RZ: a Tool for Bringing Constructive and Computable Mathematics Closer to Programming Practice

ANDREJ BAUER, *Faculty of Mathematics and Physics, University of Ljubljana, Ljubljana, Slovenia*
E-mail: Andrej.Bauer@fmf.uni-lj.si

CHRISTOPHER A. STONE, *Computer Science Department, Harvey Mudd College, Claremont, CA, USA*
E-mail: stone@cs.hmc.edu

Abstract

Realizability theory is not just a fundamental tool in logic and computability. It also has direct application to the design and implementation of programs, since it can produce code interfaces for the data structure corresponding to a mathematical theory. Our tool, called RZ, serves as a bridge between the worlds of constructive mathematics and programming. By using the realizability interpretation of constructive mathematics, RZ translates specifications in constructive logic into annotated interface code in Objective Caml. The system supports a rich input language allowing descriptions of complex mathematical structures. RZ does not extract code from proofs, but allows any implementation method, from handwritten code to code extracted from proofs by other tools.

1 Introduction

Given a mathematical description of a mathematical structure (constants, functions, relations and axioms), what should a computer implementation look like?

For simple cases, the answer is obvious. A group would have a type whose values represent group elements, as well as an associative binary operation a constant neutral element and a unary inverse operator.

But for more interesting structures, especially those arising in constructive mathematical analysis, the answer is less clear. How do we implement the real numbers (a Cauchy-complete Archimedean ordered field)? Or choose the operations for a compact metric space or a space of smooth functions? Significant research goes into finding satisfactory representations [12, 15, 36, 37], while implementations of exact real arithmetic [24, 29] show that theory can be put into practice quite successfully.

The theory of realizability provides guidance in development of computable mathematics. Our work shows that realizability is not only a fundamental tool in logic and computability, but also has direct application to the design and implementation of programs: it can produce a description of the data structure (a code interface) directly corresponding to a mathematical specification.

However, doing this manually is tedious. Worse, different but logically equivalent sets of axioms correspond to different, although interdefinable, interfaces for code. One might then want to compare several variations, since some interfaces will be more useful than others in practice. And few programmers—even those with strong backgrounds in mathematics and classical logic—are familiar

with constructive logic or realizability. Programmers are used to language constructs describing interfaces (e.g. C++ header files, ML signatures or Java interfaces) and to using logical assertions (e.g. preconditions and postconditions).

We have therefore implemented a system, called RZ, to serve as a bridge between the logical world and the programming world.[1] RZ translates specifications in constructive logic into standard interface code in a programming language (currently Objective Caml [25], but other languages could be used).

The constructive part of the original specification turns into interface code, listing types and values to be implemented. The rest becomes assertions about these types and values. As these assertions have no computational content, their constructive and classical meanings agree, and they can be understood by programmers and mathematicians accustomed to classical logic.

RZ was designed as a lightweight system supporting a rich input language. Although transforming complete proofs into complete code is possible [23], we have not implemented this. Other excellent systems, including Coq [8] and Minlog [7], can extract programs from proofs. But they work best managing the entire task, from specification to code generation. In contrast, interfaces generated by RZ can be implemented in any fashion as long as the assertions are satisfied. Code can be written by hand, using imperative, concurrent and other language features rather than a 'purely functional' subset. At the other extreme, the output of RZ can be viewed as a possible *input* to a program extraction tool, where the distinction between computational and non-computational parts (in Coq these are Set and Prop, respectively) has been automatically determined; a corresponding implementation would then be provided via theorem-proving and program extraction.

The article is organized as follows. In Section 2 we present a version of realizability which is most suitable for our purposes. Sections 3 and 4 describe the input and the output language of RZ, while in Section 5 we explain how RZ translates from one to the other. Various implementation issues are discussed in Section 6, and examples of RZ at work are shown in Section 7. We conclude with remarks on related work in Section 8.

An earlier description of our RZ work appears in [4]; since then, the input syntax and underlying implementation has been significantly revised and improved, and the support for dependent types and hoisting is completely new.

2   Typed realizability

RZ is based on *typed realizability* by Longley [26]. This variant of realizability corresponds most directly to programmers' intuition about implementations.

We approach typed realizability and its relationship to real-world programming by way of example. Suppose we are asked to design a data structure for the set $\mathcal{G}$ of all finite simple[2] directed graphs with vertices labelled by distinct integers. An exemplar directed graph G is shown in Figure 1. One common representation of such graphs uses a pair of lists (ℓ_V, ℓ_A), where ℓ_V is the list of vertex labels and ℓ_A is the *adjacency list* representing the arrows by pairing the labels of each source and target. For the above graph G, $\ell_V = [1; 2; 3; 4]$ and $\ell_A = [(1,2); (2,2); (2,3); (3,2); (3,1)]$. Thus we define the datatype of graphs as[3]

$$\texttt{type graph} = \texttt{int list} \; * \; \texttt{(int * int) list}$$

---

[1]RZ is publicly available for download at http://math.andrej.com/rz/, together with the abridged version of this article presented at CiE 2007 [5].

[2]Simple means at most one arrow between any two vertices.

[3]We use OCaml notation in which t list classifies finite lists of elements of type t, and $t_1 * t_2$ classifies pairs containing a value of type t_1 and a value of type t_2.

FIGURE 1. A finite directed graph G

However, this is not a complete description of the intended representation, as there are representation invariants and conditions not expressed by the type, e.g.

1. The order in which vertices and arrows are listed is not important; e.g. $[1; 2; 3; 4]$ and $[4; 1; 2; 3]$ represent the same vertices.
2. Each vertex and arrow must be listed exactly once.
3. The source and target of each arrow must appear in the list of vertices.

Thus, to implement the mathematical set $\mathcal{G}$, we must not only decide on the underlying datatype `graph`, but also determine what values of that type represent which elements of $\mathcal{G}$. As we shall see next, this can be expressed either using a *realizability relation* or a *partial equivalence relation (per)*.

2.1 Modest sets and pers

We now define typed realizability as it applies to OCaml. Other general-purpose programming languages could be used instead, as long as they provide the usual ground types, product and function types.[4]

Let Type be the collection of all (non-parametric) OCaml types. To each type $t \in$ Type we assign the set $[\![t]\!]$ of values of type t that behave *functionally* in the sense of Longley [27]. Such values are represented by terminating expressions that do not throw exceptions or return different results on different invocations. They may *use* exceptions, store and other computational effects, provided they appear functional from the outside; a useful example using computational effects is presented in Section 7.6. A functional value of function type may diverge as soon as it is applied. The collection Type with the assignment of functional values $[\![t]\!]$ to each $t \in$ Type forms a *typed partial combinatory algebra (TPCA)*, which provides a theoretical basis for the definition of a realizability model that suits our needs.

Going back to our example, we see that an implementation of directed graphs $\mathcal{G}$ specifies a datatype $|\mathcal{G}| = \text{graph}$ together with a *realizability relation* $\Vdash_{\mathcal{G}}$ between $\mathcal{G}$ and $[\![\text{graph}]\!]$. The meaning of $(\ell_V, \ell_A) \Vdash_{\mathcal{G}} G$ is 'OCaml value (ℓ_V, ℓ_A) represents/realizes/implements graph G'. There are two natural conditions that $\Vdash_{\mathcal{G}}$ ought to satisfy: (i) for every $G \in \mathcal{G}$ there should be at least one realizer (ℓ_V, ℓ_A) representing it, and (ii) if (ℓ_V, ℓ_A) represents both G and G' then $G = G'$.[5] If (ℓ_V, ℓ_A) and (ℓ_V', ℓ_A') represent the same graph (e.g. because ℓ_V is a permutation of ℓ_V', and similarly for ℓ_A and ℓ_A') we say that they are *equivalent* and write $(\ell_V, \ell_A) \approx_{\mathcal{G}} (\ell_V', \ell_A')$. The relation $\approx_{\mathcal{G}}$ is a *partial* equivalence relation (symmetric and transitive, but not reflexive) because not every $(\ell_V, \ell_A) \in [\![\text{graph}]\!]$ represents a graph.

---

[4]It is also convenient to work with a language that supports sum types, as this allows a more natural representation of disjoint unions.

[5]The latter condition is called *modesty* and is not strictly necessary for the development of the theory, though programmers would naturally expect it to hold.

A general definition is in order. A *modest set* is a triple $A = (\langle A \rangle, |A|, \Vdash_A)$ where $\langle A \rangle$ is the *underlying set*, $|A| \in \mathsf{Type}$ is the *underlying type*, and $\Vdash_A$ is a *realizability relation* between $[\![|A|]\!]$ and $\langle A \rangle$, satisfying

1. *totality:* for every $x \in \langle A \rangle$ there is $v \in [\![|A|]\!]$ such that $v \Vdash_A x$, and
2. *modesty:* if $u \Vdash_A x$ and $u \Vdash_A y$ then $x = y$.

The *support* of A is the set $\|A\| = \{v \in [\![|A|]\!] \mid \exists x \in \langle A \rangle . v \Vdash_A x\}$ of those values that realize something. We define the relation $\approx_A$ on $[\![|A|]\!]$ by

$$u \approx_A v \iff \exists x \in \langle A \rangle . \left(u \Vdash_A x \wedge v \Vdash_A x \right).$$

From totality and modesty of $\Vdash_A$ it follows that $\approx_A$ is a per, i.e. symmetric and transitive. Observe that $\|A\| = \{v \in [\![|A|]\!] \mid v \approx_A v\}$, whence $\approx_A$ restricted to $\|A\|$ is an equivalence relation. In fact, we may recover a modest set up to isomorphism from $|A|$ and $\approx_A$ by taking $\langle A \rangle$ to be the set of equivalence classes of $\approx_A$, and $v \Vdash_A x$ to mean $v \in x$.

The two views of implementations, as modest sets $(\langle A \rangle, |A|, \Vdash_A)$, and as pers $(|A|, \approx_A)$, are equivalent.[6] We concentrate on the view of modest sets as pers. They are more convenient to use in RZ because they refer only to types and values, as opposed to arbitrary sets. Nevertheless, it is useful to understand how modest sets and pers arise from natural programming practice.

Modest sets form a category whose objects are modest sets and morphisms are the realized functions. A *realized function* $f : A \to B$ is a function $f : \langle A \rangle \to \langle B \rangle$ for which there exists $v \in [\![|A| \to |B|]\!]$ such that, for all $x \in \langle A \rangle$ and $u \in |A|$,

$$u \Vdash_A x \implies v\, u \Vdash_B f(x). \tag{1}$$

This condition is just a mathematical expression of the usual idea that v is an implementation of f if it does to realizers what f does to the elements they represent.

The equivalent category of pers has as objects pairs $A = (|A|, \approx_A)$ where $|A| \in \mathsf{Type}$ and $\approx_A$ is a per on $[\![|A|]\!]$. A morphism $A \to B$ is represented by a function $v \in [\![|A| \to |B|]\!]$ such that, for all $u, u' \in \|A\|$,

$$u \approx_A u' \implies v\, u \approx_B v\, u'. \tag{2}$$

Values v and v' that both satisfy (2) represent the same morphism if, for all $u, u' \in \|A\|$, $u \approx_A u'$ implies $v\, u \approx_B v'\, u'$.

The category of pers has a very rich structure. For example, we may form a Cartesian product $A \times B$ of pers A and B by

$$|A \times B| = |A| * |B|,$$

$$(u_1, v_1) \approx_{A \times B} (u_2, v_2) \iff u_1 \approx_A u_2 \wedge v_1 \approx_B v_2.$$

The projections $\pi_1 : A \times B \to A$ and $\pi_2 : A \times B \to B$ are realized by `fst` and `snd`, respectively.

---

[6]And there is a third view, as a partial surjection $\delta_A : {\subseteq} [\![|A|]\!] \twoheadrightarrow \langle A \rangle$, with $\delta_A(v) = x$ when $v \Vdash_A x$. This is how realizability is presented in Type Two Effectivity [37].

The morphisms between pers A and B again form a per B^A, also written as $A \to B$, called the *exponential* of A and B, with

$$|B^A| = |A| \to |B|,$$

$$\|B^A\| = \left\{ v \in [\![|A| \to |B|]\!] \mid \forall u, u' \in [\![|A|]\!] . \left(u \approx_A u' \implies vu \approx_B vu' \right) \right\}$$

$$u \approx_{B^A} v \iff u, v \in \|A\| \wedge \forall w \in \|A\| . uw \approx_B vw.$$

The evaluation map $e : B^A \times A \to B$ is realized by OCaml application, $\mathtt{fun}\,(u, v) \to u\,v$. If a function $f : C \times A \to B$ is realized by v, then its transpose $\tilde{f} : C \to B^A$, $\tilde{f}(z)(x) = f(z, x)$, is realized by $\mathtt{fun}\,z\,x \to v(z, x)$. This shows that the category of pers is Cartesian closed. In Section 5.1 we review other canonical constructions on modest sets.

As an example we consider the cyclic group on seven elements $(\mathbb{Z}_7, 0, -, +)$. To implement the group, we must give a representation of $\mathbb{Z}_7$ as a per $Z = (|Z|, \approx_Z)$, and provide realizers for the neutral element 0, negation $-$ and addition $+$.

One possibility is to choose $\mathtt{int}$ as the underlying type $|Z|$, and to let $\|Z\|$ be only the integers 0 through 6. Then negation and addition must work modulo 7 (i.e. must return an integer in the range 0–6 when given integers in this range). The neutral element would be the integer constant 0, and the equivalence $\approx_Z$ would be integer equality.

Alternatively, we could take $\mathtt{int}$ as the underlying type $|Z|$, but let $\|Z\|$ include all values of type $\mathtt{int}$. In this case, negation and addition could be simply integer addition and negation.[7] Here the neutral element could be implemented as any integer multiple of 7, and the equivalence $\approx_Z$ would be equivalence-modulo-7.

Both of these pers happen to be *decidable*, i.e. it can be algorithmically decided whether two values in $\|Z\|$ represent the same element of $\mathbb{Z}_7$.

Not all pers are decidable. Examples include implementations of semigroups with an undecidable word problem [31], implementations of computable sets of integers (which might be realized by membership functions of type $\mathtt{int} \to \mathtt{bool}$), and implementations of computable real numbers (which might be realized by infinite Cauchy sequences). There is no presupposition that pers are computable (implementable). We can require decidable equivalence by adding a suitable axiom; see Section 7.1.

2.2 Interpretation of logic

In the realizability interpretation of logic, each formula ϕ is assigned a set of *realizers*, which can be thought of as computations that witness the validity of ϕ. The situation is somewhat similar, but not equivalent, to the propositions-as-types translation of logic into type theory, where proofs of a proposition correspond to terms of the corresponding type. More precisely, to each formula ϕ we assign an underlying type $|\phi|$ of realizers, but unlike the propositions-as-types translation, not all terms of type $|\phi|$ are necessarily valid realizers for ϕ, and some terms that are realizers may not correspond to any proofs (e.g. if they denote partial functions or use computational effects).

It is customary to write $t \Vdash \phi$ when $t \in [\![|\phi|]\!]$ is a realizer for ϕ. The underlying types and the realizability relation $\Vdash$ are defined inductively on the structure of ϕ; an outline is shown in Figure 2. We say that a formula ϕ is *valid* if it has at least one realizer.

---

[7]Taking care to prevent integer overflow.

Underlying types of realizers:

$$
\begin{array}{llll}
|\top| & = \texttt{unit} & |\bot| & = \texttt{unit} \\
|x=y| & = \texttt{unit} & |\phi \wedge \psi| & = |\phi| \times |\psi| \\
|\phi \Rightarrow \psi| & = |\phi| \rightarrow |\psi| & |\phi \vee \psi| & = \texttt{`or}_0 \texttt{ of } |\phi_0| + \texttt{`or}_1 \texttt{ of } |\phi_1| \\
|\forall x{:}A.\ \phi| & = |A| \rightarrow |\phi| & |\exists x{:}A.\ \phi| & = |A| \times |\phi|
\end{array}
$$

Realizers:

$$
\begin{array}{lll}
() \Vdash \top \\
() \Vdash x=y & \text{iff} & x=y \\
(t_1,t_2) \Vdash \phi \wedge \psi & \text{iff} & t_1 \Vdash \phi \text{ and } t_2 \Vdash \psi \\
t \Vdash \phi \Rightarrow \psi & \text{iff} & \text{for all } u \in |\phi|, \text{ if } u \Vdash \phi \text{ then } t\,u \Vdash \psi \\
\texttt{`or}_0\, t \Vdash \phi \vee \psi & \text{iff} & t \Vdash \phi \\
\texttt{`or}_1\, t \Vdash \phi \vee \psi & \text{iff} & t \Vdash \psi \\
t \Vdash \forall x{:}A.\ \phi(x) & \text{iff} & \text{for all } u \in |A|, \text{ if } u \Vdash_A x \text{ then } t\,u \Vdash \phi(x) \\
(t_1,t_2) \Vdash \exists x{:}A.\ \phi(x) & \text{iff} & t_1 \Vdash_A x \text{ and } t_2 \Vdash \phi(x)
\end{array}
$$

FIGURE 2. Realizability interpretation of logic (outline)

In classical mathematics, a predicate on a set X may be viewed as a subset of X or a (possibly non-computable) function $X \rightarrow \texttt{bool}$, where $\texttt{bool} = \{\bot, \top\}$ is the set of truth values. Accordingly, since in realizability propositions are witnessed by realizers, a predicate ϕ on a modest set A may be viewed as a subset of $\langle A \rangle \times [[|\phi|]]$, or a (possibly non-computable) function $\langle A \rangle \times [[|\phi|]] \rightarrow \{\bot, \top\}$. In terms of pers, which is what RZ uses, a predicate ϕ on a per $A = (|A|, \approx_A)$ is a (possibly non-computable) function $\phi : [[|A|]] \times [[|\phi|]] \rightarrow \texttt{bool}$, i.e.

- *strict:* if $\phi(u,v)$ then $u \in \|A\|$ and
- *extensional:* if $\phi(u_1,v)$ and $u_1 \approx_A u_2$ then $\phi(u_2,v)$.

We illustrate how the realizability interpretation extracts the computational content of a proposition. To make an interesting example, suppose we have implemented the real numbers $\mathbb{R}$ as a per $R = (\texttt{real}, \approx_R)$, and consider the statement that every cubic $x^3 + ax + b$ has a root,

$$\forall a{:}R.\ \forall b{:}R.\ \exists x{:}R.\ x^3 + ax + b = 0. \tag{3}$$

By computing according to Figure 2, we see that a realizer for this proposition is a value r of type $\texttt{real} \rightarrow \texttt{real} \rightarrow \texttt{real} \times \texttt{unit}$ such that, if t realizes $a \in \mathbb{R}$ and u realizes $b \in \mathbb{R}$, then $r\,t\,u = (v,w)$ with v realizing a real number x such that $x^3 + ax + b = 0$, and with w being trivial. (This can be 'thinned' to a realizer of type $\texttt{real} \rightarrow \texttt{real} \rightarrow \texttt{real}$ that does not bother to compute w.) In essence, the realizer r computes a root of the cubic equation. Note that r is *not* extensional, i.e. different realizers t and u for the same a and b may result in different roots. To put this in another way, r realizes a *multi-valued* function[8] rather than a per morphism. It is well known in computable mathematics that certain operations, such as equation solving, are only computable if we allow them to be multi-valued. They arise naturally in RZ as translations of $\forall \exists$ statements.

There are propositions whose realizers are 'irrelevant' or free of computational content. For example, realizers for $\top$ and equality have type $\texttt{unit}$. Another example is a negation $\neg\phi$, which is defined to be the same as $\phi \Rightarrow \bot$, whose realizers have type $|\phi| \rightarrow \texttt{unit}$. Such realizers do not

---

[8]The multi-valued nature of the realizer comes from the fact that it computes *any one* of many values, not that it computes *all* of the many values.

compute anything useful, and we may as well throw them away. Sometimes only a part of a realizer is computationally irrelevant, as we saw in the last example. Propositions that are free of computational content are characterized as the $\neg\neg$-*stable propositions*. A proposition ϕ is said to be $\neg\neg$-stable, or just *stable* for short, when $\neg\neg\phi \Rightarrow \phi$ is valid. Any *negative* proposition, i.e. one built from $\top$, $\bot$, $=$, $\wedge$, $\Rightarrow$ and $\forall$ is stable, but there may be other propositions that are stable and are not written in the negative form.

It would be unproductive to bother the programmer with requirements for useless code. On input, one can specify whether abstract predicates have computational content. On output, extracted realizers go through a *thinning* phase, which removes irrelevant realizers.

2.3   Uniform families of modest sets

Many structures are naturally viewed as families of sets, or sets depending on parameters, or *dependent types* as they are called in type theory. For example, the n-dimensional Euclidean space $\mathbb{R}^n$ depends on the dimension $n \in \mathbb{N}$, the Banach space $\mathcal{C}([a,b])$ of uniformly continuous real functions on the closed interval $[a,b]$ depends on $a,b \in \mathbb{R}$ such that $a < b$, etc. In general, a family of sets $\{A_i\}_{i \in I}$ is an assignment of a set A_i to each $i \in I$ from an *index set I*.

In the category of modest sets the appropriate notion is that of a *uniform* family $\{A_i\}_{i \in I}$, which is an assignment of a modest set $A_i = (\langle A_i \rangle, |A|, \Vdash_{A_i})$ to each $i \in \langle I \rangle$, where I is an index modest set [19, 6.3]. The uniformity comes from the requirement that all the A_i's share the same underlying type $|A_i| = |A|$. It is a desirable restriction from the implementation point of view, because it removes dependencies at the level of types. Note also that there is no dependency on the realizers, only on the elements of the underlying set.

We may express uniform families in terms of pers, too. A uniform family of pers $\{A_i\}_{i \in I}$ indexed by a per I is given by an underlying type $|A|$ and a family of pers $(\approx_{A_i})_{i \in [\![|I|]\!]}$, i.e.

- *strict:* if $u \approx_{A_i} v$ then $i \in \|I\|$, and
- *extensional:* if $u \approx_{A_i} v$ and $i \approx_I j$ then $u \approx_{A_j} v$.

We may form the *sum* $\Sigma_{i \in I} A_i$ of a uniform family $\{A_i\}_{i \in I}$ as

$$|\Sigma_{i \in I} A_i| = |I| \times |A|$$

$$(i_1, u_1) \approx_{\Sigma_{i \in I} A_i} (i_2, u_2) \iff i_1 \approx_I i_2 \wedge u_1 \approx_{A_{i_1}} u_2$$

and the *product* $\Pi_{i \in I} A_i$ as

$$|\Pi_{i \in I} A_i| = |I| \to |A|$$

$$\|\Pi_{i \in I} A_i\| = \{v \in [\![|I| \to |A|]\!] \mid \forall i,j \in [\![|I|]\!] . (i \approx_I j \implies v i \approx_{A_i} v j)\}$$

$$u \approx_{\Pi_{i \in I} A_i} v \iff u, v \in \|\Pi_{i \in I} A_i\| \wedge \forall i,j \in [\![|I|]\!] . (i \approx_I j \implies u i \approx_{A_i} v j).$$

These constructions allow us to interpret (extensional) dependent type theory in the category of modest sets.

As an example of a uniform family we consider the cyclic group $(\mathbb{Z}_n, 0, -, +)$ of order n. To keep things simple, we assume that n ranges over natural numbers that can be represented by type `int` (i.e. $|N| = \text{int}$), and that $\approx_N$ is equality. The uniform family $\{Z_n\}_{n \in N}$ is then like the cyclic group of order 7, with 7 replaced by n. Ignoring overflow, the underlying type would be $|Z_n| = \text{int}$. Any of the implementations suggested for $\mathbb{Z}_7$ would work here, with 7 replaced by the parameter n; in

one case we would have $u \approx_{Z_n} v \iff u = v$ and in the other $u \approx_{Z_n} v \iff u \bmod n = v \bmod n$. Negation would be specified as a constant of dependent type $\Pi_{n \in N} Z_n \to Z_n$. Its realizer `neg` would then have type $|N| \to |Z_n| \to |Z_n|$, i.e. $\mathtt{int} \to \mathtt{int} \to \mathtt{int}$, so that $\mathtt{neg}(n)$ would be a realizer for negation on $\mathbb{Z}_n$. The realizer for addition would similarly take an extra argument n.

3 Specifications as signatures with assertions

In programming we distinguish between *implementation* and *specification* of a structure. In OCaml these two notions are expressed with modules and module types, respectively.[9] A module defines types and values, while a module type simply lists the types, type definitions and values provided by a module. For a complete specification, a module type must also be annotated with *assertions* which specify the required properties of declared types and values. For example, if we look at the definition of module type Ab in Figure 3, we might guess that Ab is a signature for an Abelian group. However, Ab requires an implementation that satisfies the signature of an Abelian group, but does not guarantee an Abelian group. A complete description would contain the following further *assertions*:

1. there is a per $\approx_t$ on $[\![\mathtt{t}]\!]$,
2. $\mathtt{zero} \in \|\mathtt{t}\|$.
3. for all $u, v \in [\![\mathtt{t}]\!]$, if $u \approx_t v$ then $\mathtt{neg}\, u \approx_t \mathtt{neg}\, v$,
4. for all $u_1, u_2, v_1, v_2 \in [\![\mathtt{t}]\!]$, if $u_1 \approx_t v_1$ and $u_2 \approx_t v_2$ then $\mathtt{add}(u_1, u_2) \approx_t \mathtt{add}(v_1, v_2)$,
5. for all $u \in \|\mathtt{t}\|$, $\mathtt{add}(\mathtt{zero}, u) \approx_t u$,
6. for all $u \in \|\mathtt{t}\|$, $\mathtt{add}(u, \mathtt{neg}\, u) \approx_t \mathtt{zero}$,
7. for all $u, v, w \in \|\mathtt{t}\|$, $\mathtt{add}(\mathtt{add}(u, v), w) \approx_t \mathtt{add}(u, \mathtt{add}(v, w))$,
8. for all $u, v \in \|\mathtt{t}\|$, $\mathtt{add}(u, v) \approx_t \mathtt{add}(v, u)$.

Assertions 2–4 state that `zero`, `neg` and `add` realize a constant, a unary and a binary operation, respectively, while assertions 5–8 correspond to axioms for Abelian groups.

The output of RZ consists of *module specifications*, module types plus assertions about their components. More specifically, a specification may contain value declarations, type declarations and definitions, module declarations, specification definitions, proposition declarations and assertions. The language of specifications is summarized in Figure 4.

A specification can describe an OCaml structure (a collection of definitions for types and values) or an OCaml functor (a parameterized module, i.e. a function mapping modules to modules). The latter would be appropriate, e.g. when describing a uniform implementation of the real numbers that works given any implementation of natural numbers.

```
module type Ab =
sig
  type t
  val zero : t
  val neg : t -> t
  val add : t * t -> t
end
```

FIGURE 3. The module type Ab

Types
$$\tau ::= \quad T \mid M.T \qquad\qquad\qquad\qquad\qquad \text{Type names}$$
$$\mid \quad \texttt{unit} \mid \tau_1 \times \tau_2 \qquad\qquad\qquad \text{Unit and cartesian product}$$
$$\mid \quad \tau_1 \to \tau_2 \qquad\qquad\qquad\qquad \text{Function type}$$
$$\mid \quad {}^{\backprime}l_1 \,\texttt{of}\, \tau_1 + \cdots + {}^{\backprime}l_n \,\texttt{of}\, \tau_n \qquad \text{Disjoint sum}$$
$$\mid \quad \alpha \qquad\qquad\qquad\qquad\qquad \text{Polymorphic types}$$

Terms
$$e ::= \quad x \mid M.x \qquad\qquad\qquad\qquad\qquad \text{Term names}$$
$$\mid \quad \texttt{fun}\, x{:}\tau_1 \to e \mid e_1\, e_2 \qquad\qquad \text{Functions and application}$$
$$\mid \quad () \mid (e_1,\ldots,e_n) \mid \texttt{p}_n\, e \qquad \text{Tuples and projection}$$
$$\mid \quad {}^{\backprime}l\, e \mid (\texttt{match}\, e \,\texttt{with}\, {}^{\backprime}l_1\, x_1 \to e_1 \mid \cdots \mid {}^{\backprime}l_n\, x_n \to e_n) \qquad \text{Injection and projection from a sum}$$
$$\mid \quad \texttt{assure}\, x{:}\tau,\, p \,\texttt{in}\, e \mid \texttt{assure}\, p \,\texttt{in}\, e \qquad \text{Obligations}$$
$$\mid \quad \texttt{let}\, x{=}e_1 \,\texttt{in}\, e_2 \qquad\qquad\qquad \text{Local definitions}$$

Propositions (negative fragment)
$$p ::= \quad P \mid M.P \qquad\qquad\qquad\qquad\qquad \text{Atomic proposition}$$
$$\mid \quad \top \mid \bot \mid \neg p_1 \mid p_1 \wedge p_2 \mid p_1 \to p_2 \mid p_1 \leftrightarrow p_2 \qquad \text{Predicate logic}$$
$$\mid \quad \texttt{match}\, e \,\texttt{with}\, {}^{\backprime}l_1\, x_1 \Rightarrow p_1 \mid \cdots \mid {}^{\backprime}l_n\, p_n \Rightarrow e_n \qquad \text{Propositional case}$$
$$\mid \quad \texttt{fun}\, x{:}\tau \to p \mid p\, e \qquad\qquad \text{Propositional functions and application}$$
$$\mid \quad e_1 \approx_s e_2 \mid e{:}\|s\| \qquad\qquad \text{Pers and support}$$
$$\mid \quad e_1{=}e_2 \qquad\qquad\qquad\qquad \text{(Observational) term equality}$$
$$\mid \quad \forall x{:}\tau.\, p \mid \forall x{:}\|s\|.\, p \qquad\qquad \text{Term quantifiers}$$

Basic modest sets
$$s ::= \quad S \mid s\, e$$

Modules
$$m ::= \quad M \mid m.M \qquad\qquad\qquad\qquad \text{Model names}$$
$$\mid \quad m_1\, m_2 \qquad\qquad\qquad\qquad \text{Application of parameterized model}$$

Proposition Kinds
$$\Pi ::= \quad \texttt{bool} \qquad\qquad\qquad\qquad\qquad \text{Classifier for propositions}$$
$$\mid \quad \tau \to \Pi \qquad\qquad\qquad\qquad \text{Classifier for a predicate/relation}$$

Specification elements
$$\theta ::= \quad \texttt{val}\, x : \tau \qquad\qquad\qquad\qquad \text{Value declaration}$$
$$\mid \quad \texttt{type}\, T \qquad\qquad\qquad\qquad \text{Type declaration}$$
$$\mid \quad \texttt{type}\, T = \tau \qquad\qquad\qquad \text{Type definition}$$
$$\mid \quad \texttt{module}\, m : \Sigma \qquad\qquad\qquad \text{Module declaration}$$
$$\mid \quad \texttt{module type}\, S = \Sigma \qquad\qquad \text{Specification definition}$$
$$\mid \quad \texttt{predicate}\, P : \Pi \qquad\qquad \text{Proposition declaration}$$
$$\mid \quad \texttt{assertion}\, A : p \qquad\qquad\qquad \text{Assertion}$$

Specifications (module types with assertions)
$$\Sigma ::= \quad S \mid m.S \qquad\qquad\qquad\qquad \text{Specification names}$$
$$\mid \quad \texttt{sig}\, \theta_1 \ldots \theta_n \,\texttt{end} \qquad\qquad \text{Specification elements}$$
$$\mid \quad \texttt{functor}\, (m{:}\Sigma_1) \to \Sigma_2 \qquad \text{Parameterized specification}$$

FIGURE 4. The syntax of specifications (simplified)

Assertions within module specifications (which appear as code comments) are expressed in the *negative fragment* of first-order logic, which contains constants for truth and falsehood, negation, conjunction, implication, equivalence and universal quantification (but no disjunction or existential). For convenience we also introduce the propositional case

$$\texttt{match}\, e \,\texttt{with}\, {}^{\backprime}l_1 x_1 \Rightarrow p_1 \mid \cdots \mid {}^{\backprime}l_n x_n \Rightarrow p_n$$

which is read 'if e is of the form ${}^{\backprime}l_i x_i$ then p_i holds $(i = 1,\ldots,n)$', with the understanding that the expression is false if e does not match any case. This can be expressed with the negative formula

$$(\forall x_1.e = {}^{\backprime}l_1 x_1 \to p_1) \wedge \cdots \wedge (\forall x_n.e = {}^{\backprime}l_n x_n \to p_n) \wedge$$
$$\neg((\forall x_1.e \neq {}^{\backprime}l_1 x_1) \wedge \cdots \wedge (\forall x_n.e \neq {}^{\backprime}l_n x_n)).$$

The negative fragment is the part of first-order logic that has no computational content in the realizability interpretation. Consequently, the classical and constructive interpretations of assertions agree. This is quite desirable, since RZ acts as a bridge between constructive mathematics and real-world programmers, who typically are not familiar with constructive logic.

RZ ever produces only a small subset of OCaml types (the unit type, products, function types, polymorphic variant types and parameter types). Correspondingly, the language of terms produced is fairly simple (tuples, functions, polymorphic variants and local definitions). However, the programmer is free to implement a specification using any types and terms that exist in OCaml.

A special kind of term is an *obligation* `assure` $x{:}\tau, p$ `in` e, which means 'in term e, let x be any element of $[\![\tau]\!]$ that satisfies p'. An obligation is equivalent to a combination of Hilbert's indefinite description operator and a local definition, `let` $x{=}(\varepsilon x{:}\tau.p)$`in` e, where $\varepsilon x{:}\tau.p$ means 'any $x\in[\![\tau]\!]$ such that p'. The alternative form `assure` p `in` e stands for `assure` _`:unit`, p `in` e.

Obligations arise from the fact that well-formedness of the input language is undecidable; see Section 4. In such cases the system computes a realizability translation, but also produces obligations to be checked. The programmer must replace each obligation with a value satisfying the obligation (i.e. demonstrate that the obligation can be satisfied). If such values do not exist, the specification is unimplementable.

4   The input language

The input to RZ consists of one or more theories. A RZ *theory* is a generalized logical signature with associated axioms, similar to a Coq module signature. Theories describe *models*, or implementations. A summary of the input language appears in Figure 5.

The simplest theory Θ is a list of *theory element*s `thy` $\theta_1\ldots\theta_n$ `end`. A theory element may specify that a certain set, set element, proposition or predicate or model must exist (using the `Parameter` keyword). It may also provide a definition of a set, term, proposition, predicate or theory (using the `Definition` keyword). Finally, a theory element can be a named axiom (using the `Axiom` keyword).

We allow model parameters in theories; typical examples in mathematics include the theory of a vector space parameterized by a field of scalars, or the theory of the real numbers parameterized by a model of the natural numbers.

Following Sannella, Sokolowski, and Tarlecki[10] [32] RZ supports two forms of parameterization. A theory of a parameterized implementation $[m{:}\Theta_1]{\rightarrow}\Theta_2$ describes a uniform family of models (i.e. a single implementation; a functor in OCaml) that maps every model m satisfying Θ_1 to a model of Θ_2. In contrast, a theory $\lambda m{:}\Theta_1.\ \Theta_2$ maps models to theories; if T is such a theory, then $T(M_1)$ and $T(M_2)$ are theories whose implementations might be completely unrelated.[11]

Propositions and predicates appearing in theories may use full first-order constructive logic, not just the negative fragment. The grammar for logical inputs is shown in Figure 5. Most of this should be familiar, including the use of lambda abstraction to define predicates.

The language of sets is rich, going well beyond the type systems of typical programming languages. In addition to any base sets postulated in a theory, one can construct dependent Cartesian products and dependent function spaces. We also supports disjoint unions (with labelled tags), quotient spaces (a set modulo a stable equivalence relation), subsets (elements of a set satisfying a predicate). RZ even permits explicit references to sets of realizers.

---

[10] 'parameterized (program specification) $\neq$ (parameterized program) specification'.

[11] Further, in some cases $T(M_1)$ might be implementable while $T(M_2)$ is not.

Propositions

$\varphi, \rho ::= \quad p \mid M.p$ — Predicate names

$\mid \quad \top \mid \bot \mid \neg\varphi_1 \mid \varphi_1{\wedge}\varphi_2 \mid \varphi_1{\vee}\varphi_2 \mid \varphi_1{\Rightarrow}\varphi_2$ — Predicate logic

$\mid \quad \mathtt{match}\, e\, \mathtt{with}\, l_1\, x_1{\Rightarrow}\varphi_1 \mid \cdots \mid l_n\, x_n{\Rightarrow}\varphi_n$ — Propositional case

$\mid \quad \lambda x{:}s.\ \varphi \mid \varphi\ e$ — Predicates and application

$\mid \quad e_1 = e_2$ — Term equality

$\mid \quad \forall x{:}s.\ \varphi \mid \exists x{:}s.\ \varphi \mid \exists! x{:}s.\ \varphi$ — Term quantifiers

Sets

$s ::= \quad \alpha \mid M.\alpha$ — Set names

$\mid \quad 1 \mid [x{:}s_1] \times s_2$ — Unit and (dependent) Cartesian product

$\mid \quad 0 \mid l_1{:}s_1 + l_2{:}s_2$ — Void and disjoint union

$\mid \quad [x{:}s_1]{\to}s_2$ — (Dependent) function space

$\mid \quad \lambda x{:}s_1.\ s_2 \mid s\ e$ — Dependent set and application

$\mid \quad s/\rho$ — Set quotient by an equivalence relation

$\mid \quad \{x{:}s \mid \rho\}$ — Subset satisfying a predicate

$\mid \quad \mathtt{rz}\ s$ — Realizers of a set

Terms

$e ::= \quad x \mid M.x$ — Term names

$\mid \quad \lambda x{:}s_1.\ e \mid e_1\ e_2$ — Function and application

$\mid \quad (e_1, \ldots, e_2) \mid \pi_n\ e$ — Tuple and projection

$\mid \quad l\ e \mid (\mathtt{match}\, e_0\, \mathtt{with}\, l_1\, x_1{\Rightarrow}e_1 \mid l_2\, x_2{\Rightarrow}e_2)$ — Injection and projection from a union

$\mid \quad [e]_\rho \mid \mathtt{let}\ [x]_\rho = e_1\ \mathtt{in}\ e_2$ — Equivalence class and picking a representative

$\mid \quad \mathtt{rz}\ e \mid \mathtt{let}\ \mathtt{rz}\ x = e_1\ \mathtt{in}\ e_2$ — Realized value and picking a realizer

$\mid \quad e : s$ — Type coercion (e.g., in and out of a subset)

$\mid \quad \iota\, x{:}s.\ \varphi$ — Definite description

$\mid \quad \mathtt{let}\ x = e_1\ \mathtt{in}\ e_2$ — Local definition

Models

$M ::= \quad m \mid M.m$ — Model names

$\mid \quad M_1\ M_2$ — Application of parameterized model

Proposition Kinds

$\Pi ::= \quad \mathtt{Prop} \mid \mathtt{Stable}$ — Classifiers for all propositions/stable propositions

$\mid \quad \mathtt{Equiv}(s)$ — Classifier for stable equivalences on s

$\mid \quad [x{:}s]{\to}\Pi$ — Classifier for a predicate/relation

Set Kinds

$\kappa ::= \quad \mathtt{Set}$ — Classifier for a proper set

$\mid \quad [x{:}s]{\to}\kappa$ — Classifier for a dependent set

Theory Elements

$\theta ::= \quad \mathtt{Definition}\ x := e. \mid \mathtt{Definition}\ \alpha := s.$ — Give a name to a term or set

$\mid \quad \mathtt{Definition}\ p := \varphi. \mid \mathtt{Definition}\ T := \Theta.$ — Give a name to a predicate or theory

$\mid \quad \mathtt{Parameter}\ x : s. \mid \mathtt{Parameter}\ \alpha : \kappa.$ — Require an element in the given set or kind

$\mid \quad \mathtt{Parameter}\ p : \Pi. \mid \mathtt{Parameter}\ m : \Theta.$ — Require a predicate or model of the given sort

$\mid \quad \mathtt{Axiom}\ p : \varphi.$ — Axiom that must hold

Theories

$\Theta ::= \quad T$ — Theory name

$\mid \quad \mathtt{thy}\ \theta_1, \ldots, \theta_n\ \mathtt{end}$ — Theory of a model

$\mid \quad [m{:}\Theta_1]{\to}\Theta_2$ — Theory of a uniform family of models

$\mid \quad \lambda m{:}\Theta_1.\ \Theta_2 \mid \Theta\ m$ — Parameterized theory and application

FIGURE 5. Input syntax (simplified)

The term language includes introduction and elimination constructs for the set level. For product sets we have tuples and projections ($\pi_1 e$, $\pi_2 e$, …), and for function spaces we have lambda abstractions and application. One can inject a term into a tagged union, or do case analyses on the members of a union. We can produce an equivalence class or pick a representative from a equivalence class (as long as what we do with it does not depend on the choice of representative). We can produce a set of realizers or choose a representative from a given set of realizers (as long as what we do with it

does not depend on the choice of representative). We can inject a term into a subset (if it satisfies the appropriate predicate), or project an element out of a subset. Finally, the term language also allows local definitions of term variables, and definite descriptions (as long as there is a unique element satisfying the predicate in question).

From the previous paragraph, it is clear that checking the well-formedness of terms is not decidable. RZ checks what it can, but does not attempt serious theorem proving. Uncheckable constraints remain as obligations in the final output, and should be verified by other means before the output can be used.

5   Translation

Having shown the input and output languages for RZ, we now explain how the translation from one to the other works. A theory is translated to a specification, where the theory elements are translated as follows.

5.1   *Translation of sets and terms*

A set declaration `Parameter` s:`Set` is translated to

```
type s
predicate (≈ₛ) : s → s → bool
assertion symmetric_s :  ∀ x:s, y:s, x ≈ₛ y → y ≈ₛ x
assertion transitive_s : ∀ x:s, y:s, z:s, x ≈ₛ y ∧ y ≈ₛ z → x ≈ₛ z
predicate ‖s‖ : s → bool
assertion support_def_s :  ∀ x:s,  x : ‖s‖ ↔ x ≈ₛ x
```

This says that the programmer should define a type s and a per $\approx_s$ on $[\![s]\!]$. Here $\approx_s$ is *not* an OCaml value of type $s \to s \to$ `bool`, but an abstract relation on the set $[\![s]\!] \times [\![s]\!]$. The relation may be uncomputable.

The translation of the declaration of a dependent set `Parameter t:s→Set` follows the interpretation of dependent sets as uniform families (Section 2.3):

```
type t
predicate ≈ₜ : s → t → t → bool
assertion strict_t :  ∀ x:s, y:t, z:t, y ≈ₜ ₓ z → x : ‖s‖
assertion extensional_t :
  ∀ x:s, y:s, z:t, w:t, x ≈ₛ y → z ≈ₜ ₓ w → z ≈ₜ ᵧ w
assertion symmetric_t : ∀ x:s, y:t, z:t, y ≈ₜ ₓ z → z ≈ₜ ₓ y
assertion transitive_t :
  ∀ x:s, y:t, z:t, w:t, y ≈ₜ ₓ z ∧ z ≈ₜ ₓ w → y ≈ₜ ₓ w
predicate ‖t‖ : s → t → bool
assertion support_def_t : ∀ x:s, y:t, y : ‖t x‖ ↔ y ≈ₜ ₓ y
```

The underlying output type `t` is still non-dependent, but the per is parameterized by `s`.

A value declaration `Parameter` x:s is translated to

```
val x : s
assertion x_support : x : ‖s‖
```

which requires the definition of a value `x` of type `s` which is in the support of `s`. When `s` is not a basic set, RZ computes the interpretation of the underlying type and support.

A value definition `Definition x:=e` where e is an expression denoting an element of `s` is translated to

```
val x : s
assertion x_def : x ≈ₛ e
```

The assertion does *not* force x to be defined as e, only to be equivalent to it with respect to $\approx_s$. This is useful, as often the clearest or easiest ways to define a value are not the most efficient ways to compute it.

Constructions of sets in the input language are translated to corresponding constructions of modest sets. In Section 2.1 we saw how products, exponentials and their dependent versions are formed. We briefly review the remaining constructions of modest sets. We only consider those constructions of terms that are not entirely straightforward.

Disjoint union. A disjoint union of modest sets $l_1{:}A + l_2{:}B$ is the modest set whose underlying type is the sum of underlying types,

$$|l_1{:}A + l_2{:}B| = {}^{\backprime}l_1 \,\texttt{of}\, |A| + {}^{\backprime}l_2 \,\texttt{of}\, |B|,$$

and the per is the disjoint union of pers $\approx_A$ and $\approx_B$, so that we have

$$
\begin{aligned}
{}^{\backprime}l_1 u \approx_{l_1{:}A+l_2{:}B} {}^{\backprime}l_1 v &\Longleftrightarrow u \approx_{s_1} v, \\
{}^{\backprime}l_2 u \approx_{l_1{:}A+l_2{:}B} {}^{\backprime}l_2 v &\Longleftrightarrow u \approx_{s_2} v.
\end{aligned}
$$

Subsets. The construction of subsets may look surprising at first, but it makes sense computationally. Given a predicate ϕ on a per A, the sub-per $\{x{:}A \,|\, \phi\}$ has underlying type $|A| \times |\phi|$ where $(u_1, v_1) \approx_{\{x:A|\phi\}} (u_2, v_2)$ when $u_1 \approx_A u_2$, $v_1 \Vdash \phi(u_1)$ and $v_2 \Vdash \phi(u_2)$. The point is that a realizer for an element of $\{x{:}A \,|\, \phi\}$ carries information about *why* the element belongs to the subset, just like a predicate ϕ carries information of type $|\phi|$ about why a particular instance of ϕ holds.

A type coercion $e{:}t$ can convert an element of the subset $s = \{x{:}t \,|\, \phi(x)\}$ to an element of t. At the level of realizers this is achieved by the first projection, which keeps a realizer for the element but forgets the one for $\phi(e)$. The opposite type coercion $e'{:}s$ takes an $e' \in t$ and converts it to an element of the subset. This is only well formed when $\phi(e')$ is valid. Then, if $u \Vdash_t e'$ and $v \Vdash \phi(e')$, a realizer for $e'{:}s$ is (u, v). However, since RZ cannot in general know a v which validates $\phi(e')$, it emits the pair $(u, (\texttt{assure}\ v{:}|\phi|, \phi uv\ \texttt{in}\ v))$.

Quotients. The category of modest sets has coequalizers, hence a quotient modest set A/ρ may be constructed for an any equivalence relation ρ on A. However, because equality does not carry any computational content, equality of equivalence classes $[x]_\rho = [y]_\rho$ implies only $\neg\neg\rho(x,y)$, not the usual $\rho(x,y)$. As this may cause confusion and mistakes, it is better to permit only quotients by *stable* equivalence relations, which behave as expected.

A stable equivalence relation on a per A is the same thing as a partial equivalence relation ρ on $|A|$ which satisfies $x \approx_A y \Longrightarrow \rho(x,y)$. Then the quotient A/ρ is the per with $|A/\rho| = |A|$ and $x \approx_{A/\rho} y \Longleftrightarrow \rho(x,y)$.

Luckily, it seems that many equivalence relations occurring in computable mathematics are stable, or can be made stable. For example, Cauchy sequences (of rational numbers) $(a_n)_{n\in\mathbb{N}}$ and $(b_n)_{n\in\mathbb{N}}$ represent the same real number when

$$\forall i \in \mathbb{N}.\, \exists j \in \mathbb{N}.\, \forall m, n \geq j.\, |a_m - b_n| \leq 2^{-i}. \tag{4}$$

This defines an equivalence relation on the set of Cauchy sequences which does not seem to be stable; intuitively a realizer for this equivalence would be a computation telling us at what point in the sequence the terms will be within 2^{-i} of each other. However, if we restrict attention just to the *rapid* Cauchy sequences, i.e. those satisfying $\forall i \in \mathbb{N}.\, |a_{i+1} - a_i| \leq 2^{-i}$, then the equivalence relation becomes

$$\forall i \in \mathbb{N}.\, |a_i - b_i| \leq 2^{-i+3},$$

which is a negative formula; the above realizer is rendered unnecessary. It is interesting that most practical implementations of real numbers follow this line of reasoning and represent real numbers in a way that avoids annotating every sequence with its rate of convergence.

Translation of an equivalence class $[e]_\rho$ is quite simple, since a realizer for e also realizes its equivalence class $[e]_\rho$. The elimination term `let` $[x]_\rho = \xi$ `in` e, means 'let x be any element of ρ-equivalence class ξ in e'. It is only well formed when e does not depend on the choice of x, but this is something RZ cannot check. Therefore, if u realizes ξ, RZ uses u as a realizer for x and emits an obligation saying that the choice of a realizer for x does not affect e.

The underlying set of realizers. Another construction on a per A is the underlying per of realizers `rz` A, defined by

$$|\texttt{rz}\, A| = |A|$$
$$u \approx_{\texttt{rz}\, A} v \Longleftrightarrow u \in \|A\| \wedge u = v,$$

where by $u = v$ we mean observational equality of values u and v. An element $r \in \texttt{rz}\, A$ realizes a unique element $\texttt{rz}\, r \in A$. The elimination term `let rz` $x = e_1$ `in` e_2, which means 'let x be any realizer for e_1 in e_2', is only well formed if e_2 does not depend on the choice of x. This is an uncheckable condition, hence RZ emits a suitable obligation in the output, and uses for x the same realizer as for e_1.

The construction `rz` A validates the Presentation Axiom (see Section 7.5). In the input language it gives us access to realizers, which is useful because many constructions in computable mathematics, such as those in Type Two Effectivity [37], are explicitly expressed in terms of realizers.

Definite description. Russell's definite description operator $\iota x{:}s.\phi(x)$ denotes the unique element of $\{x{:}s \mid \phi(x)\}$. In case such an x does not exist, or if there are several, the term is not well formed. The RZ translation essentially just asks the programmer to provide suitable realizers for x and for $\phi(x)$, and to check uniqueness,

$$\texttt{assure } x{:}s, b{:}|\phi|, (x{:}\|s\| \wedge \phi x b \wedge \forall x'{:}s.\forall c{:}|\phi|.(\phi x' c \to x \approx_s x')) \texttt{ in } (x, b).$$

This is the best RZ can do, since in general it can check neither that x exists, nor that it is unique.

5.2   *Translation of propositions*

The driving force behind the translation of logic is a theorem [34, 4.4.10] that says that under the realizability interpretation every formula ϕ is equivalent to one that says, informally speaking, 'there exists $u \in |\phi|$, such that u realizes ϕ'. Furthermore, the formula 'u realizes ϕ' is computationally trivial.

The translation of a predicate ϕ then consists of its underlying type $|\phi|$ and the relation $u \Vdash \phi$, expressed as a negative formula.

Thus an axiom `Axiom A:`ϕ in the input is translated to

```
val u : |φ|
assertion A : u ⊩ φ
```

which requires the programmer to validate ϕ by providing a realizer for it. When ϕ is a compound statement RZ computes the meaning of $u \Vdash \phi$ as described in Figure 2.

In RZ we avoid the explicit realizer notation $u \Vdash \phi$ in order to make the output easier to read. A basic predicate declaration `Parameter` p:s→`Prop` is translated to

```
type ty_p
predicate p : s → ty_p → bool
assertion strict_p : ∀ x:s, a:ty_p, p x a → x : ‖s‖
assertion extensional_p :
    ∀ x:s, y:s, a:ty_p, x ≈ₛ y → p x a → p y a
```

We see that the predicate p has gained an additional argument of type `ty_p` (which the programmer is supposed to define in an implementation), and we write pxa instead of $a \Vdash px$. The two assertions require that p be strict and extensional with respect to $\approx_s$.

Frequently we know that a predicate is stable, which can be taken into account when computing its realizability interpretation. For this purpose the input language has the subkind `Stable` of `Prop`. When RZ encounters a predicate which is declared to be stable, such as p:s→`Stable`, it does not generate a declaration of `ty_p` and does not give p an extra argument.

Another special kind in the RZ input language is the kind `Equiv`(s) of stable equivalence relations on a set s. When an equivalence relation is declared with `Parameter` p:`Equiv`(s), RZ will output assertions stating that p is strict, extensional, reflexive, symmetric and transitive.

6 Implementation

The RZ implementation consists of several sequential passes.

After the initial parsing, a *type reconstruction* phase checks that the input is well-typed (and checks for well-formedness to the extent that it is easily decidable), and if successful produces an annotated result with all variables explicitly tagged with types. The type checking phase uses a system of dependent types, with limited subtyping (implicit coercions) for sum types and subset types. The details are fairly standard, so are omitted here. One non-obvious consequence of the realizability translation, however, is that the subset types with logically equivalent predicates, e.g. $\{x{:}\alpha \mid \rho_1(x) \wedge \rho_2(x)\}$ and $\{x{:}\alpha \mid \rho_2(x) \wedge \rho_1(x)\}$ are isomorphic but not equal in general. An explicit coercion is required to go from one type to the other, because subset values will be pairs containing realizers for $\rho_1(x) \wedge \rho_2(x)$ and $\rho_2(x) \wedge \rho_1(x)$, and these realizers have potentially different types $|\rho_1(x)| * |\rho_2(x)|$ and $|\rho_2(x)| * |\rho_1(x)|$, respectively.

Next the realizability translation is performed as described in Section 5, producing interface code. The flexibility of the full input language (e.g. *n*-ary sum types and dependent product types) makes the translation code fairly involved, and so it is performed in a 'naive' fashion whenever possible. The immediate result of the translation is not easily readable.

Thus, up to four more passes simplify the output before it is displayed to the user. A *thinning* pass removes all references to trivial realizers produced by stable formulas. For example, direct translation

of the `free` axiom in the output for Kuratowski-finite sets, see Figure 7 and Section 7.3, yields a value specification for `free` of type

$$(\texttt{A.a} \rightarrow \texttt{S.s}) \rightarrow (\texttt{fin} \rightarrow \texttt{S.s}) * (\texttt{unit} * (\texttt{A.a} \rightarrow \texttt{unit}) * (\texttt{fin} \rightarrow \texttt{fin} \rightarrow \texttt{unit}))$$

where `unit` is the unit (terminal) type classifying the trivial realizer. Thinning replaces this by the isomorphic type

$$(\texttt{A.a} \rightarrow \texttt{S.s}) \rightarrow \texttt{fin} \rightarrow \texttt{S.s}$$

and appropriately modifies references to `free` in the assertions to account for this change in type.

An *optimization* pass applies an ad hoc collection of basic logical and term simplifications in order to make the output more readable. Logical simplifications include applications of truth table rules ($\top \wedge \varphi$ becomes φ), detection of syntactically identical premises and conclusions ($\varphi_1 \Rightarrow \varphi_1 \wedge \varphi_2$ becomes $\varphi_1 \Rightarrow \varphi_2$), and optimization of other common patterns we have seen arise ($\forall x{:}s.\ (x = e) \Rightarrow \rho(x)$ becomes $\rho(e)$). Some redundancy may remain, but in practice the optimization pass helps significantly.

Finally, the user can specify two optional steps that occur. RZ can perform a *phase-splitting* pass [18]. This is an experimental implementation of a transformation that can replace a functor (a relatively heavyweight language construct) by parameterized types and/or polymorphic values. The idea is that although functors map modules containing types and terms to other modules containing types and terms, constraints on the programming language ensure that output types depend only on input types (and not input terms). Thus, we can split each functor into a mapping from input types to output types, and then a separate (polymorphic) term mapping input types and terms to an output term. See Section 7.3 for an example.

The other optional transformation is a *hoisting* pass which moves obligations in the output to top-level positions. Obligations appear in the output inside assertions, at the point where an uncheckable property was needed. Moving these obligations to the top-level make it easier to see exactly what one is obliged to verify, and can sometimes make them easier to read, at the cost of losing information about why the obligation was required at all. See Section 7.2 for an example of hoisting.

7 Examples

In this section, we look at several small examples which demonstrate various points of RZ. For a serious case study from computable mathematics see the implementation of real numbers with RZ by Bauer and Kavkler [3].

The main theme is that constructively reasonable axioms yield computationally reasonable operations.

7.1 Decidable sets

A set S is said to be decidable when, for all $x, y \in S$, $x = y$ or $\neg(x = y)$. In classical mathematics all sets are decidable, because decidability of equality is just an instance of the law of excluded middle. But RZ requires an axiom

```
Parameter s : Set.
Axiom eq: ∀ x y : s, x = y ∨ ¬ (x = y).
```

to produce a realizer for equality

```
val eq : s → s → [`or0 | `or1]
assertion eq : ∀ (x:‖s‖, y:‖s‖), (match eq x y with
                                        `or0 ⇒ x ≈ₛ y
                                      | `or1 ⇒ ¬ (x ≈ₛ y) )
```

We read this as follows: eq is a function which takes arguments x and y of type s and returns `or0 or `or1. If it returns `or0, then x≈ₛy, and if it returns `or1, then ¬(x≈ₛy). In other words eq is a decision procedure which tells when values x and y represent the same element of the modest set.

7.2   Examples with obligations

In this section we show how RZ produces obligations, is sometimes able to optimize them away, and show the effect of hoisting.

Consider how we might define division of real numbers. Assuming the set of real numbers real, constants zero and one and multiplication operation ∗ have already been declared and axiomatized, we might write:

```
Definition nonZeroReal := {x : real | ¬ (x = zero)}.
Parameter inv : nonZeroReal → real.
Axiom inverse : ∀ x : real, ¬ (x = zero) -> x * (inv x) = one.
Definition (/) (x : real) (y : nonZeroReal) := x * (inv y).
```

We have defined the set of non-zero reals nonZeroReal and the inverse operation inv on it. Division x/y is defined as x∗inv y. This does *not* mean that the programmer must necessarily implement division this way, only that the implementation of x/y must be equivalent to x∗inv y.

In the axiom inverse, RZ encounters the subexpression inv x. Because x is quantified as an element of real rather than nonZeroReal, the typechecking phase inserts a coercion that makes the expression well-typed. Translation sees inv(x:nonZeroReal) instead of inv x and translates this to

```
inv (assure u:unit, ¬ (x ≈_real zero) in (x, u))
```

If this were the final output, the programmer would have to verify that x is not zero, and provide a trivial realizer for it. However, in this case the thinning phase first removes the trivial realizer,

```
inv (assure ¬ (x ≈_real zero) in x)
```

and then the optimizer determines that the obligation is not needed because the whole expression appears under the hypothesis that x is not zero. So in the end the programmer sees

```
assertion inverse :
   ∀ (x:‖real‖),  ¬ (x ≈_real zero) → (x * inv x) ≈_real one
```

Assuming further that a strict linear order < on real has been axiomatized, we might proceed by relating it to inv:

```
Axiom inv_positive: ∀ x : real, zero < x → zero < inv x.
```

Once again `inv x` appears in the input, but this time the optimizer is unable to remove the obligation, so the output is

```
assertion inv_positive: ∀ (x:‖real‖),
   zero < x → zero < inv (assure (not (x ≈real zero)) in x)
```

Local obligations can sometimes be hard to read, but if we activate the hoisting phase (see Section 6), the obligation can be moved to the top level. As this is done, the hypotheses under which the obligation appears are collected, and we get

```
assertion inv_positive:
  assure (∀ (x:‖real‖),  zero < x → not (x ≈real zero))
    in ∀ (x:‖real‖),  zero < x → zero < inv x
```

Now it is easier to understand what must be checked, namely that positive reals are not zero—an easy consequence of irreflexivity of $<$, but not something that the RZ optimizer is aware of.

Lastly, we could define the golden ratio as the positive solution of $x^2 = x + 1$,

```
the x : real, (zero < x ∧ x*x = x + one)
```

Not surprisingly, RZ cannot determine that there is a unique such x, so it outputs an obligation:

```
assure x:real,
  (x : ‖real‖ ∧ zero < x ∧ x * x =real= x + one ∧
    (∀ (x':‖real‖),  zero < x' ∧ x' * x' ≈real x' + one → x ≈real x'))
  in x
```

7.3   Finite sets

There are many characterizations of finite sets, but the one that works best constructively is due to Kuratowski, who identified the finite subsets of *A* as the least family *K(A)* of subsets of *A* that contains the empty set and is closed under unions with singletons. This characterization relies on powersets, which are not available in RZ. But the gist of it, namely that *K(A)* is an *initial* structure a suitable sort, can be expressed as follows.

Recall that a ∨-*semilattice* is a set *S* with a constant $0 \in S$ and an associative, commutative and idempotent operation 'join' ∨ on *S* such that 0 is the neutral element for ∨, see Figure 6 for RZ axiomatization of semilattices. The Kuratowski finite sets *K(A)* are the *free* semilattice generated by a set *A*, where ∨ is union and 0 is the empty set. This is formalized in RZ as shown in Figure 7. The theory K is parameterized by a model A which contains a set a. In the first line we include the theory

```
Definition Semilattice :=
thy
  Parameter s : Set.
  Parameter zero : s.
  Parameter join : s → s → s.
  Implicit Type x y z : s.
  Axiom commutative: ∀ x y,   join x y = join y x.
  Axiom associative: ∀ x y z, join (join x y) z = join x (join y z).
  Axiom idempotent:  ∀ x,     join x x = x.
  Axiom neutral:     ∀ x,     join x zero = x.
end.
```

FIGURE 6. The theory of a semilattice

of semilattices. Then we postulate an operation `singleton` which injects the generators into the semilattice. The three definitions are just a convenience, so that we can refer to the parts of K(A) by their natural names, e.g. `emptyset` instead of `zero`. The axiom `free` expresses the fact that K(A) is the free semilattice on A.a: for every semilattice S and a map $f : A.a \to S.s$ from the generators to the underlying set of S, there exists a unique semilattice homomorphism $g : fin \to S.s$ such that $f(x) = g(singleton\ x)$.

The output for `Semilattice` and `K` specifies values of suitable types for each declared constant and operation. All axioms but the last one are equations and have straightforward translations in terms of underlying pers. The output for the axiom `free` is shown in Figure 8. Because the axiom quantifies over all models S of the theory `Semilattice` its translation is a functor `Free` which accepts an implementation of a semilattice *S* and yields a realizer `free` validating the axiom. The computational meaning of `free` is a combination map and fold operation, taking a map $f : A.a \to S.s$ and a finite set $u = \{x_1, \ldots, x_n\}$, and return $f(x_1) \vee \cdots \vee f(x_n)$, where $\vee$ is the join operation on the semilattice *S*.

```
Definition K (A : thy
                      Parameter a : Set.
              end) :=
thy
  include Semilattice.
  Parameter singleton : A.a → s.
  Definition fin := s.
  Definition emptyset := zero.
  Definition union := join.

  Axiom free :
    ∀ S : Semilattice, ∀ f : A.a → S.s,
    ∃! g : fin → S.s,
      g emptyset = S.zero ∧
        (∀ x : A.a, f x = g (singleton x)) ∧
        (∀ u v : fin, g (union u v) = S.join (g u) (g v)).
end.
```

FIGURE 7. Kuratowski finite sets

```
module Free : functor (S : Semilattice) →
sig
  val free : (A.a → S.s) → fin → S.s
  assertion free :
    ∀ (f:‖A.a → S.s‖), let g = free f in
      g : ‖fin → S.s‖ ∧ g emptyset ≈_{S.s} S.zero ∧
      (∀ (x:‖A.a‖),  f x ≈_{S.s} g (singleton x)) ∧
      (∀ (u:‖fin‖, v:‖fin‖), g (union u v) ≈_{S.s} S.join (g u) (g v)) ∧
      (∀ h:fin → S.s,  h : ‖fin → S.s‖ ∧ h emptyset ≈_{S.s} S.zero ∧
         (∀ (x:‖A.a‖), f x ≈_{S.s} h (singleton x)) ∧
         (∀ (u:‖fin‖, v:‖fin‖), h (union u v) ≈_{S.s} S.join (h u) (h v)) →
         ∀ x:fin, y:fin,  x ≈_{fin} y → g x ≈_{S.s} h y)
end
```

FIGURE 8. Output of axiom `free`

Applying phase-splitting to this axiom yields the even simpler specification

$$\texttt{val free}: \alpha \to (\alpha \to \alpha \to \alpha) \to (\texttt{A.a} \to \alpha) \to \texttt{fin} \to \alpha$$

(with an appropriate assertion) that replaces the module parameter `S` by two extra term arguments (corresponding to the module components `S.zero` and `S.join`) and a type argument α for the type of lattice elements (corresponding to the module input `S.s`). This is even more recognizable as a folding operation over the set.

It is important to note that, in contrast to `fold` operators found in typical functional languages, `free` is only expected to work for suitable `join` arguments (e.g. idempotent and order independent). These sets are not the typical finite set data structure: there is no membership predicate, nor is there a way to compute the size of a set. There is no assumption that equality is decidable for set elements; this permits finite sets of exact real numbers, for example. Decidable equality is required both for membership and for detecting whether the same element has been added twice to the same set.[12]

Some operations are nevertheless computable. Using `free` carefully, one can determine whether a finite set is empty. In the case of a set of exact real numbers, we cannot compute the sum of a set (since there might be duplicate elements), but we could compute maximum or minimum.

More common set implementations (e.g. the `Set` module in the OCaml standard library) implement sets over values with decidable total order; these could also be formalized in RZ.

7.4   *Inductive types*

To demonstrate the use of dependent types we show how RZ handles general inductive types, also known as W-types or general trees [30]. Recall that a W-type is a set of well-founded trees, where the branching types of trees are described by a family of sets $B = \{T(x)\}_{x \in S}$. Each node in a tree has a *branching type* $x \in S$, which determines that the successors of the node are labelled by the elements of $T(x)$. For example, to get non-empty binary trees whose leaves are labelled by natural numbers, define

$$S = \{\texttt{cons}\} \cup \{\texttt{leaf}(n) \,|\, n \in \mathbb{N}\}$$
$$T(\texttt{cons}) = \{\texttt{left}, \texttt{right}\}$$
$$T(\texttt{leaf}(n)) = \emptyset.$$

Then a node of type `cons` has two successors, indexed by constants `left` and `right`, while a node of type `leaf`(n) does not have any successors.

Figure 9 shows an RZ axiomatization of W-types. The theory `Branching` describes that a branching type consists of a set `s` and a set `t` depending on `s`. The theory `W` is parameterized by a branching type `B`. It specifies a set `w` of well-founded trees and a tree-forming operation `tree` with a dependent type $\Pi_{x \in \texttt{B.s}}(\texttt{B.t}(x) \to \texttt{w}) \to \texttt{w}$. Given a branching type `x` and a map `f:B.t(x)→w`, `tree x f` is the tree whose root has branching type `x` and whose successor labelled by $\ell \in \texttt{B.t}(x)$ is the tree $f(\ell)$. The inductive nature of `w` is expressed with the axiom `induction`, which states that for every property `M.p`, if `M.p` is an inductive property then every tree satisfies it. A property is said to be *inductive* if a tree `tree x f` satisfies it whenever all its successors satisfy it.

In the translation (see Appendix A for the complete output) dependencies at the level of types and terms disappear. A branching type is determined by a pair of non-dependent types `s` and `t`

---

[12]The natural implementation would thus be an unordered collection of elements, possibly with duplicates.

```
Definition Branching :=
thy
  Parameter s : Set.      (* branching types *)
  Parameter t : s -> Set. (* branch labels *)
end.

Parameter W : [B : Branching] →
thy
  Parameter w : Set.
  Parameter tree : [x : B.s] → (B.t x → w) → w.
  Axiom induction:
    ∀ M : thy Parameter p : w → Prop. end,
    (∀ x : B.s, ∀ f : B.t x → w,
       ((∀ y : B.t x, M.p (f y)) → M.p (tree x f))) →
    ∀ t : w, M.p t.
end.
```

FIGURE 9. General inductive types

but the per $\approx_t$ depends on $[\![s]\!]$. The theory W turns into a signature for a functor receiving a branching type B and returning a type w, and an operation `tree` of type $B.s \to (B.t \to w) \to w$. One can use phase-splitting to translate axiom `induction` into a specification of a polymorphic function

$$\texttt{induction:}(B.s \to (B.t \to w) \to (B.t \to \alpha) \to \alpha) \to w \to \alpha,$$

which is a form of recursion on well-founded trees. Instead of explaining `induction`, we show a surprisingly simple, handwritten implementation of W-types in OCaml. The reader may enjoy figuring out how it works:

```
module W (B : Branching) = struct
  type w = Tree of B.s * (B.t -> w)
  let tree x y = Tree (x, y)
  let rec induction f (Tree (x, g)) =
    f x g (fun y -> induction f (g y))
end
```

7.5 Axiom of choice

RZ can help explain why a generally accepted axiom is not constructively valid. Consider the Axiom of Choice:

```
Parameter a b : Set.
Parameter r : a → b → Prop.
Axiom ac: (∀ x : a, ∃ y : b, r x y) →
          (∃ c : a → b, ∀ x : a, r x (c x)).
```

The relevant part of the output is

```
val ac : (a → b * ty_r) → (a → b) * (a → ty_r)
assertion ac :
  ∀ f:a → b * ty_r,
    (∀ (x:‖a‖),  let (p,q) = f x in p : ‖b‖ ∧ r x p q) →
    let (g,h) = ac f in
      g : ‖a → b‖ ∧ (∀ (x:‖a‖),  r x (g x) (h x))
```

This requires a function `ac` which accepts a function `f` and computes a pair of functions (`g`,`h`). The input function `f` takes an `x`:‖`a`‖ and returns a pair (`p`,`q`) such that `q` realizes the fact that `r x p` holds. The output functions `g` and `h` taking `x`:‖`a`‖ as input must be such that `h x` realizes `r x (g x)`. Crucially, the requirement `g`:‖`a`→`b`‖ says that `g` must be extensional, i.e. map equivalent realizers to equivalent realizers. We could define `h` as the second component of `f`, but we cannot hope to implement `g`, in general, because the first component of `f` is not assumed to be extensional.

The *Intensional* Axiom of Choice allows the choice function to depend on the realizers:

```
Axiom iac: (∀ x : a, ∃ y : b, r x y) →
           (∃ c : rz a → b, ∀ x : rz a, r (rz x) (c x)).
```

Now the output is

```
val iac : (a → b * ty_r) → (a → b) * (a → ty_r)
assertion iac :
  ∀ f:a → b * ty_r,
    (∀ (x:‖a‖),  let (p,q) = f x in p : ‖b‖ ∧ r x p q) →
    let (g,h) = iac f in
      (∀ x:a, x : ‖a‖ → g x : ‖b‖) ∧ (∀ (x:‖a‖),  r x (g x) (h x))
```

This is exactly the same as before *except* the troublesome requirement `g`:‖`a`→`b`‖ was weakened to ∀`x`:`a`.(`x`:‖`a`‖→`g x`:‖`b`‖). We can implement `iac` in OCaml as

```
let iac f = (fun x -> fst (f x)), (fun x -> snd (f x))
```

The Intensional Axiom of Choice is in fact just an instance of the usual Axiom of Choice applied to `rz A` and `B`. Combined with the fact that `rz A` covers `A`, this establishes the validity of *Presentation Axiom* [1], which states that every set is an image of one satisfying the axiom of choice.

7.6 *Modulus of continuity*

As a last example we show how certain constructive principles require the use of computational effects. To keep the example short, we presume that we are already given the set of natural numbers `nat` with the usual structure.

A *type 2 functional* is a map $f:(\text{nat}\to\text{nat})\to\text{nat}$. It is said to be continuous if the output of $f(a)$ depends only on an initial segment of the sequence a. We can express the (non-classical) axiom that all type 2 functionals are continuous in RZ as follows:

```
Axiom continuity: ∀ f : (nat → nat) → nat, ∀ a : nat → nat,
   ∃ k, ∀ b : nat → nat, (∀ m, m ≤ k → a m = b m) → f a = f b.
```

The axiom says that for any `f` and `a` there exists `k`∈`nat` such that `f(b)`=`f(a)` when sequences `a` and `b` agree on the first `k` terms. It translate to:

```
val continuity : ((nat → nat) → nat) → (nat → nat) → nat
assertion continuity :
  ∀ (f:‖(nat → nat) → nat‖, a:‖nat → nat‖),
    let p = continuity f a in p : ‖nat‖ ∧
    (∀ (b:‖nat → nat‖),
       (∀ (m:‖nat‖),  m ≤ p → a m ≈_nat b m) → f a ≈_nat f b)
```

i.e. that `continuity f a` is a number `p` such that `f(a) = f(b)` whenever `a` and `b` agree on the first `p` terms. In other words, `continuity` is a *modulus of continuity* functional. It cannot be implemented in a purely functional language,[13] but with the use of store we can implement it in OCaml as

```
let continuity f a = let p = ref 0 in
                     let a' n = (p := max !p n; a n) in
                       f a' ; !p
```

To compute a modulus for `f` at `a`, the program creates a function `a'` that is just like `a` except that it stores in `p` the largest argument at which it has been called. Then `f a'` is computed, its value is discarded, and the value of `p` is returned. The program works because `f` is assumed to be extensional and therefore must not distinguish between extensionally equal sequences `a` and `a'`.

8 Related work

8.1 Coq and other tools

Coq provides complete support for theorem-proving and creating trusted code. Often one writes code in Coq's functional language (values whose types are `Sets`), states and proves theorems that the code behaves correctly (where the theorems are Coq values whose types are `Props`), and has Coq extract correct code. In such cases RZ is complementary to Coq; it can clarify the constructive content of mathematical structures and hence suggests the appropriate division between code and theorems. We hope RZ will soon be able to produce output in Coq's input syntax.

In general, RZ is a smaller and more lightweight system and thus more flexible where it applies. It is not always practical or necessary to do theorem proving in order to provide an implementation; interfaces generated by RZ can be implemented in any manner. And, RZ provides a way to talk with programmers about constructive mathematics without bringing in full theorem proving.

Komagata and Schmidt [23] describe a system that uses a realizability in a way similar to RZ. Like Coq, it extracts code from proofs. An interesting implementation difference is that the algorithm they use (attributed to John Hatcliff) does thinning as it goes along, rather than making a separate pass as RZ does. (For example, the translation of the conjunction–introduction rule has four cases, depending on whether the left and/or right propositions being proved are almost negative, in which case the trivial contribution can be immediately discarded.) Unlike RZ, their system needs full formal proofs as input; it checks the proofs, and generates executable code. RZ also handles a much richer input language (including function, subset, quotient and dependent types; quantification over theories; and parameterized theories) that goes well beyond simple predicate logic over integers and lists.

The idea of annotating ML signatures with assertions is not new (e.g. [20]).

8.2 Other models of computability

Many formulations of computable mathematics are based on realizability models [2], even though they were not initially developed, (nor are they usually presented) within the framework of realizability: Recursive Mathematics [16] is based on the original realizability by Turing machines [21]; Type Two Effectivity [37] on function realizability [22] and relative function

---

[13]There are models of λ-calculus which validate the choice principle $AC_{2,0}$, but this contradicts the existence of a modulus of continuity functional, see [35, 9.6.10].

realizability [9], while topological and domain representations [6, 11] are based on realizability over the graph model $\mathcal{P}\omega$ [33]. A common feature is that they use models of computation which are well suited for the theoretical studies of computability.

Other approaches are based on simple programming languages augmented with datatypes for real numbers [17, 28] and topological algebras [36], or machine models augmented with (suitably chosen subsets of) real numbers such as Real RAM [14], the Blum–Shub–Smale model [13] and the Exact Geometric Computation model [38]. The motivation behind these ranges from purely theoretical concerns about computability and complexity to practical issues in the design of programming languages and algorithms in computational geometry. RZ attempts to improve practicality by using an actual real-world programming language, and by providing an input language which is rich enough to allow descriptions of involved mathematical structures that go well beyond the real numbers.

Finally, we hope that RZ and, hopefully, its forthcoming applications, give plenty of evidence for the *practical* value of Constructive Mathematics [10].

References

[1] J. Barwise. *Admissible Sets and Structures*. Springer, Berlin, 1975.

[2] A. Bauer. *The Realizability Approach to Computable Analysis and Topology*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, USA, 2000.

[3] A. Bauer and I. Kavkler. Implementing real numbers with RZ. In *Fourth International Conference on Computability and Complexity in Analysis*, K. Weihrauch and N. Zhong, eds. Electronic Notes in Theoretical Computer Science, 2007. Forthcoming.

[4] A. Bauer and C. A. Stone. Specifications via realizability. In *Proceedings of the Workshop on the Constructive Logic for Automated Software Engineering (CLASE 2005) in Edinburgh, Scotland*, Vol. 153 of *Electronic Notes in Theoretical Computer Science*, pp. 77–92, 2006.

[5] A. Bauer and C. A. Stone. RZ: a tool for bringing constructive and computable mathematics closer to programming practice. In *Computation and Logic in the Real World, Third Conference on Computability in Europe, in Siena, Italy,*, Vol. 4497 of *LNCS*, pp. 28–42, Springer, Berlin, 2007.

[6] A. Bauer, L. Birkedal, and D. S. Scott. Equilogical spaces. *Theoretical Computer Science*, **1**, Elsevier, Amsterdam, 35–59, 2004.

[7] H. Benl, U. Berger, H. Schwichtenberg, M. Seisenberger, and W. Zuber. Proof theory at work: program development in the Minlog system. In *Automated Deduction: A Basis for Applications. Volume II, Systems and Implementation Techniques*, W. Bibel and P. H. Schmidt, eds. Kluwer Academic Publishers, Dordrecht, 1998.

[8] Y. Bertot and P. Castéran. *Interactive Theorem Proving and Program Development*. Springer, Berlin, 2004.

[9] L. Birkedal. *Developing Theories of Types and Computability*. PhD thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA, 1999.

[10] E. Bishop and D. Bridges. *Constructive Analysis*, Vol. 279 of *Grundlehren der math. Wissenschaften*. Springer, Berlin, 1985.

[11] J. Blanck. *Computability on Topological Spaces by Effective Domain Representations*. PhD thesis, Uppsala University, Department of Mathematics, Uppsala, Sweden, 1997.

[12] J. Blanck. Domain representability of metric spaces. *Annals of Pure and Applied Logic*, **83**, 225–247, 1997.

[13] L. Blum, F. Cucker, M. Shub, and S. Smale. *Complexity and Real Computation*. Springer, New York, 1998.

[14] A. Borodin and J. I. Monro. *The Computational Complexity of Algebraic and Numeric Problems*. Number 1 in Elsevier computer science library: Theory of computation series. American Elsevier, New York, London, Amsterdam, 1975.

[15] A. Edalat and A. Lieutier. Domain theory and differential calculus (functions of one variable). *Mathematical Structures in Computer Science*, **14**, 771–802, 2004.

[16] Y. L. Ershov, S. S. Goncharov, A. Nerode, and J. B. Remmel, eds. *Handbook of Recursive Mathematics*. Elsevier, Amsterdam, 1998.

[17] M. H. Escardó. *PCF Extended with Real numbers*. PhD thesis, Department of Computer Science, University of Edinburgh, Edinburgh, Scotland, UK, 1997.

[18] R. Harper, J. C. Mitchell, and E. Moggi. Higher-order modules and the phase distinction. In *Proceedings of the 17th ACM Symposium on Principles of Programming Languages (POPL '90) in San Francisco*, pp. 341–354, ACM, Berlin, 1990.

[19] B. Jacobs. *Categorical Logic and Type Theory*. Elsevier Science, North-Holland, Amsterdam, 1999.

[20] S. Kahrs, D. Sannella, and A. Tarlecki. The definition of Extended ML: a gentle introduction. *Theoretical Computer Science*, **173**, 445–484, 1997.

[21] S. C. Kleene. On the interpretation of intuitionistic number theory. *Journal of Symbolic Logic*, **10**, 109–124, 1945.

[22] S. C. Kleene and R. E. Vesley. *The Foundations of Intuitionistic Mathematics, Especially in Relation to Recursive Functions*. North-Holland Publishing Company, North-Holland, Amsterdam, 1965.

[23] Y. Komagata and D. A. Schmidt. Implementation of intuitionistic type theory and realizability theory. *Technical Report TR-CS-95-4*. Kansas State University, Manhattan, KS, USA, 1995.

[24] B. Lambov. RealLib: an efficient implementation of exact real arithmetic. *Mathematical Structures in Computer Science*, **17**, 81–98, 2007.

[25] X. Leroy, D. Doligez, J. Garrigue, D. Rémy, and J. Vouillon. The Objective Caml system, documentation and user's manual - release 3.08. *Technical report*. INRIA, Rocquencourt, France, 2004.

[26] J. Longley. Matching typed and untyped realizability. *Electronic Notes in Theoretical Computer Science*, **23**, 74–100, 1999.

[27] J. Longley. When is a functional program not a functional program? In *International Conference on Functional Programming (ICFP '99) in Paris*, pp. 1–7, ACM, 1999.

[28] J. R. Marcial-Romero and M. H. Escardó. Semantics of a sequential language for exact real-number computation. In *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science (LICS '04) in Turku, Finland*, pp. 426–435, IEEE, 2004.

[29] N. Müller. The iRRAM: exact arithmetic in C++. In *Computability and Complexity in Analysis: 4th International Workshop, CCA 2000 Swansea, UK, September 17, 2000, Selected Papers*, Number 2064 in *Lecture Notes in Computer Science*, J. Blanck, V. Brattka, and P. Hertling, eds. pp. 222–252. Springer, Berlin, 2001.

[30] B. Nordström, K. Petersson, and J. M. Smith. *Programming in Martin-Löf's Type Theory*. Oxford University Press, USA, 1990.

[31] E. Post. Recursive unsolvability of a problem of Thue. *The Journal of Symbolic Logic*, **12**, 1–11, 1947.

[32] D. Sannella, S. Sokolowski, and A. Tarlecki. Toward formal development of programs from algebraic specifications: parameterisation revisited. *Acta Informatica*, **29**, 689–736, 1992.

[33] D. S. Scott. Data types as lattices. *SIAM Journal of Computing*, **5**, 522–587, 1976.

[34] A. S. Troelstra and D. van Dalen. *Constructivism in Mathematics, An Introduction*, Vol. 1. Number 121 in Studies in Logic and the Foundations of Mathematics. North-Holland, Amsterdam, 1988.

[35] A. S. Troelstra and D. van Dalen. *Constructivism in Mathematics, An Introduction*, Vol. 2. Number 123 in Studies in Logic and the Foundations of Mathematics. North-Holland, 1988.

[36] J. V. Tucker and J. I. Zucker. Computable functions and semicomputable sets on many-sorted algebras. In *Handbook of Logic in Computer Science*, Vol. 5. S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum, eds, Clarendon Press, Oxford, 1998.

[37] K. Weihrauch. *Computable Analysis*. Springer, Berlin, 2000.

[38] C. K. Yap. Theory of real computation according to EGC, 2006. In *Dagstuhl Seminar "Reliable Implementation of Real Number Algorithms: Theory and Practice"*, *LNCS in 2006 in Dagstuhl, Germany*, January 8–13, 2006.

Appendix

A   An interface for the theory of general inductive types

To give at least one complete example, we include here an unabridged output for the theory of inductive types shown in Figure 9.

```
module type Branching =
 sig
   type s

   (** predicate (=s=) : s -> s -> bool *)
   (**  assertion symmetric_s :  forall x:s, y:s,  x =s= y -> y =s= x

        assertion transitive_s :
          forall x:s, y:s, z:s,  x =s= y /\ y =s= z -> x =s= z
   *)

   (** predicate ||s|| : s -> bool *)
   (**  assertion total_def_s :  forall x:s,  x : ||s|| <-> x =s= x
   *)

   (** branching types *)

   type t

   (** predicate (=t=) : s -> t -> t -> bool *)
   (**  assertion strict_t :  forall x:s, y:t, z:t,  y =(t x)= z -> x : ||s||

        assertion extensional_t :
          forall x:s, y:s, z:t, w:t,  x =s= y -> z =(t x)= w -> z =(t y)= w

        assertion symmetric_t :
          forall x:s, y:t, z:t,  y =(t x)= z -> z =(t x)= y

        assertion transitive_t :
          forall x:s, y:t, z:t, w:t,  y =(t x)= z /\ z =(t x)= w ->
            y =(t x)= w
```

```
    *)

    (** predicate ||t|| : s -> t -> bool *)
    (**  assertion total_def_t :
           forall x:s, y:t,  y : ||t x|| <-> y =(t x)= y
    *)

    (** branch labels *)
  end

module W : functor (B : Branching) ->
sig
  type w

  (** predicate (=w=) : w -> w -> bool *)
  (**  assertion symmetric_w :
          forall x:w, y:w,  x =w= y -> y =w= x

       assertion transitive_w :
          forall x:w, y:w, z:w,  x =w= y /\ y =w= z -> x =w= z
  *)

  (** predicate ||w|| : w -> bool *)
  (**  assertion total_def_w :  forall x:w,  x : ||w|| <-> x =w= x
  *)

  val tree : B.s -> (B.t -> w) -> w
  (**  assertion tree_support :
          forall x:B.s, y:B.s,  x =B.s= y ->
            forall f:B.t -> w, g:B.t -> w,
              (forall z:B.t, t:B.t,  z =(B.t x)= t -> f z =w= g t) ->
              tree x f =w= tree y g
  *)

  val induction : (B.s -> (B.t -> w) -> (B.t -> 'ty_p) -> 'ty_p) -> w -> 'ty_p
  (**  assertion 'ty_p [p:w -> 'ty_p -> bool] induction :
          (forall x:w, a:'ty_p,  p x a -> x : ||w||) ->
          (forall x:w, y:w, a:'ty_p,  x =w= y -> p x a -> p y a) ->
          forall f:B.s -> (B.t -> w) -> (B.t -> 'ty_p) -> 'ty_p,
            (forall (x:||B.s||),
               forall f':B.t -> w,
                 (forall y:B.t, z:B.t,  y =(B.t x)= z ->
                    f' y =w= f' z) ->
                 forall g:B.t -> 'ty_p,
                   (forall y:B.t,  y : ||B.t x|| -> p (f' y) (g y)) ->
                   p (tree x f') (f x f' g)) ->
            forall (t:||w||),  p t (induction f t)
  *)
end
```