

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

Andrej Bauer¹ and Christopher A. Stone²

¹ Faculty of Mathematics and Physics, University of Ljubljana, Slovenia
`Andrej.Bauer@fmf.uni-lj.si`

² Computer Science Department, Harvey Mudd College, USA
`stone@cs.hmc.edu`

Abstract. Realizability theory can produce code interfaces for the data structure corresponding to a mathematical theory. Our tool, called RZ, serves as a bridge between constructive mathematics and programming by translating 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, like groups, the answer is obvious. 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 [1–4], while implementations of exact real arithmetic [5, 6] show that theory can be put into practice quite successfully.

Realizability theory can be used to produce a description of the data structure (a code interface) directly corresponding to a mathematical specification. But few programmers — even those with strong backgrounds in mathematics and classical logic — are familiar with constructive logic or realizability.

We have therefore implemented a system, called RZ, to serve as a bridge between the logical world and the programming world.³ RZ translates specifications in constructive logic into standard interface code in a programming language (currently Objective Caml [7], 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

³ RZ is publicly available for download at <http://math.andrej.com/rz/>, together with an extended version of this paper .

these types and values. As these assertions have no computational content, so 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 [8], we have not implemented this. Other excellent systems, including Coq [9] and Minlog [10], 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. Or, the output can serve as a basis for theorem-proving and code extraction using another system.

An earlier description of our RZ work appears in [11]; 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 John Longley [12]. 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 directed graphs with vertices labeled by distinct integers. 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. Thus we define the datatype of graphs as⁴

```
type graph = int list * (int * int) list
```

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., the order in which vertices and arrows are listed is not important, each vertex and arrow must be listed exactly once, and 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.

⁴ We use OCaml notation in which `t list` classifies finite lists of elements of type `t`, and `t1 * t2` classifies pairs containing a value of type `t1` and a value of type `t2`.

Let \mathbf{Type} be the collection of all (non-parametric) OCaml types. To each type $t \in \mathbf{Type}$ we assign the set $\llbracket t \rrbracket$ of values of type t that behave *functionally* in the sense of Longley [13]. 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 7.4. A functional value of function type may diverge as soon as it is applied. The collection \mathbf{Type} with the assignment of functional values $\llbracket t \rrbracket$ to each $t \in \mathbf{Type}$ forms a *typed partial combinatory algebra (TPCA)*.

Going back to our example, we see that an implementation of directed graphs \mathcal{G} specifies a datatype $|\mathcal{G}| = \mathbf{graph}$ together with a *realizability relation* $\Vdash_{\mathcal{G}}$ between \mathcal{G} and $\llbracket \mathbf{graph} \rrbracket$. The meaning of $(\ell_V, \ell_A) \Vdash_{\mathcal{G}} G$ is “OCaml value (ℓ_V, ℓ_A) represents/realizes/implements graph G ”. Generalizing from this, we define a *modest set* to be a triple $A = (\langle A \rangle, |A|, \Vdash_A)$ where $\langle A \rangle$ is the *underlying set*, $|A| \in \mathbf{Type}$ is the *underlying type*, and \Vdash_A is a *realizability relation* between $\llbracket |A| \rrbracket$ and $\langle A \rangle$, satisfying (1) *totality*: for every $x \in \langle A \rangle$ there is $v \in \llbracket |A| \rrbracket$ 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 $\|A\| = \{v \in \llbracket |A| \rrbracket \mid \exists x \in \langle A \rangle . v \Vdash_A x\}$ of those values that realize something. We define the relation \approx_A on $\llbracket |A| \rrbracket$ by

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

From totality and modesty of \Vdash_A it follows that \approx_A is a per, i.e., symmetric and transitive. Observe that $\|A\| = \{v \in \llbracket |A| \rrbracket \mid v \approx_A v\}$, 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.⁵ 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.

Pers form a category whose objects are pairs $A = (|A|, \approx_A)$ where $|A| \in \mathbf{Type}$ and \approx_A is a per on $\llbracket |A| \rrbracket$. A morphism $A \rightarrow B$ is represented by a function $v \in \llbracket |A| \rrbracket \rightarrow \llbracket |B| \rrbracket$ such that, for all $u, u' \in \|A\|$, $u \approx_A u' \implies v u \approx_B v u'$. Two such functions v and v' 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, namely that of a regular locally cartesian closed category [14]. This suffices for the interpretation of first-order logic and (extensional) dependent types [15].

Not all pers are *decidable*, i.e., there may be no algorithm for deciding when two values are equivalent. Examples include implementations of semigroups with an undecidable word problem [16] and implementations of computable real numbers (which might be realized by infinite Cauchy sequences).

⁵ And there is a third view, as a partial surjection $\delta_A : \subseteq \llbracket |A| \rrbracket \rightarrow \langle A \rangle$, with $\delta_A(v) = x$ when $v \Vdash_A x$. This is how realizability is presented in Type Two Effectivity [1].

Underlying types of realizers:

$ \top $	$= \mathbf{unit}$	$ \perp $	$= \mathbf{unit}$
$ x = y $	$= \mathbf{unit}$	$ \phi \wedge \psi $	$= \phi \times \psi $
$ \phi \Rightarrow \psi $	$= \phi \rightarrow \psi $	$ \phi \vee \psi $	$= \text{'or}_0 \text{ of } \phi_0 + \text{'or}_1 \text{ of } \phi_1 $
$ \forall x:A. \phi $	$= A \rightarrow \phi $	$ \exists x:A. \phi $	$= A \times \phi $

Realizers:

$() \Vdash \top$	
$() \Vdash x = y$	iff $x = y$
$(t_1, t_2) \Vdash \phi \wedge \psi$	iff $t_1 \Vdash \phi$ and $t_2 \Vdash \psi$
$t \Vdash \phi \Rightarrow \psi$	iff for all $u \in \phi $, if $u \Vdash \phi$ then $t u \Vdash \psi$
$\text{'or}_0 t \Vdash \phi \vee \psi$	iff $t \Vdash \phi$
$\text{'or}_1 t \Vdash \phi \vee \psi$	iff $t \Vdash \psi$
$t \Vdash \forall x:A. \phi(x)$	iff for all $u \in A $, if $u \Vdash_A x$ then $t u \Vdash \phi(x)$
$(t_1, t_2) \Vdash \exists x:A. \phi(x)$	iff $t_1 \Vdash_A x$ and $t_2 \Vdash \phi(x)$

Fig. 1. Realizability interpretation of logic (outline)

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 (for example, if they denote partial functions or use computational effects).

It is customary to write $t \Vdash \phi$ when $t \in \llbracket |\phi| \rrbracket$ 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 1. We say that a formula ϕ is *valid* if it has at least one realizer.

In classical mathematics, a predicate on a set X may be viewed as a subset of X or a (possibly non-computable) function $X \rightarrow \mathbf{bool}$, where $\mathbf{bool} = \{\perp, \top\}$ is the set of truth values. Accordingly, since in realizability propositions are witnessed by realizers, a predicate ϕ on a per $A = (|A|, \approx_A)$ is a (possibly non-computable) function $\phi : \llbracket |A| \rrbracket \times \llbracket |\phi| \rrbracket \rightarrow \mathbf{bool}$ that is *strict* (if $\phi(u, v)$ then $u \in \llbracket |A| \rrbracket$) and *extensional* (if $\phi(u_1, v)$ and $u_1 \approx_A u_2$ then $\phi(u_2, v)$).

Suppose we have implemented the real numbers \mathbb{R} as a per $R = (\mathbf{real}, \approx_R)$, and consider $\forall a:R. \forall b:R. \exists x:R. x^3 + ax + b = 0$. By computing according to Figure 1, we see that a realizer for this proposition is a value r of type $\mathbf{real} \rightarrow \mathbf{real} \rightarrow \mathbf{real} \times \mathbf{unit}$ 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 $\mathbf{real} \rightarrow \mathbf{real} \rightarrow \mathbf{real}$ 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* function⁶ 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.

Some propositions, such as equality and negation, have “irrelevant” realizers free of computational content. Sometimes only a part of a realizer is computationally irrelevant. 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. On input, one can specify whether abstract predicates have computational content. On output, extracted realizers go through a *thinning* phase, which removes irrelevant realizers.

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 pers the appropriate notion is that of a *uniform* family. 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 $(\approx_{A_i})_{i \in \llbracket I \rrbracket}$ that is strict (if $u \approx_{A_i} v$ then $i \in \llbracket I \rrbracket$) and extensional (if $u \approx_{A_i} v$ and $i \approx_I j$ then $u \approx_{A_j} v$).

We can also form the *sum* $\Sigma_{i \in I} A_i$ or *product* $\Pi_{i \in I} A_i$ of a uniform family, allowing an interpretation of (extensional) dependent type theory.

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. 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.

The output of RZ consists of *module specifications*, module types plus assertions about their components. More specifically, a typical specification may contain value declarations, type declarations and definitions, module declarations, specification definitions, proposition declarations, and assertions. RZ only outputs assertions that are free of computational content, and do not require knowledge of constructive mathematics to be understood.

A special construct is the *obligation* `assume $x:\tau$, p in e` which means “in term e , let x be any element of $\llbracket \tau \rrbracket$ that satisfies p ”. An obligation is equivalent to a combination of Hilbert’s indefinite description operator and a local definition,

⁶ 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.

`let $x=(\varepsilon x:\tau.p)$ in e` , where $\varepsilon x:\tau.p$ means “any $x \in \llbracket \tau \rrbracket$ such that p ”. The alternative form `assume p in e` stands for `assume $_unit, p$ 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. The programmer must replace each obligation with a value satisfying the obligation. 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.

The simplest theory Θ is a list of *theory elements* `thy $\theta_1 \dots \theta_n$ 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.

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.

Propositions and predicates appearing in theories may use full first-order constructive logic, not just the negative fragment.

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 labeled 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.

The term language includes introduction and elimination constructs for the set level. For product sets we have tuples and projections $(\pi_1 e, \pi_2 e, \dots)$, 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

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 $\llbracket s \rrbracket$. Here \approx_s is *not* an OCaml value of type `s → s → bool`, but an abstract relation on the set $\llbracket s \rrbracket \times \llbracket s \rrbracket$. The relation may be uncomputable.

The translation of the declaration of a dependent set `Parameter t : s → Set` uses uniform families (Section 2.2). The underlying type `t` is non-dependent, but the per \approx_t receives an additional parameter `x : ||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`.

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. We comment on those that are least familiar.

Subsets. Given a predicate ϕ on a per A , the sub-per $\{x : A \mid \phi\}$ has underlying type $|A| \times |\phi|$ where $(u_1, v_1) \approx_{\{x:A \mid \phi\}} (u_2, v_2)$ 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 $\{x : A \mid \phi\}$ carries information about *why* the element belongs to the subset.

A type coercion $e : t$ can convert an element of the subset $s = \{x : t \mid \phi(x)\}$ 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 $(u, (\text{assure } v : |\phi|, \phi u v \text{ in } v))$.

Quotients. Even though we may form quotients of pers by arbitrary equivalence relations, only quotients by $\neg\neg$ -stable relations behave as expected.⁷ A stable equivalence relation on a per A is the same thing as a partial equivalence relation ρ on $|A|$ which satisfies $\rho(x, y) \implies x \approx_A y$. Then the quotient A/ρ is the per with $|A/\rho| = |A|$ and $x \approx_{A/\rho} y \iff \rho(x, y)$.

Luckily, it seems that many equivalence relations occurring in computable mathematics are stable, or can be made stable. For example, the coincidence relation on Cauchy sequences is expressed by a $\forall\exists\forall$ formula, but if we consider *rapid* Cauchy sequences (those sequences a satisfying $\forall i \in \mathbb{N}. |a_{i+1} - a_i| \leq 2^{-i}$), it becomes a (negative) \forall formula. 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 $\text{let } [x]_\rho = \xi \text{ in } e$, 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 $\mathbf{rz} A$, defined by $|\mathbf{rz} A| = |A|$ and $u \approx_{\mathbf{rz} A} v \iff \|A\| \wedge u = v$, where by $u = v$ we mean observational equality of values u and v . An element $r \in \mathbf{rz} A$ realizes a unique element $\mathbf{rz} r \in A$. The elimination term $\text{let } \mathbf{rz} x = e_1 \text{ in } e_2$, 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 $\mathbf{rz} A$ validates the Presentation Axiom (see Section 7.3). 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 [1], are explicitly expressed in terms of realizers.

5.2 Translation of propositions

The driving force behind the translation of logic is a theorem [17, 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 ”.

⁷ The trouble is that from equality of equivalence classes $[x]_\rho = [y]_\rho$ we may conclude only $\neg\neg\rho(x, y)$ rather than the expected $\rho(x, y)$.

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 :  $|\phi|$ 
assertion A : u  $\Vdash$   $\phi$ 
```

which requires the programmer to validate ϕ by providing a realizer for it. When ϕ is a compound statement RZ computes the meaning as described in Figure 1.

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** $p : s \rightarrow \mathbf{Prop}$ is translated to a type declaration **type** `ty_p` and a predicate declaration **predicate** $p : s \rightarrow \mathbf{ty_p} \rightarrow \mathbf{bool}$ together with assertions that p is strict and extensional.

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 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.

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. An *optimization* pass applies an ad-hoc collection of basic logical and term simplifications in order to make the output more readable. Some redundancy may remain, but in practice the optimization pass helps significantly.

Finally, the user can specify two optional steps occur. RZ can perform a *phase-splitting* pass [18]. This is an experimental implementation of an transformation that can replace a functor (a relatively heavyweight language construct) by parameterized types and/or polymorphic values.

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.

7 Examples

In this section we look at several examples which demonstrate various points of RZ. Unfortunately, serious examples from computable mathematics take too

much space⁸ and will have to be presented separately. 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, but RZ requires an axiom

```
Parameter s : Set.
Axiom eq:  $\forall x y : s, x = y \vee \neg(x = y)$ .
```

to produce a realizer for equality

```
val eq : s → s → ['or0 | 'or1]
assertion eq :  $\forall (x:|s|, y:|s|), (\text{match eq } x \text{ y with}$ 
                                      $\text{'or0} \Rightarrow x \approx_s y$ 
                                      $| \text{'or1} \Rightarrow \neg(x \approx_s y) )$ 
```

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 $x \approx_s y$, and if it returns `'or1`, then $\neg(x \approx_s y)$. In other words `eq` is a decision procedure.

7.2 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 [19]. 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 labeled by the elements of $T(x)$. Figure 2 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 $\prod_{x \in B.s} (B.t(x) \rightarrow w) \rightarrow w$. 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 dependencies at the level of types and terms disappear. A branching type is determined by a pair of non-dependent types `s` and `t` but the per \approx_t depends on $\llbracket s \rrbracket$. 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 $B.s \rightarrow (B.t \rightarrow w) \rightarrow w$. One can use phase-splitting to translate axiom **induction** into a specification of a polymorphic function

$$\text{induction} : (B.s \rightarrow (B.t \rightarrow w) \rightarrow (B.t \rightarrow \alpha) \rightarrow \alpha) \rightarrow w \rightarrow \alpha,$$

⁸ The most basic structure in analysis (the real numbers) alone requires several operations and a dozen or more axioms.

```

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:
     $\forall M : \text{thy}$  Parameter p : w -> Prop. end,
    ( $\forall x : B.s, \forall f : B.t\ x \rightarrow w,$ 
      ( $\forall y : B.t\ x, M.p\ (f\ y)) \rightarrow M.p\ (tree\ x\ f)$ ) ->
     $\forall t : w, M.p\ t.$ 
end.

```

Fig. 2. General inductive types

which is a form of recursion on well-founded trees. Instead of explaining induction, we show a surprisingly simple, hand-written 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

```

7.3 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: ( $\forall x : a, \exists y : b, r\ x\ y$ ) ->
          ( $\exists c : a \rightarrow b, \forall x : a, r\ x\ (c\ x)$ ).

```

The relevant part of the output is

```

val ac : (a -> b * ty_r) -> (a -> b) * (a -> ty_r)
assertion ac :
   $\forall f : a \rightarrow b * ty_r,$ 
  ( $\forall (x : \|a\|), \text{let } (p, q) = f\ x \text{ in } p : \|b\| \wedge r\ x\ p\ q$ ) ->
  let (g, h) = ac f in
  g :  $\|a \rightarrow b\| \wedge (\forall (x : \|a\|), r\ x\ (g\ x)\ (h\ x))$ 

```

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 $g:\|a \rightarrow b\|$ says that g must be extensional, i.e., map equivalent realizers to equivalent realizers. We could define h as the first component of f , but we cannot hope to implement g in general because the second component of f is not assumed to be extensional.

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

$$\text{Axiom iac: } (\forall x : a, \exists y : b, r\ x\ y) \rightarrow (\exists c : rz\ a \rightarrow b, \forall x : rz\ a, r\ (rz\ x)\ (c\ x)).$$

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 was weakened to $\forall x:a. (x:\|a\| \rightarrow g\ x:\|b\|)$. 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* [20], which states that every set is an image of one satisfying the axiom of choice.

7.4 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 $f : (\text{nat} \rightarrow \text{nat}) \rightarrow \text{nat}$. 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:

$$\text{Axiom continuity: } \forall f : (\text{nat} \rightarrow \text{nat}) \rightarrow \text{nat}, \forall a : \text{nat} \rightarrow \text{nat}, \\ \exists k, \forall b : \text{nat} \rightarrow \text{nat}, (\forall m, m \leq k \rightarrow a\ m = b\ m) \rightarrow f\ a = f\ b.$$

The axiom says that for any f and a there exists $k \in \text{nat}$ 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,⁹ 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'` which 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, states and proves theorems that the code behaves correctly, and has Coq extract correct code. In such cases RZ is complementary to Coq; it can suggest the appropriate division between code and theorems. We hope RZ will soon be able to produce output in Coq's input syntax.

Komagata and Schmidt [8] 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. 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., [22]).

8.2 Other Models of Computability

Many formulations of computable mathematics are based on realizability models [14], even though they were not initially developed, (nor are they usually presented) within the framework of realizability: Recursive Mathematics [23] is based on the original realizability by Turing machines [24]; Type Two Effectivity [1] on function realizability [25] and relative function realizability [26], while topological and domain representations [27, 28] are based on realizability over the

⁹ 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 [21, 9.6.10].

graph model $\mathcal{P}\omega$ [29]. A common feature is that they use models of computation which are well suited for the theoretical studies of computability.

Approaches based on simple programming languages with datatypes for real numbers [30,31] and topological algebras [2], and machines augmented with (suitably chosen subsets of) real numbers [32–34] are motivated by issues ranging from theoretical concerns about computability/complexity to practical questions in computational geometry. RZ attempts to improve practicality by using a real-world language, and by providing an input language rich enough for descriptions of mathematical structures going 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 [35].

References

1. Weihrauch, K.: *Computable Analysis*. Springer, Berlin (2000)
2. Tucker, J., Zucker, J.I.: Computable functions and semicomputable sets on many-sorted algebras. In Abramsky, S., Gabbay, D., Maibaum, T., eds.: *Handbook of Logic in Computer Science, Volume 5*, Oxford, Clarendon Press (1998)
3. Blanck, J.: Domain representability of metric spaces. *Annals of Pure and Applied Logic* **83** (1997) 225–247
4. Edalat, A., Lieutier, A.: Domain of differentiable functions. In Blanck, J., Brattka, V., Hertling, P., Weihrauch, K., eds.: *Computability and Complexity in Analysis*. (2000) CCA2000 Workshop, Swansea, Wales, September 17–19, 2000.
5. Müller, N.: The iRRAM: Exact arithmetic in C++. In Blanck, J., Brattka, V., Hertling, P., Weihrauch, K., eds.: *Computability and Complexity in Analysis*. (2000) 319–350 CCA2000 Workshop, Swansea, Wales, September 17–19, 2000.
6. Lambov, B.: RealLib: an efficient implementation of exact real arithmetic. In Grubba, T., Hertling, P., Tsuiki, H., Weihrauch, K., eds.: *Computability and Complexity in Analysis*. (2005) 169–175 *Proceedings, Second International Conference, CCA 2005*, Kyoto, Japan, August 25–29, 2005.
7. Leroy, X., Doligez, D., Garrigue, J., Rémy, D., Vouillon, J.: *The Objective Caml system, documentation and user’s manual - release 3.08*. Technical report, INRIA (July 2004)
8. Komagata, Y., Schmidt, D.A.: *Implementation of intuitionistic type theory and realizability theory*. Technical Report TR-CS-95-4, Kansas State University (1995)
9. Bertot, Y., Castéran, P.: *Interactive Theorem Proving and Program Development*. Springer (2004)
10. Benl, H., Berger, U., Schwichtenberg, H., Seisenberger, M., Zuber, W.: Proof theory at work: Program development in the Minlog system. In Bibel, W., Schmidt, P.H., eds.: *Automated Deduction: A Basis for Applications. Volume II, Systems and Implementation Techniques*. Kluwer Academic Publishers, Dordrecht (1998)
11. Bauer, A., Stone, C.A.: Specifications via realizability. In: *Proceedings of the Workshop on the Constructive Logic for Automated Software Engineering (CLASE 2005)*. Volume 153 of *Electronic Notes in Theoretical Computer Science*. (2006) 77–92
12. Longley, J.: Matching typed and untyped realizability. *Electr. Notes Theor. Comput. Sci.* **23**(1) (1999)
13. Longley, J.: When is a functional program not a functional program? In: *International Conference on Functional Programming*. (1999) 1–7

14. Bauer, A.: The Realizability Approach to Computable Analysis and Topology. PhD thesis, Carnegie Mellon University (2000)
15. Jacobs, B.: Categorical Logic and Type Theory. Elsevier Science (1999)
16. Post, E.: Recursive unsolvability of a problem of Thue. *The Journal of Symbolic Logic* **12** (1947) 1–11
17. Troelstra, A.S., van Dalen, D.: Constructivism in Mathematics, An Introduction, Vol. 1. Number 121 in *Studies in Logic and the Foundations of Mathematics*. North-Holland (1988)
18. Harper, R., Mitchell, J.C., Moggi, E.: Higher-order Modules and the Phase Distinction. In: *Proceedings of the 17th ACM Symposium on Principles of Programming Languages (POPL '90)*. (1990) 341–354
19. Nordström, B., Petersson, K., Smith, J.M.: *Programming in Martin-Löf's Type Theory*. Oxford University Press (1990)
20. Barwise, J.: *Admissible Sets and Structures*. Springer-Verlag (1975)
21. Troelstra, A.S., van Dalen, D.: Constructivism in Mathematics, An Introduction, Vol. 2. Number 123 in *Studies in Logic and the Foundations of Mathematics*. North-Holland (1988)
22. Kahrs, S., Sannella, D., Tarlecki, A.: The definition of Extended ML: A gentle introduction. *Theoretical Computer Science* **173**(2) (1997) 445–484
23. Ershov, Y.L., Goncharov, S.S., Nerode, A., Remmel, J.B., eds.: *Handbook of Recursive Mathematics*. Elsevier, Amsterdam (1998)
24. Kleene, S.C.: On the interpretation of intuitionistic number theory. *Journal of Symbolic Logic* **10** (1945) 109–124
25. Kleene, S.C., Vesley, R.E.: *The Foundations of Intuitionistic Mathematics, especially in relation to recursive functions*. North-Holland Publishing Company (1965)
26. Birkedal, L.: *Developing Theories of Types and Computability*. PhD thesis, School of Computer Science, Carnegie Mellon University (December 1999)
27. Blanck, J.: *Computability on topological spaces by effective domain representations*. PhD thesis, Uppsala University, Department of Mathematics, Uppsala, Sweden (1997)
28. Bauer, A., Birkedal, L., Scott, D.S.: Equilogical spaces. *Theoretical Computer Science* **1**(315) (2004) 35–59
29. Scott, D.S.: Data types as lattices. *SIAM Journal of Computing* **5**(3) (1976) 522–587
30. Escardó, M.H.: *PCF extended with real numbers*. PhD thesis, Department of Computer Science, University of Edinburgh (December 1997)
31. Marcial-Romero, J.R., Escardó, M.H.: Semantics of a sequential language for exact real-number computation. In: *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science*. (July 2004) 426–435
32. Borodin, A., Monro, J.I.: The computational complexity of algebraic and numeric problems. Number 1 in *Elsevier computer science library : Theory of computation series*. New York, London, Amsterdam : American Elsevier (1975)
33. Blum, L., Cucker, F., Shub, M., Smale, S.: *Complexity and Real Computation*. Springer-Verlag, New York (1998)
34. Yap, C.K.: *Theory of real computation according to EGC* (2006) To appear in LNCS Volume based on the Dagstuhl Seminar “Reliable Implementation of Real Number Algorithms: Theory and Practice”, Jan 8-13, 2006.
35. Bishop, E., Bridges, D.: *Constructive Analysis*. Volume 279 of *Grundlehren der math. Wissenschaften*. Springer-Verlag (1985)