

GMETA: A Generic Formal Metatheory Framework for First-Order Representations

Bruno C. d. S. Oliveira
Seoul National University
bruno@ropas.snu.ac.kr

Gyesik Lee
Seoul National University
gslee@ropas.snu.ac.kr

Sungkeun Cho
Seoul National University
skcho@roaps.snu.ac.kr

Kwangkeun Yi
Seoul National University
kwang@ropas.snu.ac.kr

April 12, 2011

Abstract

This paper presents GMETA: a generic framework for *first-order representations* of variable binding that provides *once and for all* many of the so-called infrastructure lemmas and definitions required in mechanizations of formal metatheory. The key idea is to employ *datatype-generic programming* (DGP) and *modular programming* techniques to deal with the infrastructure overhead. Using a generic *universe* for representing a large family of object languages we define datatype-generic libraries of infrastructure for first-order representations such as *locally nameless* or *de Bruijn* indices. Modules are then used to provide *templates*: a convenient interface between the datatype-generic libraries and the end-users of GMETA. We conducted case studies based on the POPLmark challenge, and showed that dealing with challenging binding constructs, like the ones found in System $F_{<}$, is possible with GMETA. All of GMETA's generic infrastructure is implemented in the Coq theorem prover. Furthermore, due to GMETA's modular design, the libraries can be easily used, extended, and customized by users.

1 Introduction

A key issue in mechanical developments of formal metatheory for programming languages concerns the representation and manipulation of terms with variable binding. There are two main approaches to address this issue: *first-order* and *higher-order* approaches. In first-order approaches variables are typically encoded using names or natural numbers, whereas higher-order approaches such as higher-order abstract syntax (HOAS) use the function space in the meta-language to encode binding of the object language. Higher-order approaches are appealing because issues like capture-avoidance and alpha-equivalence can be handled once and for all by the meta-logic. This is why such approaches are used in logical frameworks such as Abella (Gacek 2008), Hybrid (Momigliano et al. 2008) or Twelf (Pfenning and Schürmann 1999); and have also been advocated as an interesting alternative (Despeyroux et al. 1995; Chlipala 2008) for formalizing metatheory in general-purpose theorem provers like Coq (Coq Development Team 2009).

The main advantage of first-order approaches, and the reason why they are so popular in theorem provers like Coq, is that they are close to pen-and-paper developments (based on the nominal approach) and terms with binders are easy to manipulate.

However, the main drawback of first-order approaches is that the tedious infrastructure required for handling variable binding has to be repeated each time for a new object language. For each binding construct in the language, there is typically a set of *infrastructure operations*

		trm	typ
trm	Variables	$bsubst_{trm \times trm}$	$bsubst_{trm \times typ}$
	Parameters	$fst_{trm \times trm}$	$fst_{trm \times typ}$
typ	Variables	$bsubst_{typ \times trm}$	$bsubst_{typ \times typ}$
	Parameters	$fst_{typ \times trm}$	$fst_{typ \times typ}$

Figure 1: Possible variations of substitutions for parameters and variables for a language with two syntactic sorts (`trm` and `typ`) in the locally nameless style.

and associated *lemmas* that should be implemented. In the locally nameless style (Aydemir et al. 2008) and locally named (Mckinna and Pollack 1993) styles, for example, we usually need operations like substitutions for parameters (free variables) and for (bound) variables as well some associated lemmas. For de Bruijn indices (de Bruijn 1972) we need similar infrastructure, but for operations such as substitution and shifting instead.

The problem is that, in many cases, the majority of the total number of lemmas and definitions in a formalization consists of basic infrastructure. Figure 1 illustrates the issue using a simple language with two syntactic sorts (types and terms) supporting binding constructs for both type and term variables and assuming a locally nameless style. In the worst case scenario, 8 different types of substitutions will be needed. Basically we need substitutions for parameters and variables, and for each of these we need to consider all four combinations of substitutions using types and terms. While not all operations are usually needed in formalizations, many of them are. For example, System $F_{<}$, which is the language formalized in the POPLMark challenge (Aydemir et al. 2005), requires 6 out of the 8 substitutions. Because for each operation we need to also prove a number of associated lemmas, solutions to the POPLMark challenge typically have a large percentage of lemmas and definitions just for infrastructure. In the solution by Aydemir et al. (2008), infrastructure amounts to 65% of the total number of definitions and lemmas (see also Figure 10).

Importantly, it is insufficient to consider only *homogeneous* operations like $bsubst_{trm \times trm}$, which perform substitutions of variables on terms of the same sort (`trm`). In the general case we also have to consider *heterogeneous* operations, such as $bsubst_{typ \times trm}$, in which the sort of variables being substituted (`typ`) is not of the same sort as the terms which are being substituted (`trm`). In languages of the System F family, operations like $bsubst_{typ \times trm}$ and $fst_{typ \times trm}$ are needed because those languages support type abstractions in terms ($\Lambda X.e$) and, consequently, we need to provide a way to substitute type variables in terms.

As the number of syntactic sorts and binding constructs increases, there is a combinatorial explosion of the number of infrastructure lemmas and operations. Indeed, reports from larger formalizations confirm that the problem only gets worse with size: in a type-directed translation from an ML-module language to System F_{ω} using locally nameless style (Rossberg et al. 2010) the authors report that “*Out of a total of around 550 lemmas, approximately 400 were tedious infrastructure lemmas*”.

1.1 Our solution

To deal with the combinatorial explosion of infrastructure operations and lemmas we propose the use of *datatype-generic programming* (DGP) and *modular programming* techniques. The key idea is that, with DGP, we can define once and for all the tedious infrastructure lemmas and operations in a generic way and, with modules, we can provide a convenient interface for users to instantiate such generic infrastructure to their object languages.

Our realization of this idea is GMETA: a generic framework for first-order representations of variable binding, implemented in the Coq theorem prover¹. In GMETA a DGP *uni-*

¹We also have an experimental Agda implementation.

```

(*@Iso typ_iso {
  Parameter typ_fvar,
  Variable  typ_bvar,
  Binder    typ_all _
}*)
Inductive typ :=
| typ_fvar  :  $\mathbb{N} \rightarrow typ$ 
| typ_bvar  :  $\mathbb{N} \rightarrow typ$ 
| typ_top   :  $typ$ 
| typ_arrow :  $typ \rightarrow typ \rightarrow typ$ 
| typ_all   :  $typ \rightarrow typ \rightarrow typ$ .
(*@Iso trm_iso {
  Parameter trm_fvar,
  Variable  trm_bvar,
  Binder    trm_abs _,
  Binder    trm_tabs _ binds typ
}*)
Inductive trm :=
| trm_fvar  :  $\mathbb{N} \rightarrow trm$ 
| trm_bvar  :  $\mathbb{N} \rightarrow trm$ 
| trm_app   :  $trm \rightarrow trm \rightarrow trm$ 
| trm_abs   :  $typ \rightarrow trm \rightarrow trm$ 
| trm_tapp  :  $trm \rightarrow typ \rightarrow trm$ 
| trm_tabs  :  $typ \rightarrow trm \rightarrow trm$ .

```

Figure 2: Syntax definitions and GMETA isomorphism annotations for a locally nameless style version of System $F_{<}$ in Coq.

verse (Martin-Löf 1984) is used to represent a large family of object languages and includes constructs for representing the binding structure of those languages. The universe itself is independent of the particular choice of first-order representations: it can be instantiated, for example, to *locally nameless* or *de Bruijn* representations. GMeta uses that universe to provide libraries with the infrastructure operations and lemmas for the various first-order representations.

The infrastructure is reused by users through *templates*. Templates are functors parameterized by isomorphisms between the object language and the corresponding representation of that language in the universe. By instantiating templates with isomorphisms, users get access to a module that provides infrastructure tailored for a particular binding construct in their own object language. For example, for System $F_{<}$, the required infrastructure is provided by 3 modules which instantiate GMeta's locally nameless template:

Module $M_{\text{trm} \times \text{trm}}$:= *LNTemplate trm_iso trm_iso*.

Module $M_{\text{typ} \times \text{typ}}$:= *LNTemplate typ_iso typ_iso*.

Module $M_{\text{typ} \times \text{trm}}$:= *LNTemplate typ_iso trm_iso*.

Essentially, each module corresponds to one of the 3 combinations needed in System $F_{<}$, and contains the relevant lemmas and operations. By using this scheme we can deal with the general case of object languages with N syntactic sorts, just by expressing the combinations needed in that language. Moreover GMeta can also provide some more specialized templates for additional reuse and it is easy for users to define their own types of infrastructure and customized templates.

Since isomorphisms can be mechanically generated from the inductive definition of the object language, provided a few annotations, GMeta also includes optional tool support for generating such isomorphisms automatically. Figure 2 illustrates these annotations for System $F_{<}$. Essentially, the keyword **Is** introduces an isomorphism annotation, while the keywords **Parameter**, **Variable** and **Binder** provide the generator with information about which constructors correspond, respectively, to the parameters, variables or binders. Therefore, at the cost of just a few annotations or explicitly creating an isomorphism by hand, GMeta provides much of the tedious infrastructure boilerplate that would constitute a large part of the whole development otherwise.

1.2 Contributions

Our main contribution is to investigate how DGP techniques can deal with the infrastructure overhead required by formalizations using first-order representations. GMeta shows that DGP is a viable approach to address the problem and that it has some important advantages over alternatives like LGen (Aydemir and Weirich 2009) (see Section 7). Nevertheless, GMeta is still a proof-of-concept tool and, as discussed in Section 6.4, a bit more support from theorem provers is desirable for improving the usability of the tool.

More concretely, the contributions of this paper are:

- *Sound, generic, reusable and extensible infrastructure for first-order representations*: The main advantages of using DGP are that it allows a library-based approach in which 1) the infrastructure can be defined and verified once and for all *within* the meta-logic itself; and 2) extending the infrastructure is easy since it just amounts to extending the library. Approaches based on code generation can only provide reuse, but they do not guarantee soundness and they are hard to extend.
- *Heterogeneous generic operations and lemmas*: Of particular interest is the ability of GMeta to deal with challenging binding constructs, involving multiple syntactic sorts (such as binders found in the System F family of languages), using heterogeneous generic operations and lemmas.

- *Case studies using the POPLmark challenge:* To validate our approach in practice, we conducted case studies using the POPLmark challenge. Compared to other solutions, our approach shows significant savings in the number of definitions and lemmas required by formalizations.
- *Coq implementation and other resources:* The GMETA framework Coq implementation is available online² along with other resources such as tutorials and more case studies.

2 Different Types of Infrastructure in GMETA

The purpose of this section is to define what should be considered infrastructure lemmas and definitions and to distinguish between the different varieties of these lemmas and definitions. A classification of infrastructure is desirable for two reasons:

1. Some infrastructure is easier to reuse than other in GMETA. Also some infrastructure is more general than other.
2. There is an informal understanding of what infrastructure is, but the existing literature is not very clear as to what actually constitutes infrastructure (Aydemir et al. 2008; Aydemir and Weirich 2009).

To help with the characterization of infrastructure, we propose two classifications, by difficulty to reuse and by generality of the infrastructure.

2.1 Classification by difficulty to reuse

Type 1. Common operations and lemmas: Operations such as free and bound variables, term size or substitution-like operations (such as substitutions for parameters and variables in the locally nameless style; or shifting in the de Bruijn style) are of this type. For example:

$$\begin{aligned} fsubst_{\text{trm} \times \text{trm}} &: \mathbb{N} \rightarrow \text{trm} \rightarrow \text{trm} \rightarrow \text{trm} \\ bsubst_{\text{trm} \times \text{trm}} &: \mathbb{N} \rightarrow \text{trm} \rightarrow \text{trm} \rightarrow \text{trm} \\ bsubst_{\text{typ} \times \text{trm}} &: \mathbb{N} \rightarrow \text{typ} \rightarrow \text{trm} \rightarrow \text{trm} \end{aligned}$$

Also, basic lemmas about the operations (such as several forms of permutation lemmas about substitutions, or lemmas involving some simple conditions about variables) fall in this category. For example:

$$\begin{aligned} tfsbst_lemma &: \forall (T U V : \text{trm}) (a b : \mathbb{N}), \\ & a \neq b \Rightarrow \\ & a \notin (fv_{\text{trm} \times \text{trm}} V) \Rightarrow \\ & fsubst_{\text{trm} \times \text{trm}} (fsubst_{\text{trm} \times \text{trm}} T a U) b V = \\ & fsubst_{\text{trm} \times \text{trm}} (fsubst_{\text{trm} \times \text{trm}} T b V) a (fsubst_{\text{trm} \times \text{trm}} U b V) \end{aligned}$$

Type 2. Lemmas involving unary inductive relations (predicates): When formalizing metatheory it is common to use inductive relations. For example several different forms of well-formedness are usually described in this way. For unary inductive relations, which are essentially predicates, we can use the predicate form instead to provide the associated lemmas in GMETA. An example is the closed terms predicate:

$$closed_{\text{trm} \times \text{trm}} : \text{trm} \rightarrow \text{Prop}$$

²<http://ropas.snu.ac.kr/gmeta/>

There are many permutations that only hold for closed terms. For example:

$$\begin{aligned} & \text{tbsubst_var_twice_wf} : \forall (T : \text{trm}) k (U V : \text{trm}), \\ & \text{closed}_{\text{trm} \times \text{trm}} V \Rightarrow \\ & \text{bsubst}_{\text{trm} \times \text{trm}} T k V = \text{bsubst}_{\text{trm} \times \text{trm}} (\text{bsubst}_{\text{trm} \times \text{trm}} T k V) k U \end{aligned}$$

Type 3. Lemmas involving general inductive relations: General inductive relations can be dealt in GMETA with a corresponding generic inductive relation. However, there is a cost in using these generic inductive relations because they require a mapping between the concrete relation defined directly over the object language and the generic relation. This mapping requires additional work from the user as well as some basic knowledge about DGP, so they are not as easy to reuse as type 1 and 2. The typical example is well-formed terms in an environment:

$$\text{envT}_{\text{typ} \times \text{trm}} : (\text{env typ}) \rightarrow \text{trm} \rightarrow \text{Prop}$$

For which there are many interesting lemmas that are useful in formalizations. For example:

$$\begin{aligned} & \text{envT_Twf} : \forall (E : \text{env typ}) (T : \text{trm}), \\ & \text{envT}_{\text{typ} \times \text{trm}} E T \Rightarrow \\ & \forall (k : \mathbb{N}) (U : \text{trm}), T = \text{bsubst}_{\text{trm} \times \text{trm}} T k U \end{aligned}$$

In short, it is easy to reuse infrastructure of types 1 and 2 in GMETA, but reusing infrastructure of type 3 requires a bit more work.

2.2 Classification by generality

Some infrastructure is general whereas other is domain-specific (that is, only exist for particular types of languages). We propose the following classification to distinguish between these two types of infrastructure:

Type A. General infrastructure Lemmas and definitions involving only one binder and a single combination of sorts are pervasive: they are definable for any object language with some binding structure. The examples that we gave for infrastructure of type 1 and 2 in Section 2.1 are of this type.

Type B. Domain-specific infrastructure Lemmas and definitions involving multiple binders or multiple combinations of sorts tend to be specific to particular languages. For example the lemma envT_Twf uses two different combinations ($\text{trm} \times \text{trm}$ and $\text{typ} \times \text{trm}$). This lemma makes sense when we are talking about some typed languages, but it does not make sense for many other languages. Another example is:

$$\begin{aligned} & \text{fbfst_subst_perm_core}_{\text{typ} \times \text{trm}} : \forall (t : \text{trm}) (u, v : \text{typ}) (m k : \mathbb{N}), \\ & \text{closed}_{\text{typ} \times \text{typ}} u \Rightarrow \\ & \text{bst}_{\text{typ} \times \text{trm}} k (\text{fst}_{\text{typ} \times \text{typ}} m u v) (\text{fst}_{\text{typ} \times \text{trm}} m u t) \\ & = \text{fst}_{\text{typ} \times \text{trm}} m u (\text{bst}_{\text{typ} \times \text{trm}} k v t) \end{aligned}$$

This lemma involves substitutions on different kind of binders ($\text{fst}_{\text{typ} \times \text{typ}}$ and $\text{fst}_{\text{typ} \times \text{trm}}$) and it makes sense when the language has 2 different binders on 2 sorts.

The 2 classifications proposed by us are mostly orthogonal. We will use a combination between the letter A or B and the numbers 1, 2 and 3 to refer to a particular type on which a lemma or definition falls into. For example, fst_subst_lemma is of type A1 (that is, it is a general common lemma), whereas $\text{fbfst_subst_perm_core}_{\text{typ} \times \text{trm}}$ is of type B2 (that is, it is a domain-specific

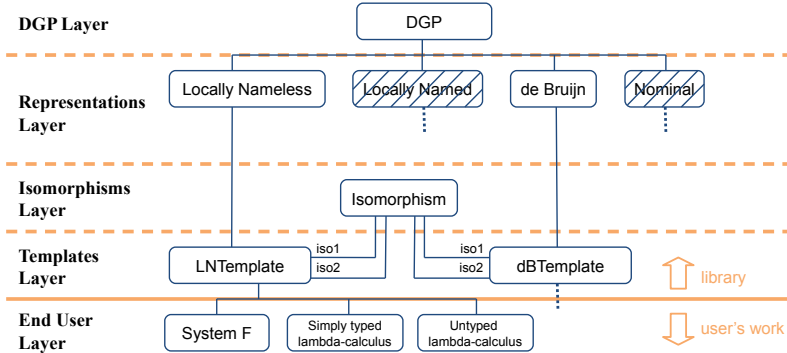


Figure 3: A simplified modular structure overview of GMETA.

lemma using a predicate). As a remark, we noticed that in practice lemmas that fall in the category A3 are rare and we do not have any examples of these in GMETA.

3 GMeta Design

This section gives a general overview of GMETA’s design and discusses the techniques used by us to make GMETA convenient to use. The GMETA framework is structured into 5 layers. In this paper we will discuss two layers in detail in Sections 4 and 5. These layers are related to DGP and are the most interesting from a technical point of view. The other layers will not be introduced in detail due to the lack of space, although the key ideas are discussed in this section. These other layers are important because of the convenience aspect of GMETA, but they are less interesting from a technical point of view. More information about these layers is available in GMETA’s webpage.

3.1 GMeta’s modular structure

Figure 3 shows an overview of the modular structure of GMETA. The structure is hierarchical, with the more general modules at the top and the more specific modules at the bottom.

- **DGP Layer:** The core DGP infrastructure is defined at the top-most layer. The main component is a universe that acts as a generic language that the lower-level modules use to define the infrastructure lemmas and definitions.
- **Representation Layer:** This layer is where the generic infrastructure lemmas and definitions for particular first-order representations are defined. GMETA currently supports locally nameless and de Bruijn representations, and we are planning to support more representations in the future.
- **Isomorphism Layer:** This layer provides simple module signatures for isomorphisms that serve as interfaces between the object language and its representation in the generic language. The adequacy of the object language representation follows from the isomorphism laws.
- **Templates Layer:** This layer provides templates for the basic infrastructure lemmas and definitions required by particular meta-theoretical developments. Templates are ML-style functors parameterized by isomorphisms between the syntactic sorts of object language and their corresponding representations in the generic language. In the Figure 3 we show

only *LNTemplate* and *dBTemplate*, which are the fundamental templates providing reuse for the general infrastructure (type A). In GMETA there are also other templates which can be *optionally* used for reusing additional infrastructure. However, those templates tend to be more domain-specific (that is they cover type B infrastructure).

- **End User Layer:** End users will use GMETA’s libraries to develop metatheory for particular object languages such as for example, the simply typed lambda calculus (STLC) or System $F_{<}$. used in our case studies.

3.2 Making GMETA Convenient to Use

The main design goal of GMETA is to make it useful for users wishing to conduct formalizations of metatheory in Coq using conventional first-order representations, without significantly increasing the *cost-of-entry* to users. Importantly, although DGP plays a fundamental role in the definition of the core libraries of GMETA (at the DGP and representation layers), end users should not need knowledge about DGP for basic uses of GMETA. However this is not trivial to achieve because, among other things, end-user proofs (such as soundness) generally require unfolding infrastructure operations like substitution and those operations are written in a datatype-generic way, in a form which is alien to users that do not know about DGP. Therefore, to provide convenience to the user and to avoid that users need to know about DGP, GMETA employs the techniques described next.

Automatically generated isomorphisms: Isomorphisms are a well-know technique to allow using DGP with conventional user-defined datatypes (or inductive definitions) such as the ones presented in Figure 2. Like in most libraries for DGP in functional programming languages (Rodriguez et al. 2009), GMETA uses automatically generated isomorphisms between the user-defined object language and a corresponding representation of that language the generic universe. However, unlike the isomorphisms used in DGP libraries in functional programming languages, GMETA’s isomorphisms cannot be generated automatically without additional information about the binding structure of the language. To deal with this problem GMETA uses a small annotation language to describe the binding structure of the object language. This allows the generation of isomorphisms which are aware of the correct binding structure for that language (see Figure 2 for an example of the annotation language).

Templates: GMETA uses templates to solve the problem of interfacing with the infrastructure DGP libraries. The definitions and lemmas in templates are very simple. For example, the definition of parameter substitution in the locally nameless template is:

Module *LNTemplate* (*iso*_{S₁} : **Iso**, *iso*_{S₂} : **Iso**).

...

$$[\cdot \rightarrow \cdot]_T \cdot : \mathbb{N} \rightarrow S_1 \rightarrow S_2 \rightarrow S_2$$

$$[k \rightarrow u]_T t = to_{S_2} ([k \rightarrow (from_{S_1} u)] (from_{S_2} t))$$

Essentially, S_1 and S_2 are supposed to be the types of the syntactic sorts used in object language. These types come from the isomorphisms iso_{S_1} and iso_{S_2} , which are the parameters of *LNTemplate*. Definitions like $[\cdot \rightarrow \cdot]_T \cdot$ are simply using the isomorphism (through the operations to_{S_2} , $from_{S_1}$ and $from_{S_2}$) to interface with generic operations like $[\cdot \rightarrow \cdot] \cdot$ (see Figure 8) defined in the representations layer.

Because of the isomorphisms between the user’s object language and the representation of that language in the universe, users do not need to interact directly with the generic universe. Instead all that a user needs to do is to instantiate the templates with the automatically generated isomorphisms. In Section 1.1, we already described how this technique is used to generate the infrastructure for System $F_{<}$.

$$\begin{aligned} \text{tos}_2 (\text{from}_{S_2} t) &= t & (1) \\ \text{from}_{S_2} (\text{tos}_2 t) &= t & (2) \\ \text{from}_{S_2} ([k \rightarrow u]_T t) &= [k \rightarrow (\text{from}_{S_2} u)] (\text{from}_{S_2} t) & (3) \end{aligned}$$

Figure 4: Isomorphism and adequacy laws.

Special tactics: Templates alone are not enough to avoid dealing with DGP concepts directly. When proving lemmas for their own formalizations users may need to unfold operations, which are defined in terms of corresponding generic operations. For example, the following lemma – which states that typing is preserved by substitution – is a core lemma in formalization of in the solution to the POPLMark challenge by Aydemir et al. (2008).

Lemma *typing_subst* : $\forall E F U t T z u,$
 $(E \text{ ++ } (z, U) :: F) \vdash t : T \Rightarrow F \vdash u : U \Rightarrow$
 $(E \text{ ++ } F) \vdash ([z \rightarrow u]_T t) : T.$

Proof.

intros; dependent induction H; g_{simpl}.

...
g_{rewrite} tb_{fsubst}_permutation_var_wf; eauto.

...
Qed.

The details of the Coq proof are not relevant. What is important to note is: 1) the key difference to the original proof by Aydemir et al. (2008) is that two different tactics (*g_{simpl}* and *g_{rewrite}*) are used; and 2) the lemma *tb_{fsubst}_permutation_var_wf* and the operation $[\cdot \rightarrow \cdot]_T$ are provided by GMETA’s templates.

If the user would try to use *simpl* (the standard Coq tactic to unfold and simplify definitions) directly, the definition of $[\cdot \rightarrow \cdot]_T \cdot$ would be unfolded and he would be presented with parts of the definition of $[\cdot \rightarrow \cdot]$ (see Figure 8). This is a generic function defined in terms of GMeta’s universe. However this is clearly undesirable, since the expected definition at this point is one similar to a manually defined operation for the object language in hand.

Our solution to this problem is define some Coq tactics (such as *g_{simpl}* and *g_{rewrite}*) that *specialize* operations and lemmas such as $[\cdot \rightarrow \cdot]_T$ and *tb_{fsubst}_permutation_var_wf* using the isomorphisms provided by the user, and the isomorphism and adequacy laws shown in Figure 4.

Different use modes GMETA can be used in different use modes, depending on how much knowledge a user has and what the user wishes to accomplish with GMETA. The most common scenario of use of GMETA is to formalize metatheory using the *existing* libraries provided by GMETA. In this scenario users can use GMETA in two modes: basic or advanced. The distinction between this two modes essentially depends on which types of infrastructure can be dealt by GMETA :

- **Basic user:** A basic user can simply use the isomorphism generator, instantiate templates like *LNTemplate* and use the special tactics provided by GMETA to deal with infrastructure of type A1 and A2. It may be even possible to reuse some infrastructure of type B. If some domain-specific templates exist in the libraries for the particular kind of object-language that the user is defining, then that template can be used. There is no knowledge of DGP required and no need to understand the technical details involved in GMETA ’s libraries to use this mode.

- **Advanced user:** It is also possible to use GMETA to deal with infrastructure of type B provided some additional knowledge about DGP and understanding of how templates work.

Finally, GMETA can also be used by *library writers* who wish to develop their own libraries of infrastructure:

- **Library writer:** A library writer is someone that wishes to use GMETA to develop his own set of infrastructure (for example writing infrastructure for the locally named approach or for dealing with languages which support dynamic binding). The big advantage of GMETA in comparison with generative approaches such as LGen is that, provided some basic knowledge about DGP, developing libraries of infrastructure is easy.

4 DGP for Datatypes with First-Order Binders

This section briefly introduces DGP using inductive families to define universes of datatypes, and shows how to adapt a conventional universe of datatypes to support binders and variables. In our presentation we assume a type theory extended with inductive families, such as the Calculus of Inductive Constructions (CIC) (Paulin-Mohring 1996) or extensions of Martin-Löf type-theory (Martin-Löf 1984) with inductive families (Dybjer 1997).

4.1 Inductive Families

Functional languages like ML or Haskell support datatype declarations for simple inductive structures such as the natural numbers:

DATA Nat = z | s Nat

Inductive families are a generalization of conventional datatypes that has been introduced in dependently typed languages such as Epigram (McBride and McKinna 2004), Agda (Norell 2007) or the Coq theorem prover. They are also one of the inspirations for Generalized Algebraic Datatypes (GADTs) (Peyton Jones et al. 2006) in Haskell.

We adopt a notation similar to the one used by Epigram to describe inductive families. In general, such notation is of the form:

DATA *type-constructor-sig* WHERE *data-constructor-sigs*

Using this notation, the definition of natural numbers is:

$$\text{DATA } \frac{}{\text{Nat} : \star} \text{ WHERE } \frac{}{z : \text{Nat}} \quad \frac{n : \text{Nat}}{s \ n : \text{Nat}}$$

Essentially the same information as the conventional datatype definition is described (albeit in a more detailed form). Both the type and data constructors are written using *sans serif*. Signatures use natural deduction rules, with the context with all the constructor arguments and their types above the line, and the type of the fully applied constructor below the line. Finally, note that the notation \star (in $\text{Nat} : \star$) means the ‘type’ (or kind) of types.

The expressiveness of this notation reveals itself when we begin defining whole families of datatypes. For example we can define a family of vectors of size n as follows:

$$\text{DATA } \frac{A : \star \quad n : \text{Nat}}{\text{Vector}_A \ n : \star} \text{ WHERE } \frac{}{vz : \text{Vector}_A \ z} \quad \frac{n : \text{Nat} \quad a : A \quad as : \text{Vector}_A \ n}{vs \ a \ as : \text{Vector}_A \ (s \ n)}$$

$$\begin{array}{c}
\text{DATA } \text{Rep} = 1 \mid \text{Rep} + \text{Rep} \mid \text{Rep} \times \text{Rep} \mid \text{K Rep} \mid \text{R} \\
\\
\text{DATA } \frac{r, s : \text{Rep}}{\llbracket s \rrbracket_r : \star} \text{ WHERE} \quad \frac{}{() : \llbracket 1 \rrbracket_r} \quad \frac{s : \text{Rep} \quad v : \llbracket s \rrbracket}{k \ v : \llbracket \text{K } s \rrbracket_r} \\
\\
\frac{s_1, s_2 : \text{Rep} \quad v : \llbracket s_1 \rrbracket_r}{i_1 \ v : \llbracket s_1 + s_2 \rrbracket_r} \quad \frac{s_1, s_2 : \text{Rep} \quad v : \llbracket s_2 \rrbracket_r}{i_2 \ v : \llbracket s_1 + s_2 \rrbracket_r} \\
\\
\frac{s_1, s_2 : \text{Rep} \quad v_1 : \llbracket s_1 \rrbracket_r \quad v_2 : \llbracket s_2 \rrbracket_r}{(v_1, v_2) : \llbracket s_1 \times s_2 \rrbracket_r} \quad \frac{v : \llbracket r \rrbracket}{r \ v : \llbracket \text{R} \rrbracket_r} \\
\\
\text{DATA } \frac{s : \text{Rep}}{\llbracket s \rrbracket : \star} \text{ WHERE} \quad \frac{s : \text{Rep} \quad v : \llbracket s \rrbracket_s}{\text{in } v : \llbracket s \rrbracket}
\end{array}$$

Figure 5: A simple universe of types.

In this definition the type constructor for vectors has two type arguments. The first argument specifies the type A of elements of the vector, while the second argument n is the size of the vector. The type of the elements A is *parametric*; that is, it is a (globally visible) parameter of all the constructors (both type and data constructors). In contrast, the type n for the size of vectors varies for each constructor; it is z for the vz constructor and $s \ n$ for the vs constructor. We write parametric type arguments in type constructors such as Vector_A using subscript. Also, if a constructor is not explicitly applied to some arguments (for example $\text{vs } a \ as$ is not applied to n), then those arguments are implicitly passed.

4.2 Datatype Generic Programming

The key idea behind DGP is that many functions can be defined generically for whole families of datatype definitions. Inductive families are useful to DGP because they allow us to define universes (Martin-Löf 1984) representing whole families of datatypes. By defining functions over this universe we obtain generic functions that work for any datatypes representable in that universe.

A Simple Universe The universe that underlies GMETA can be viewed as a simplified version of the universe for regular tree types by Morris et al. (2004). Essentially, Morris et al.’s universe is expressive enough to represent recursive types using μ -types (Pierce 2002). However, instead of the nominal approach traditionally used with recursive type binders, the universe of regular tree types uses a well-scoped de Bruijn indices representation (Altenkirch and Reus 1999; McBride and McKinna 2004). As a consequence the presentation of that universe is somehow complicated by the use of telescopes needed in a well-scoped de Bruijn indices representation.

For presentation purposes and to avoid distractions related to the use of telescopes (which are orthogonal to our purposes), we will use instead a simplified version of regular tree types in which only a single top-level recursive type binder is allowed. This reduces the expressive power of the universe (in particular, no mutually inductive definitions are allowed), but it also avoids the use of telescopes and leads to a simpler presentation of the universe.

The complete universe of regular tree types by Morris et al. (which includes the ability to represent mutually-inductive definitions), is used in our experimental versions of GMETA (both in Coq and in Agda) available in the online web page.

Figure 5 shows our simple universe. The datatype `Rep` (defined using the simpler ML-style notation for datatypes) describes the “grammar” of types that can be used to construct the datatypes representable in the universe. The first three constructs represent unit, sum and product types. The `K` constructor allows the representation of constants of some representable type. The `R` constructor is the most interesting construct: it is a reference to the recursive type that we are defining. For example, the type representations for naturals and lists of naturals are defined as follows:

```
RNat : Rep
RNat = 1 + R
```

```
RList : Rep
RList = 1 + K RNat × R
```

Compared with the more familiar definitions of recursive types for naturals and lists using μ -types (Pierce 2002)

```
Nat =  $\mu R. 1 + R$ 
List =  $\mu R. 1 + Nat \times R$ 
```

we can see that the main difference is that there is no μ binder. This is because, as mentioned earlier, since we can only express a single-top level binders, a single variable (R) is enough and there is no need to have a separate binding construct.

The interpretation of the universe is given by two mutually inductive families $\llbracket \cdot \rrbracket_r$ and $\llbracket \cdot \rrbracket$, while the data constructors of these two families provide the syntax to build terms of that universe. The parametric type³ r in the subscript in $\llbracket \cdot \rrbracket_r$, is the recursive type that is used when interpreting the constructor `R`. For illustrating the data constructors of terms of the universe, we first define the constructors `nil` and `cons` for lists:

```
nil :  $\llbracket RList \rrbracket$ 
nil = in (i1 ())

cons :  $\llbracket RNat \rrbracket \rightarrow \llbracket RList \rrbracket \rightarrow \llbracket RList \rrbracket$ 
cons n ns = in (i2 (k n, r ns))
```

When interpreting $\llbracket RList \rrbracket$, the representation type r in $\llbracket \cdot \rrbracket_r$ stands for $1 + K RNat \times R$. The constructor `k` takes a value of some interpretation for a type representation s and embeds it in the interpretation for representations of type r . For example, when building values of type $\llbracket RList \rrbracket$, `k` is used to embed a natural number in the list. Similarly, the constructor `r` embeds list values in a larger list. The `in` constructor embeds values of type $\llbracket r \rrbracket_r$ into a value of inductive family $\llbracket r \rrbracket$, playing the role of a fixpoint. The remaining data constructors (for representing unit, sums and products values) have the expected role, allowing sum-of-product values to be created.

As a final remark, note that, in this universe, `in` and $\llbracket \cdot \rrbracket$ are redundant because the occurrences of $\llbracket \cdot \rrbracket$ in $\llbracket \cdot \rrbracket_r$, in the data constructors `r` and `k`, could have been replaced, respectively, by $\llbracket r \rrbracket_r$ and $\llbracket R \rrbracket_s$. However, the stratification into two inductive families plays an important role in Section 4.3, when the universe is extended with support for variables.

Generic Functions The key advantage of universes is that we can define (generic) functions that work for any representable datatypes. A simple example is a generic function counting the number of recursive occurrences on a term:

```
size :  $\forall (r : Rep). \llbracket r \rrbracket \rightarrow \mathbb{N}$ 
size (in t) = size t
```

³Recall that, as explained in Section 4.1, parametric types are visible in both the type and data constructors.

$$\begin{array}{l}
\text{DATA Rep} = \dots \mid \text{E Rep} \mid \text{B Rep Rep} \\
\\
Q : \star \quad (* \text{ Quantifier type } *) \\
V : \star \quad (* \text{ Variable type } *) \\
\\
\text{DATA } \frac{r, s : \text{Rep}}{\llbracket s \rrbracket_r : \star} \text{ WHERE } \dots \frac{s : \text{Rep} \quad v : \llbracket s \rrbracket}{e \ v : \llbracket \text{E } s \rrbracket_r} \quad \frac{s_1, s_2 : \text{Rep} \quad q : Q \quad v : \llbracket s_2 \rrbracket_r}{\lambda_{s_1} q.v : \llbracket \text{B } s_1 \ s_2 \rrbracket_r} \\
\\
\text{DATA } \frac{s : \text{Rep}}{\llbracket s \rrbracket : \star} \text{ WHERE } \dots \frac{s : \text{Rep} \quad v : V}{\text{var } v : \llbracket s \rrbracket}
\end{array}$$

Figure 6: Extending universe with representations of binders and variables.

$$\begin{array}{l}
size : \forall (r, s : \text{Rep}). \llbracket s \rrbracket_r \rightarrow \mathbb{N} \\
size () = 0 \\
size (k \ t) = 0 \\
size (i_1 \ t) = size \ t \\
size (i_2 \ t) = size \ t \\
size (t, v) = size \ t + size \ v \\
size (r \ t) = 1 + size \ t
\end{array}$$

To define such generic function, two-mutually inductive definitions are needed: one inductively defined on $\llbracket r \rrbracket$; and another inductively defined on $\llbracket s \rrbracket_r$. For convenience the same name *size* is used in both definitions. Note that r and s (bound by \forall) are implicitly passed in the calls to *size*.

When interpreted on values of type $\llbracket \text{RNat} \rrbracket$, *size* computes the value of the represented natural number. When interpreted on values of type $\llbracket \text{RList} \rrbracket$, *size* computes the length of the list. What is great about this function is that it works, not only for these types, but for any datatypes representable in the universe.

4.3 A Universe for Representing First-Order Binding

In this paper our goal is to define common infrastructure definitions and lemmas for first-order representations, once and for all, using generic functions and lemmas. However, the universe presented in Figure 5 is insufficient for this purpose because generic functions such as substitution and free variables require structural information about binders and variables. Therefore, we first need to enrich our universe to support these constructs.

Extended Universe Figure 6 shows the additional definitions required to support representations of binders, variables, and also deeply embedded terms. The data constructor B of the datatype Rep provides the type for representations of binders. The type Rep is also extended with a constructor E which is the representation type for deeply embedded terms. This constructor is very similar to K . However, the fundamental difference is that generic functions should go inside the terms represented by deeply embedded terms, whereas terms built with K should be treated as constants by generic functions.

The abstract types Q and V represent the types of quantifiers and variables. Depending on the particular first-order representations of binders these types will be instantiated differently. The following table shows the instantiations of Q and V for 4 of the most popular first-order representations:

$\text{RLambda} : \text{Rep}$
 $\text{RLambda} = \mathbb{R} \times \mathbb{R} + \mathbb{B} \mathbb{R} \mathbb{R}$

$\text{fvar} : \mathbb{N} \rightarrow \llbracket \text{RLambda} \rrbracket$
 $\text{fvar } n = \text{var} (\text{inl } n)$

$\text{bvar} : \mathbb{N} \rightarrow \llbracket \text{RLambda} \rrbracket$
 $\text{bvar } n = \text{var} (\text{inr } n)$

$\text{app} : \llbracket \text{RLambda} \rrbracket \rightarrow \llbracket \text{RLambda} \rrbracket \rightarrow \llbracket \text{RLambda} \rrbracket$
 $\text{app } e_1 e_2 = \text{in} (\text{i}_1 (\text{r } e_1, \text{r } e_2))$

$\text{lam} : \llbracket \text{RLambda} \rrbracket \rightarrow \llbracket \text{RLambda} \rrbracket$
 $\text{lam } e = \text{in} (\text{i}_2 (\lambda_{\mathbb{R}} \mathbb{1}. \text{r } e))$

Figure 7: The untyped lambda calculus using the locally nameless approach.

	Q	V	$\lambda x. x y$
Nominal	\mathbb{N}	\mathbb{N}	$\lambda x. x y$
De Bruijn	$\mathbb{1}$	\mathbb{N}	$\lambda. 0 \ 1$
Locally nameless	$\mathbb{1}$	$\mathbb{N} + \mathbb{N}$	$\lambda. 0 \ y$
Locally named	\mathbb{N}	$\mathbb{N} + \mathbb{N}$	$\lambda x. x a$

The last column of the table shows how the lambda term $\lambda x. x y$ can be encoded in the different approaches. For the nominal approach there is only one sort of variables, which can be represented by a natural number (alternatively a string could be used instead). In this representation, the binders hold information about the bound variables, thus the type Q is the same type as the type of variables V : a natural number. De Bruijn indices do not need to hold information about variables in the binders, because the variables are denoted positionally with respect to the current enclosing binder. Thus, in the de Bruijn style, the type Q is just the unit type and the type V is a natural number. The locally nameless approach can be viewed as a variant of the de Bruijn style. Like de Bruijn, no information is needed at the binders, thus the type Q is just the unit type. The difference to the de Bruijn style is that parameters and (bound) variables are distinguished: (bound) variables are represented in the same way as de Bruijn variables; but parameters belong to another sort of variables. Therefore in the locally nameless style the type V is instantiated to a sum of two natural numbers. Finally, in the locally named style, there are also two sorts of variables. However, bound variables are represented as in the nominal style instead. Thus the type Q is a natural number and the type V is a sum type of two naturals.

The inductive family $\llbracket \cdot \rrbracket_r$ is extended with two new data constructors. The constructor e is similar to the constructor k and is used to build deeply embedded terms. The other constructor uses the standard lambda notation $\lambda_{s_1} q.v$ to denote the constructor for binders. The type representation s_1 is the representation of the syntactic sort of the variables that are bound by the binder, whereas the type representation s_2 is the representation of the syntactic sort of the body of the abstraction. We use $s_1 = \mathbb{R}$ to denote that the syntactic sort of the variables to be bound is the same as that of the body. This distinction is necessary because in certain languages the syntactic sorts of variables to be bound and the body of the abstraction are not the same. For example, in System F, type abstractions in terms such as $\Lambda X.e$ bind type variables X in a term e .

The inductive family $\llbracket \cdot \rrbracket$ is also extended with one additional data constructor for variables. This constructor allows terms to be constructed using a variable instead of a concretely defined term.

Untyped lambda calculus using locally nameless representations As a simple example demonstrating the use of the universe to represent languages with binding constructs, we show how the untyped lambda calculus is encoded in Figure 7. The definition `RLambda` is the representation type for the untyped lambda calculus. Note that the variable case is automatically built in, so the representation only needs to account for the application and abstraction cases. Application consists of a constructor with two recursive arguments and it is represented by the product type on the left side of the sum. Abstraction is a binder with a recursive argument and is represented by the right side of the sum. The four definitions `fvar`, `bvar`, `app` and `lam` provide shorthands the corresponding parameters, variables, application and abstraction constructors. Terms can be built using these constructors. For example, the identity function is defined as:

```
idRLambda : [[RLambda]]
idRLambda = lam (bvar 0)
```

5 Modular Parameterization on Representations

This section shows how generic operations and lemmas defined over the universe presented in Section 4 can be used to provide much of the basic infrastructure boilerplate once and for all for the languages representable in the universe for two of the most important first-order representations: locally nameless and de Bruijn.

5.1 Locally Nameless

Figure 8 presents generic definitions for the locally nameless approach. In this approach binders do not bind names, and (bound) variables and parameters (free variables) are distinguished. Thus, as discussed in Section 4.3, the types Q and V are, respectively, the unit type⁴ and a sum of two naturals. Using these instantiations for Q and V , the *set of parameters* and *substitution* operations can be defined *generically* for the locally nameless approach. Furthermore the operation for instantiating a (bound) variable with a term is also defined in a generic way. Finally, generic lemmas can be defined using the generic operations. The statements for *subst_fresh* – which states that if a parameter does not occur in a term, then substitution of that parameter is the identity – and *bfsbst_perm* – which states that substitutions for parameters and variables can be exchanged under certain well-formedness conditions – are shown as examples of such generic lemmas.

As explained in Section 4, generic operations are defined over terms of the universe by two mutually-inductive operations defined over the $[[\cdot]]$ and $[[\cdot]]_r$ (mutually-)inductive families. Note that, for convenience, we use same function names for mutual definitions.

The operation fv_{r_1} computes the set of parameters in a term. The only interesting case happens with parameters:

$$fv_{r_1}(\text{var}(\text{inl } x)) = \text{if } r_1 \equiv r_2 \text{ then } \{x\} \text{ else } \emptyset$$

In this case a singleton set containing the parameter is returned only when the type representations r_1 and r_2 are the same. That is, the computation of parameter sets depends on the representation r_1 . For example, in System F, if r_1 is the type representation for System F types, then fv_{r_1} computes the set of type parameters which occur in a term or a type (depending on what r_2 represents). Note that the constructor `inl` arises from the instantiation of $V = \mathbb{N} + \mathbb{N}$. This constructor signals that the case under consideration is the left case of the sum type, which represents parameters. Variables, on the other hand, are represented by the right case of the sum type, which is signaled by the `inr` constructor. The other cases of fv_{r_1} are straightforward.

⁴For convenience, we use `1` for both the unit type and the unique term of that type.

Instantiation of Q and V :

$$Q = \mathbb{1}$$

$$V = \mathbb{N} + \mathbb{N}$$

Heterogeneous sets of parameters (free variable): Given $r_1 : \text{Rep}$,

$$fv_{r_1} : \forall (r_2 : \text{Rep}). \llbracket r_2 \rrbracket \rightarrow 2^{\mathbb{N}}$$

$$fv_{r_1} (\text{in } t) = fv_{r_1} t$$

$$fv_{r_1} (\text{var } (\text{inl } x)) = \text{if } r_1 \equiv r_2 \text{ then } \{x\} \text{ else } \emptyset$$

$$fv_{r_1} (\text{var } (\text{inr } y)) = \emptyset$$

$$fv_{r_1} : \forall (r_2, s : \text{Rep}). \llbracket s \rrbracket_{r_2} \rightarrow 2^{\mathbb{N}}$$

$$fv_{r_1} () = \emptyset$$

$$fv_{r_1} (\mathbf{k} t) = \emptyset$$

$$fv_{r_1} (\mathbf{e} t) = fv_{r_1} t$$

$$fv_{r_1} (\mathbf{i}_1 t) = fv_{r_1} t$$

$$fv_{r_1} (\mathbf{i}_2 t) = fv_{r_1} t$$

$$fv_{r_1} (t, v) = (fv_{r_1} t) \cup (fv_{r_1} v)$$

$$fv_{r_1} (\lambda_{r_3} \mathbb{1}.t) = fv_{r_1} t$$

$$fv_{r_1} (\mathbf{r} t) = fv_{r_1} t$$

Heterogeneous substitution for parameters:

$$[\cdot \rightarrow \cdot] \cdot : \forall (r_1 r_2 : \text{Rep}). \mathbb{N} \rightarrow \llbracket r_1 \rrbracket \rightarrow \llbracket r_2 \rrbracket \rightarrow \llbracket r_2 \rrbracket$$

$$[k \rightarrow u] (\text{in } t) = \text{in } ([k \rightarrow u] t)$$

$$[k \rightarrow u] (\text{var } (\text{inl } x)) = \text{if } r_1 \equiv r_2 \wedge k \equiv x \text{ then } u \text{ else } (\text{var } (\text{inl } x))$$

$$[k \rightarrow u] (\text{var } (\text{inr } y)) = \text{var } (\text{inr } y)$$

$$[\cdot \rightarrow \cdot] \cdot : \forall (r_1, r_2, s : \text{Rep}). \mathbb{N} \rightarrow \llbracket r_1 \rrbracket \rightarrow \llbracket s \rrbracket_{r_2} \rightarrow \llbracket s \rrbracket_{r_2}$$

$$[k \rightarrow u] () = ()$$

$$[k \rightarrow u] (\mathbf{k} t) = \mathbf{k} t$$

$$[k \rightarrow u] (\mathbf{e} t) = \mathbf{e} ([k \rightarrow u] t)$$

$$[k \rightarrow u] (\mathbf{i}_1 t) = \mathbf{i}_1 ([k \rightarrow u] t)$$

$$[k \rightarrow u] (\mathbf{i}_2 t) = \mathbf{i}_2 ([k \rightarrow u] t)$$

$$[k \rightarrow u] (t, v) = ([k \rightarrow u] t, [k \rightarrow u] v)$$

$$[k \rightarrow u] (\lambda_{r_3} \mathbb{1}.z) = \lambda_{r_3} \mathbb{1}.([k \rightarrow u] z)$$

$$[k \rightarrow u] (\mathbf{r} t) = \mathbf{r} ([k \rightarrow u] t)$$

Heterogeneous substitution for (bound) variables:

$$\{\cdot \rightarrow \cdot\} \cdot : \forall (r_1, r_2 : \text{Rep}). \mathbb{N} \rightarrow \llbracket r_1 \rrbracket \rightarrow \llbracket r_2 \rrbracket \rightarrow \llbracket r_2 \rrbracket$$

$$\{k \rightarrow u\} (\text{in } t) = \text{in } (\{k \rightarrow u\} t)$$

$$\{k \rightarrow u\} (\text{var } (\text{inl } x)) = \text{var } (\text{inl } x)$$

$$\{k \rightarrow u\} (\text{var } (\text{inr } y)) = \text{if } r_1 \equiv r_2 \wedge k \equiv y \text{ then } u \text{ else } (\text{var } (\text{inr } y))$$

$$\{\cdot \rightarrow \cdot\} \cdot : \forall (r_1, r_2, s : \text{Rep}). \mathbb{N} \rightarrow \llbracket r_1 \rrbracket \rightarrow \llbracket s \rrbracket_{r_2} \rightarrow \llbracket s \rrbracket_{r_2}$$

$$\{k \rightarrow u\} () = ()$$

$$\{k \rightarrow u\} (\mathbf{k} t) = \mathbf{k} t$$

$$\{k \rightarrow u\} (\mathbf{e} t) = \mathbf{e} (\{k \rightarrow u\} t)$$

$$\{k \rightarrow u\} (\mathbf{i}_1 t) = \mathbf{i}_1 (\{k \rightarrow u\} t)$$

$$\{k \rightarrow u\} (\mathbf{i}_2 t) = \mathbf{i}_2 (\{k \rightarrow u\} t)$$

$$\{k \rightarrow u\} (t, v) = (\{k \rightarrow u\} t, \{k \rightarrow u\} v)$$

$$\{k \rightarrow u\} (\lambda_{r_3} \mathbb{1}.t) = \text{if } (r_3 \equiv \mathbf{R} \wedge r_1 \equiv r_2) \vee (r_3 \not\equiv \mathbf{R} \wedge r_1 \equiv r_3) \\ \text{then } \lambda_{r_3} \mathbb{1}.(\{k \rightarrow u\} t) \text{ else } \lambda_{r_3} \mathbb{1}.(\{k \rightarrow u\} t)$$

$$\{k \rightarrow u\} (\mathbf{r} t) = \mathbf{r} (\{k \rightarrow u\} t)$$

Some heterogeneous lemmas:

$$\text{subst_fresh} : \forall (r_1, r_2 : \text{Rep}) (t : \llbracket r_1 \rrbracket) (u : \llbracket r_2 \rrbracket) (m : \mathbb{N}).$$

$$m \notin (fv_{r_2} t) \Rightarrow [m \rightarrow u] t = t$$

$$\text{bfsubst_perm} : \forall (r_1, r_2, r_3 : \text{Rep}) (t : \llbracket r_1 \rrbracket) (u : \llbracket r_2 \rrbracket) (v : \llbracket r_3 \rrbracket) (m k : \mathbb{N}).$$

$$(wf_{r_3} u) \Rightarrow \{k \rightarrow ([m \rightarrow u] v)\} ([m \rightarrow u] t) = [m \rightarrow u] (\{k \rightarrow v\} t)$$

Figure 8: Generic definitions for the locally nameless approach.

Instantiation of Q and V : $Q = \mathbb{1}$ and $V = \mathbb{N}$.

Heterogeneous shifting: Given $r_1 : \text{Rep}$,

$$\begin{aligned}
\uparrow \cdot & : \forall (r_2 : \text{Rep}). \mathbb{N} \rightarrow \llbracket r_2 \rrbracket \rightarrow \llbracket r_2 \rrbracket \\
\uparrow_m (\text{in } t) & = \text{in } (\uparrow_m t) \\
\uparrow_m (\text{var } n) & = \text{if } r_1 \equiv r_2 \wedge m \leq n \text{ then } (\text{var } (n + 1)) \text{ else } (\text{var } n) \\
\uparrow \cdot & : \forall (r_1, r_2, s : \text{Rep}). \mathbb{N} \rightarrow \llbracket s \rrbracket_{r_2} \rightarrow \llbracket s \rrbracket_{r_2} \\
\uparrow_m () & = () \\
\uparrow_m (\mathbf{k} t) & = \mathbf{k} t \\
\uparrow_m (\mathbf{e} t) & = \mathbf{e} (\uparrow_m t) \\
\uparrow_m (\mathbf{i}_1 t) & = \mathbf{i}_1 (\uparrow_m t) \\
\uparrow_m (\mathbf{i}_2 t) & = \mathbf{i}_2 (\uparrow_m t) \\
\uparrow_m (t, v) & = (\uparrow_m t, \uparrow_m v) \\
\uparrow_m (\lambda_{r_3} \mathbb{1}. t) & = \text{if } (r_3 \equiv \mathbf{R} \wedge r_2 \equiv r_1) \vee (r_3 \not\equiv \mathbf{R} \wedge r_3 \equiv r_1) \\
& \quad \text{then } \lambda_{r_3} \mathbb{1}. (\uparrow_{(m+1)} t) \text{ else } \lambda_{r_3} \mathbb{1}. (\uparrow_m t) \\
\uparrow_m (\mathbf{r} t) & = \mathbf{r} (\uparrow_m t)
\end{aligned}$$

Figure 9: Heterogeneous shifting for de Bruijn representations.

In the generic definitions for substitutions⁵ the interesting cases are variables and binders. In the case of variables, the condition $r_1 \equiv r_2$ is necessary to check whether the parameter (or variable) and the term to be substituted have the same representation. The binder case in the heterogeneous substitution for variables is more interesting. The subscript r_3 keeps the information about which kind of variables is to be bound. When $r_3 = \mathbf{R}$, the binding is homogeneous, that is, the variable to be bound and the body of the binder have the same representation. For example, the term-level abstraction in terms $(\lambda x : T.e)$ of System F is homogeneous. An example of heterogeneous binding is the type-level abstraction in terms $(\Lambda X.e)$ of System F. In this case r_3 is the representation for System F types. Variable shifting happens when the bound variable and the terms to be substituted have the same representation. Note that, in the case of homogeneous binding ($r_3 \equiv \mathbf{R}$), we compare r_1 with r_2 , not with r_3 , because the bound variable and the body of the binder have the same representation r_2 .

The main advantage of representing the syntax of languages with our generic universe is, of course, that all generic operations are immediately available. For instance, the 8 substitution operations mentioned in Section 1 can be recovered through suitable instantiations of the type representations r_1, r_2, r_3 in the two generic substitutions presented in this section.

5.2 De Bruijn

A key advantage of our modular approach is that we do not have to commit to using a particular first-order representation. Instead, by suitably instantiating the types Q and V , we can define the generic infrastructure for our own favored first-order representation. For example we can use GMETA to define the generic infrastructure for de Bruijn representations. In de Bruijn representations, binders do not bind any names, therefore the type Q is instantiated with the unit type. Also, because there is only one sort of (positional) variables, the type V is instantiated with natural numbers. For space reasons we only show how to define one operation, *shifting*, in Figure 9. Other generic operations like substitution are similar to the locally nameless style, and generic lemmas about operations on de Bruijn representations are also definable generically. The implementation of heterogeneous generic shifting is quite simple and follows a pattern similar to that used in the generic operations for the locally nameless style for dealing with homogeneous and heterogeneous binders. The variable and binder cases implement the expected behavior for the de Bruijn indices operations and all the other cases

⁵Note that the notation for substitutions follows Aydemir et al. (2008).

		Definitions	Infrastructure (lemmas + definitions)	Core (lemmas + definitions)	Overall		
					inf. overhead	total	ratio
STLC (locally nameless)	Aydemir et al.	11	13 + 3	4 + 0	17	31	55%
	GMETA basic	7	4 + 0	4 + 0	1	15	7%
System $F_{<}$: (locally nameless)	Aydemir et al.	20	48 + 7	17 + 1	60	93	65%
	GMETA basic	13	26 + 1	17 + 1	25	58	43%
	GMETA advanced	11	15 + 0	18 + 1	11	45	24%
System $F_{<}$: (de Bruijn)	Vouillon	27	24 + 0	50 + 0	41	101	41%
	GMETA basic	12	1 + 0	52 + 0	3	65	5%

Figure 10: Formalization of POPLmark challenge (part 1A+2A) and STLC in Coq using locally nameless approach and de Bruijn approach with and without GMETA.

		style	savings
STLC	GMETA basic vs Aydemir et al.	LN	52%
$F_{<}$:	GMETA basic vs Aydemir et al.	LN	38%
	GMETA adv. vs Aydemir et al.	LN	52%
	GMETA basic vs Vouillon	dB	35%

Figure 11: Savings in various formalizations in terms of numbers of definitions and lemmas.

are limited to traversal code. Note that we follow Vouillon (2007)’s definitions.

6 Discussion and Evaluation

In this section we present the results of the case studies we conducted. The discussion of these results is done in terms of three criteria proposed by Aydemir et al. (2005) (*reasonable overheads*, *cost of entry* and *transparency*) for evaluating mechanizations of formal metatheory. We also discuss how theorem provers like Coq could be improved with better support for GMETA.

6.1 Reasonable overheads

The biggest benefit of GMETA is that it significantly lowers the overheads required in mechanical formalizations by providing reuse of the basic infrastructure. Figure 11 shows the savings that GMETA achieved relative to the reference solutions by Aydemir et al. (2008). In all case studies more than 35% of the total numbers of definitions were saved. We conducted case studies in both System $F_{<}$: (using locally nameless and de Bruijn representations) and STLC (using locally nameless representations). We also employed GMETA in two modes, basic and advanced (see Section 3.2 for the explanation). In the STLC case there was no need to employ the advanced mode, since with the basic approach reuse was already as good as it could be. In System $F_{<}$: we need some lemmas of type B3, which cannot be dealt using just the basic approach, so we also used the advanced approach. In the de Bruijn case study, only the basic approach was employed.

Detailed results Figure 10 presents the detailed numbers obtained in our case studies. We follow Aydemir et al. by dividing the whole development into three parts: *definitions*, *infrastructure* and *core*. The numbers on those columns correspond to the number of definitions and lemmas used for each part. The definitions column presents the number of basic definitions about syntax, whereas the core column presents the number of main definitions and lemmas of the formalization (such as, for example, progress and preservation). The infrastructure column is the most interesting because this is where most of the tedious infrastructure lemmas and definitions are. The column inf. overhead counts the overall number of infrastructure definitions and lemmas used across the formalizations. Although, for the most part, infrastructure comes from what was classified by Aydemir et al. as the infrastructure part, some additional

infrastructure definitions and lemmas also exists in the definitions part. This explains why GMETA is able to reduce the number of definitions and lemmas in the two parts. The numbers in bold face, are the numbers that were presented by Aydemir et al. (2008). However those numbers did not reflect the real total number of definitions and lemmas in the solutions. For example, in the infrastructure part only the lemmas were counted. Since we are interested in all the infrastructure, our numbers reflect the total number of definitions and lemmas in each part.

In some cases we needed a few lemmas simply to slightly adapt the form of a lemma used in GMETA libraries to the form used in (Aydemir et al. 2008). This explains, for example the 1 in the inf. overhead column for STLC.

Basic vs Advanced The main difference between the basic and advanced approaches is that the basic approach allows all types of infrastructure to be reused except for infrastructure of type B3, whereas in the advanced approach all types of infrastructure can be reused. To capture some of the B1 and B2 lemmas used in the basic mode for System $F_{<}$: we used a domain-specific template for typed languages involve at least 2 syntactic sorts (note that this is a fairly common type of object language).

Proof size In comparison with Aydemir et al. reference solutions, the proofs in the basic approach follow essentially the same structure of the original proofs. One minor difference is that instead of some standard Coq tactics, a few more general tactics provided by GMETA should be used (see Section 3.2). Because this is the only difference, the proofs in the GMETA solution and Aydemir et al. solution have comparable sizes. In the advanced approach one additional difference is that the formalization of environments differs from Aydemir et al., in that we use two environments (for types and terms) and Aydemir et al. use a single environment with a disjoint union. This difference means that, while most proofs will still be comparable in size, a small number of proofs will be either shorter or larger.

6.2 Cost of entry

One important criteria for evaluating mechanical formalizations of metatheory is the associated cost of entry. That is, how much does a user need to know in order to successfully develop a formalization. We believe that the associated cost of entry of GMETA is comparable to first-order approaches like the one by Aydemir et al. (2008) (at least for the basic approach).

Essentially GMETA has a pay-as-you-go approach in terms of the cost-of-entry. The basic approach gives us a good cost/benefit ratio. Basically, with developments using the basic mode we get quite a bit of reuse almost for free. reference solution. In the advanced approach additional savings are possible, at the cost of some extra knowledge about DGP and GMETA.

One aspect of GMETA that (arguably) requires less knowledge when compared to Aydemir et al. (2008) is that the end-user does not need to know how to prove many basic infrastructure lemmas, since those are provided by GMETA's libraries.

6.3 Transparency

The transparency criteria is intended to evaluate how intuitive a formalization technique is. The issue of transparency is largely orthogonal to GMETA because it usually measures how particular representations of binding (such as locally nameless or de Bruijn), and lemmas and definitions using that approach, are easy to understand by humans. Since we do not introduce any new representation, transparency remains unchanged (the same representation, lemmas and definitions are used).

6.4 Improving Coq and Other Theorem Provers

We found a few shortcomings in Coq during our experience with GMETA. These affected the development of GMETA and also its usability. It would be nice to have the following improvements to address those shortcomings.

Improved support for dependently typed programming GMETA makes significant use of dependently typed programming. Being based on the calculus of inductive constructions, Coq has good support for dependent types. However, only very recently more attention has been given to providing better support for dependent pattern matching (Coquand 1992) as done, for example, in the Agda or Epigram dependently typed languages. In this respect the work by Sozeau (2007) on the PROGRAM tactic and the other related tactics like *dependent destruction* is certainly a good step forward. Unfortunately, support for these features is still experimental on the current version of Coq (8.3) and there is still some gap to dependently typed languages like Agda or Epigram.

Our experience with GMETA is that there are still a few bugs lying around in the current implementation of Coq which make the development of dependently typed programs that use dependent pattern matching difficult. We found *nested* dependent pattern matching and the use mutual induction with *dependent destruction* to be particularly tricky in Coq. Curiously we only need these features for defining isomorphisms, which (technically speaking) are a fairly uninteresting part of GMETA (since usually the definition of isomorphisms in DGP is supposed to be straightforward and mechanical). More concretely, the need for nested dependent pattern matching arises from the *to* function (which converts from a representation of the object language in the universe to the user defined object language) and the *to_from* lemma (which states that $to\ (from\ t) = t$).

With the very valuable help of Chlipala upcoming book on “Certified Programming with Dependent Types” (Chlipala 2009) we did manage to learn enough techniques to tame Coq and make our programs and proofs pass the type and termination checkers. However in our experimental version of GMETA which is based on the more expressive universe for regular tree types by (Morris et al. 2004) and it requires even more dependent types, we found that the techniques we used earlier no longer worked and we are still looking for a way to implement isomorphisms (although the remaining parts of the framework work as expected).

In contrast, our experience with developing a version of GMETA in Agda was much smoother due to the good support for dependent pattern matching. In Agda, the definitions of isomorphisms were essentially painless. Of course, the main reason not to use Agda in the first place is that the main purpose of GMETA is to assist with theorem proving of mechanized metatheory and Coq support for theorem proving is generally better than Agda, which is focused primarily on programming.

To make a long story short, we believe that support for dependent pattern matching in Coq could be improved in two ways. The first way is simply to polish the current implementation to make it bug free and well supported. We expect that in the next few versions of Coq this will happen. The second improvement would be to provide better assistance when developing programs using dependent pattern matching. The emacs mode for Agda and Epigram are good examples of how the “IDE” can provide valuable help to the user.

Partial evaluation of modules One disadvantage of using DGP for reusing infrastructure is that such infrastructure is written generically. When trying to prove some lemmas users may have to unfold such generic definitions. The way that GMETA currently deals with this problem is through the use of some special tactics like *gsimpl* or *grewrite* (as discussed in Section 3), which make use of isomorphisms and their laws to automatically specialize definitions into the form that the user would expect. However this approach is not ideal because it requires the user to know about the special tactics and when to use them, so a nicer alternative is desirable.

We believe some kind of module-level partial evaluation would be a better solution for specializing infrastructure for object languages and eliminating the DGP indirections. Ideally when instantiating templates like *LNTemplate* with particular isomorphisms we would like to partially evaluate the resulting module so that all the isomorphism indirections would be evaluated and removed. The work by Brady and Hammond (2010) on partial evaluation for dependently typed languages and the work by Alimarine and Smetsers (2004) on partial evaluation of generic functions should be good starting points to investigate this possibility further.

In essence, with module-level partial evaluation we would get the best of DGP and generative approaches: verified generic infrastructure that can be specialized to a form similar to hand-written definitions.

Built-in support for isomorphisms Finally, it would be better to support isomorphism generation directly in Coq, instead of using of an external tool. This could be done for example via a plugin. The idea is simply to support the isomorphism annotations; and to generate and make the isomorphism available to the user directly, rather than calling an external tool that generates a Coq module that needs to be imported later. This could be supported in very much the same way that inductive principles and recursion principles are supported in Coq for inductive definitions. Basically, Coq automatically generates the associated principles when we define an inductive definition, and those principles become immediately available to the user.

7 Related Work

Generative Approaches Closest to our work are *generative* approaches like LNgen (Aydemir and Weirich 2009), which uses an external tool, based on Ott (Sewell et al. 2010) specifications, to generate the infrastructure lemmas and definitions for a particular language automatically. One advantage of generative approaches is that the generated infrastructure is directly defined in terms of the object language. In contrast, in GMETA, the infrastructure is indirectly defined in terms of generic definitions. This is not entirely ideal, but it is possible to handle the situation in a reasonably effective way in GMETA using tactics (see Section 3.2). As we have discussed in Section 6.4, allowing some kind of module-level partial evaluation would be nicer and would allow us to recover direct definitions from the generic infrastructure.

There are two main advantages of a DGP approach over generative approaches: 1) *verifiability*; and 2) *extensibility*. Like with DGP, a generator allows defining once-and-forall the infrastructure. However, unlike DGP, we cannot verify once-and-forall that the generator *always* generates correct (well-typed) infrastructure. With a generator we can only verify whether each particular generated set of infrastructure is correct. Another disadvantage of generators is that they are hard to extend. If we wanted to add a new lemma to the set of generated infrastructure we would have to directly change the generator code. With a library approach like GMETA this is much easier because it amounts to extending a module with a new generic function.

It is also interesting to compare GMETA and LNgen in terms of which types of infrastructure they can reuse and how hard it is to reuse such infrastructure. The main advantage of LNgen is that dealing inductive relations is easy. In GMETA, type 3 definitions require some more effort to be reused. A solution for this problem would be to extend the isomorphism generator tool to deal with inductive relations as well. On the other hand the strength of GMETA lies on its extensibility. It is easy to capture some types of domain-specific infrastructure, which are still common enough that we may want to capture in a reusable form. The lemma *thbfsbst_perm_core* presented in Section 2.2 is one example. Dealing with domain-specific infrastructure like that is in conflict with the general purpose nature of LNgen.

DGP and Binding DGP techniques have been used before for dealing with binders using a well-scoped de Bruijn indices representation (Altenkirch and Reus 1999; McBride and McKinna 2004). Chlipala (2007) used an approach inspired by *proof by reflection techniques* (Boutin 1997) to provide several generic operations on well-typed terms represented by well-scoped de Bruijn indices. Licata and Harper (2009) proposed a universe in Agda that permits definitions that mix binding and computation. The obvious difference is that GMETA works with traditional (non well-scoped) first-order representations instead of well-scoped de Bruijn indices. This difference of representation, means that the universes and generic functions have to deal with significantly different issues and that they are quite different in nature. More fundamentally, Chlipala (2007) and Licata and Harper (2009) work can be viewed as trying to develop new ways to formalize metatheory, in which many of the invariants that would have to be proved otherwise hold by construction. This is different from our goal: we are not proposing new ways to formalize metatheory, rather we wish to make well-established ways to formalize metatheory with first-order representations less painful to use.

DGP techniques have also been widely used in conventional functional programming languages (Jansson and Jeuring 1997; Hinze and Jeuring 2003; Lämmel and Peyton Jones 2003; Rodriguez et al. 2009), and Cheney (2005) explored how to provide generic operations such as substitution or free variables using nominal abstract syntax.

Our work is inspired by the use of *universes* in type-theory (Martin-Löf 1984; Nordström et al. 1990). The basic universe construction presented in Figure 5 is a simple variation of the *regular tree types* universe proposed by Morris et al. (2004, 2009) in Epigram. Nevertheless the extensions for representing variables and binders presented in Figure 6 are new. Dybjer and Setzer (1999, 2001) showed universe constructions within a type-theory with an axiomatization of induction recursion. Altenkirch and McBride (2003) proposed a universe capturing the datatypes and generic operations of Generic Haskell (Hinze and Jeuring 2003) and Norell (2008) shows how to do DGP with universes in Agda (Norell 2007).

Verbruggen et al. (2008, 2009) formalized a Generic Haskell (Hinze and Jeuring 2003) DGP style in Coq, which can also be used to do generic programming. This approach allows conventional datatypes to be expressed, but it cannot be used to express meta-theoretical generic operations since there are no representations for variables or binders.

Other Techniques for First-Order Approaches Aydemir et al. (2009) investigated several variations of representing syntax with locally nameless representations aimed at reducing the amount of infrastructure overhead in languages like System $F_{<}$. One advantage of these techniques is that they are very lightweight in nature and do not require additional tool support. However, while the proposed techniques are effective at achieving significant savings, they require the abstract syntax of the object language to be encoded in a way different from the traditional locally nameless style, potentially collapsing all syntactic sorts into one. In contrast GMETA allows the syntax to be encoded in the traditional locally nameless style, while at the same time reducing the infrastructure overhead through its reusable libraries of infrastructure.

Higher-order Approaches and Nominal Logic Approaches based on higher-order abstract syntax (HOAS) (Pfenning and Elliot 1988; Harper et al. 1993) are used in logical frameworks such as Abella (Gacek 2008), Hybrid (Momigliano et al. 2008) or Twelf (Pfenning and Schürmann 1999). In HOAS, the object-language binding is represented using the binding of the meta-language. This has the important advantage that facts about substitution or alpha equivalence come for free since the binding infrastructure of the meta-language is reused. It is well-known that in Coq it is not possible to use the usual HOAS encodings, although Despeyroux et al. (1995); Chlipala (2008) have shown how weaker variations of HOAS can be encoded in Coq. Approaches like GMETA or LNgen are aimed at recovering many of the properties that one expects from a logical framework for free.

Nominal logic (Pitts 2003) is an extension of first-order logic that allows reasoning about alpha-equivalent abstract syntax in a generic way. Variants of nominal logic have been adopted in the nominal Isabelle package (Urban 2005). However, because Coq does not have a nominal package, this approach cannot be used in Coq formalizations.

8 Conclusion

There are several techniques for formalizing metatheory using first-order representations, which typically involve developing the whole of the infrastructure by hand each time for a new formalization. GMETA improves on these techniques by providing reusable generic infrastructure in libraries, avoiding the repetition of definitions and lemmas for each new formalization. The DGP approach used by GMETA not only allows an elegant and verifiable formulation of the generic infrastructure, which is appealing from the theoretical point of view; but also reveals itself useful for conducting realistic formalizations of metatheory.

Acknowledgments

This work benefited from useful comments by Hugo Herberlin, Sungwoo Park, Jaeho Shin and Stephanie Weirich.

References

- A. Alimarine and S. Smetsers. Optimizing generic functions. In Dexter Kozen, editor, *Mathematics of Program Construction*, volume 3125, pages 16–31. 2004.
- T. Altenkirch and C. McBride. Generic programming within dependently typed programming. In *IFIP TC2/WG2.1 Working Conference on Generic Programming*, 2003.
- T. Altenkirch and B. Reus. Monadic presentations of lambda terms using generalized inductive types. In *CSL '99*, 1999.
- B. Aydemir, S. Weirich, and S. Zdancewic. Abstracting syntax. Technical Report MS-CIS-09-06, University of Pennsylvania, 2009.
- B. E. Aydemir and S. Weirich. LNgen: Tool Support for Locally Nameless Representations, 2009. Unpublished manuscript.
- B. E. Aydemir, A. Bohannon, M. Fairbairn, N. Nathan Foster, B. C. Pierce, P. Sewell, D. Vytiniotis, G. Washburn, S. Weirich, and S. Zdancewic. Mechanized Metatheory for the Masses: The POPLmark Challenge. In *TPHOLs '05*, 2005.
- B. E. Aydemir, A. Charguéraud, B. C. Pierce, R. Pollack, and S. Weirich. Engineering formal metatheory. In *POPL '08*, 2008.
- S. Boutin. Using reflection to build efficient and certified decision procedures. In *TACS'97*, 1997.
- Edwin C. Brady and Kevin Hammond. Scrapping your inefficient engine: using partial evaluation to improve domain-specific language implementation. In *ICFP '10*, pages 297–308, 2010.
- J. Cheney. Scrap your nameplate: (functional pearl). *SIGPLAN Not.*, 40(9):180–191, 2005.
- A. Chlipala. A certified type-preserving compiler from lambda calculus to assembly language. *SIGPLAN Not.*, 42(6):54–65, 2007.

- A. Chlipala. Parametric higher-order abstract syntax for mechanized semantics. *SIGPLAN Not.*, 43(9):143–156, 2008.
- A. Chlipala. *Certified Programming with Dependent Types*. 2009. URL <http://adam.chlipala.net/cpdt/>.
- The Coq Development Team. The Coq Proof Assistant Reference Manual, Version 8.2, 2009. Available at <http://coq.inria.fr>.
- T. Coquand. Pattern matching with dependent types. In B. Nordström, K. Pettersson, and G. Plotkin, editors, *Informal Proceedings Workshop on Types for Proofs and Programs, Båstad, Sweden, 8–12 June 1992*, pages 71–84. 1992.
- N. G. de Bruijn. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the church-rosser theorem. *Indagationes Mathematicae (Proceedings)*, 75(5):381–392, 1972.
- J. Despeyroux, A. P. Felty, and A. Hirschowitz. Higher-order abstract syntax in coq. In *TLCA '95*, 1995.
- P. Dybjer. Inductive families. *Formal Aspects of Computing*, 6:440–465, 1997.
- P. Dybjer and A. Setzer. A finite axiomatization of inductive-recursive definitions. In *TLCA '99*, 1999.
- P. Dybjer and A. Setzer. Indexed induction-recursion. In *PTCS '01*, 2001.
- A. Gacek. The abella interactive theorem prover (system description). In *IJCAR '08*, 2008.
- R. Harper, F. Honsell, and G. Plotkin. A framework for defining logics. *J. ACM*, 40(1):143–184, 1993.
- R. Hinze and J. Jeuring. Generic haskell: Practice and theory. In *Generic Programming '03*, 2003.
- P. Jansson and J. Jeuring. Polyp—a polytypic programming language extension. In *POPL '97*, 1997.
- R. Lämmel and S. Peyton Jones. Scrap your boilerplate: a practical design pattern for generic programming. *SIGPLAN Not.*, 38(3):26–37, 2003.
- D. R. Licata and R. Harper. A universe of binding and computation. In *ICFP '09*, 2009.
- P. Martin-Löf. *Intuitionistic Type Theory*. Bibliopolis, 1984.
- C. McBride and J. McKinna. The view from the left. *J. Funct. Program.*, 14(1):69–111, 2004.
- J. Mckinna and R. Pollack. Pure type systems formalized. In *TLCA '93*, pages 289–305. Springer-Verlag, 1993.
- A. Momigliano, A. J. Martin, and A. P. Felty. Two-level hybrid: A system for reasoning using higher-order abstract syntax. *Electron. Notes Theor. Comput. Sci.*, 196:85–93, 2008.
- P. Morris, T. Altenkirch, and C. McBride. Exploring the regular tree types. In *TYPES '04*, 2004.
- P. Morris, T. Altenkirch, and N. Ghani. A universe of strictly positive families. *Int. J. Found. Comput. Sci.*, 20(1):83–107, 2009.

- B. Nordström, K. Peterson, and J. M. Smith. *Programming in Martin-Löf's Type Theory: An Introduction*. Oxford University Press, 1990.
- U. Norell. *Towards a practical programming language based on dependent type theory*. PhD thesis, Chalmers University of Technology, 2007.
- U. Norell. Dependently typed programming in agda. In *Advanced Functional Programming '08*, 2008.
- C. Paulin-Mohring. *Définitions Inductives en Théorie des Types d'Ordre Supérieur*. Habilitation à diriger les recherches, Université Claude Bernard Lyon I, 1996.
- S. Peyton Jones, D. Vytiniotis, S. Weirich, and G. Washburn. Simple unification-based type inference for gadts. *SIGPLAN Not.*, 41(9):50–61, 2006.
- F. Pfenning and C. Elliot. Higher-order abstract syntax. *SIGPLAN Not.*, 23(7):199–208, 1988.
- F. Pfenning and C. Schürmann. System description: Twelf - a meta-logical framework for deductive systems. In *CADE '99*, 1999.
- B. C. Pierce. *Types and Programming Languages*. The MIT Press, 2002.
- A. M. Pitts. Nominal logic, a first order theory of names and binding. *Inf. Comput.*, 186(2):165–193, 2003.
- A. Rodriguez, J. Jeuring, P. Jansson, A. Gerdes, O. Kiselyov, and B. C. d. S. Oliveira. Comparing libraries for generic programming in haskell. *SIGPLAN Not.*, 44(2):111–122, 2009.
- A. Rossberg, C. V. Russo, and D. Dreyer. F-ing modules. In *TLDI '10*, 2010.
- P. Sewell, F. Z. Nardelli, S. Owens, G. Peskine, T. Ridge, S. Sarkar, and R. Strniša. Ott: Effective tool support for the working semanticist. *J. Funct. Program.*, 20(01):71–122, 2010.
- M. Sozeau. Program-ing finger trees in coq. In *ICFP 2007*, 2007.
- C. Urban. Nominal techniques in isabelle/hol. In *CADE '05*, pages 38–53, 2005.
- W. Verbruggen, E. de Vries, and A. Hughes. Polytypic programming in coq. In *WGP '08*, 2008.
- W. Verbruggen, E. de Vries, and A. Hughes. Polytypic properties and proofs in coq. In *WGP '09*, 2009.
- J. Vouillon. Poplmark solutions using de bruijn indices, 2007. Available at http://alliance.seas.upenn.edu/~plclub/cgi-bin/poplmark/index.php?title=Submission_by_J%C3%A9r%C3%B4me_Vouillon.