(x) = 0$. This latter fact is proved by observing that $cos(0) > 0$ and $cos(2) < 0$. The Intermediate Value Theorem tells us that there must therefore be a zero in this range, and since $sin(x)$ is positive for $0 < x < 2$, cos is strictly decreasing there, so the zero is unique. (These proofs involve some fiddly manipulations of the first few terms of the series for sin and cos.) The zero is of course $\pi/2$, and this serves as our definition of π.

The functions ln, asn, acs and atn are defined as the inverses of their respective counterparts exp, sin, cos and tan. Their continuity and differentiability (in suitable ranges) follow from the general theorems about inverse functions, with a bit of algebraic simplification.

9 Applications

There seem to be several promising areas of application, which have only been partially investigated so far.

9.1 Verification of floating-point hardware

This seems an ideal area for theorem-proving; it is hard to see how one could verify by model-checking a circuit to calculate logarithms, for example. We have already done a verification of a toy floating-point square root circuit. (By toy we mean that it uses a simple floating point format rather than the full IEEE [14] standard with special cases and denormalized numbers. Also our circuit is probably inefficient compared with a commercial design.) It is hoped in the future to do verifications of more realistic circuits and/or circuits for more complicated functions like sin.

9.2 Numerical work

It would be quite easy to program HOL to produce mathematical tables with high assurance (for human consumption or insertion of constants into hardware or software). After all, this is what Babbage designed his Difference Engine to do! More generally, there are many areas of application in the error analysis of numerical methods [23].

9.3 Computer Algebra

Computer algebra systems are widely used by applied mathematicians and others. In view of their complexity it seems likely that they include bugs, or consciously implement rules of a theoretically dubious nature.

Theorem proving offers two possible solutions. Firstly, we can implement a computer algebra system using a system like HOL as a rigorous base. Our differentiation conversion provides a (rather trivial) example of how this can be done – however to match the complicated algorithms and heuristics of commercial computer algebra systems would be an enormous undertaking.

An alternative is simply to link a theorem prover and computer algebra system, because many problems, such as solving equations, factorizing polynomials and finding integrals, require complicated methods, but the answers can be *checked* quite easily, so it is quite feasible to do only this part in the theorem prover.

These possibilities are explored, using the example of integration, in our forthcoming paper [13].

9.4 Hybrid Systems

Hybrid systems are those which involve both analogue and digital components, or both discrete and continuous models. These are of course very important when computers are used in real-world applications such as chemical plant controllers. Various formalisms for dealing with hybrid systems have been proposed [20], and it may well be useful to have a real numbers theory.

10 Conclusion and related work

As far as we are aware, the only previous construction of the classical reals in a computer theorem prover was by Jutting [16], who translated Landau's book [19] into Automath.

The reals can also be developed in a way which is 'constructive' in the sense of Bishop [2]. The usual construction is an elaboration of Cauchy's method where the rate of convergence of a Cauchy sequence is bounded explicitly. The resulting objects do not enjoy all the properties of their classical counterparts; for example $\forall x\ y.\ x \leq y \vee y < x$ is not provable.

The definition of the constructive reals has been done in NuPRL, with a proof of their completeness, i.e. that every Cauchy sequence converges [8]. Much of the construction, as well as some work on completing a general metric space, has been done in the LEGO prover [15], which is also based on a constructive logic.

The full construction described here, from $\mathbb{N}$ to $\mathbb{R}$, took about two weeks, but it would have taken much longer without careful selection of strategy. In particular, it seems that a quotient procedure tends to be much easier than picking canonical elements. Previous constructions of the integers from the naturals have been made by others using canonical representations, and their greater complexity seems to bear out this point.

The additional work on mathematical analysis took several months, on and off. Analytical proofs tend to have quite a lot of minor details which need to be filled in, particularly tedious bits of arithmetic reasoning. In the near future, derived decision procedures will become available which will greatly ease this sort of task without compromising the security of the system.

The fact that the HOL system can be programmed easily (the presence of ML rather than just an ad-hoc macro language, together with the simplicity of the underlying term structures) is a major advantage. Implementing procedures like the quotient types function and the differentiation conversion would otherwise be very difficult. Furthermore, even though some proofs are long and tedious, one can always get there with a little patience because of the system's great flexibility.

Here are some indicators of the 'size' of the proof. The complete theories described here generated 167608 primitive inferences, and took 92 minutes to build on a 48Mb SPARCserver. The total ML source is 10080 lines, including comments. The parts leading just to the real number "axioms" generated 49017 primitive inferences, took 14 minutes to build, and consisted of 2098 lines of ML.

The nature of the underlying logic has some impact on the formalization. We have already discussed the effect of total functions on the division operation. They also mean that certain traditional notations are less useful; for example we cannot infer from $lim(x_n) = l$ that the sequence (x_n) actually tends to a limit; the lim function is *always* defined. Instead, we tend to use relational notations like $x \to l$ in preference. In most cases this is no handicap; indeed it is often clearer.

If analysis were to be taken further, some extensions to the logic would be convenient. For example, it is difficult to reason in a clean way about arbitrary n-ary Cartesian products without some simple form of dependent types. Furthermore, subtypes would allow more transparent embedding of one number system in another. It is perhaps difficult to know how difficult some parts of mathematics are to formalize without actually trying. Devices like adjoining infinities to the real line are easily waffled over, but perhaps not so easily formalized.

Formalization itself can be clarifying. One is forced to be less sloppy about things like variable binding (for example, what does $f'(x) = g(x)$ mean?). To avoid proving almost identical theorems twice, we saw the need for some common framework for limiting processes before we were actually aware that such frameworks (nets and filters) already existed. This constitutes an example of how abstraction can be driven by mundane considerations of economy, rather than of beauty.

We still have a long way to go in providing specialized proof support for reasoning about the real numbers. Certain recurring themes, such as proof by bisection, have been codified in actual HOL theorems, but more experience is required to find out what else is desirable.

Applications have only just got off the ground, but it seems to be fertile territory for

computer-aided verification by theorem proving.

11 Acknowledgements

I would like to thank Mike Gordon for his help and encouragement, including his patient debugging when I was learning HOL, and for drawing my attention to Kelley's discussion of convergence nets [17]. Thomas Forster and Larry Paulson gave me some useful advice about constructing the reals. Everyone in the group at Cambridge has contributed to making such a pleasant and stimulating work environment, and many of them, too numerous to mention here, have helped my work in one way or another. I am specially grateful to Richard Boulton, Jim Grundy, Monica Nesi and Eike Ritter, who read an early version of this paper and made many helpful comments. Thanks are also due to the Science and Engineering Research Council for financial support.

References

[1] F. A. Behrend, *A Contribution to the Theory of Magnitudes and the Foundations of Analysis*, Mathematische Zeitschrift, vol. 63, pp. 345–362, 1956.

[2] E. Bishop, D. Bridges, *Constructive analysis*, Springer-Verlag, 1985.

[3] N. Bourbaki, *Elements of mathematics, vol. 3: General Topology, part 1*, Hermann 1966.

[4] N. G. de Bruijn, *Defining Reals Without the Use of Rationals*, Indagationes Mathematicae, vol. 38, pp. 100–108, 1976.

[5] J. C. Burkill, *A First Course in Mathematical Analysis*, Cambridge 1962.

[6] J. C. Burkill, H. Burkill, *A Second Course in Mathematical Analysis*, Cambridge 1970.

[7] C. W. Burrill, *Foundations of Real Numbers*, McGraw-Hill 1967.

[8] J. Chirimar, D. J. Howe, *Implementing Constructive Real Analysis*, preprint 1992.

[9] L. W. Cohen, G. Ehrlich, *Structure of the real number system*, Van Nostrand 1963.

[10] J. H. Conway, *On Numbers and Games*, Academic Press 1976.

[11] F. Faltin, N. Metropolis, B. Ross, G.-C. Rota, *The Real Numbers as a Wreath Product*, Advances in Mathematics, vol. 16, pp. 278–304, 1975.

[12] J. S. Golan, *The Theory of Semirings with Applications in Mathematics and Computer Science*, Longman 1992.

[13] J. R. Harrison, L. Théry, *Reasoning About the Reals: The Marriage of HOL and Maple*, Proceedings of the 4th International Conference on Logic Programming and Automated Reasoning (LPAR 93), Springer Lecture Notes in Artificial Intelligence vol. 698, pp. 351–353, 1993.

[14] IEEE, *Standard for Binary Floating Point Arithmetic*, ANSI/IEEE Standard 754-1985.

[15] C. Jones, *Completing the Rationals and Metric Spaces in LEGO*, Proceedings of the 2nd Workshop on Logical Frameworks, Edinburgh 1991.

[16] L. S. Jutting, *Checking Landau's "Grundlagen" in the Automath System*, PhD thesis, Eindhoven University of Technology 1977.

[17] J. L. Kelley, *General Topology*, Van Nostrand 1955.

[18] S. Kuhn, *The Derivative à la Carathéodory*, American Mathematical Monthly, vol. 98, pp. 40–44, 1991.

[19] E. Landau, *Foundations of analysis*, Chelsea 1951.

[20] A. Pnueli, *From Timed To Hybrid Systems*, preprint 1992.

[21] R. R. Stoll, *Set theory and logic*, Dover 1979.

[22] H. A. Thurston, *The number system*, Blackie 1956.

[23] P. J . L. Wallis, *Improving floating-point programming*, Wiley 1990.