

Reasoning with Higher-Order Abstract Syntax and Contexts: A Comparison

Amy Felty¹ and Brigitte Pientka²

¹ SITE, University of Ottawa, Ottawa, Canada
afelty@site.uottawa.ca

² School of Computer Science, McGill University, Montreal, Canada
bpientka@cs.mcgill.ca

Abstract. A variety of logical frameworks support the use of higher-order abstract syntax (HOAS) in representing formal systems given via axioms and inference rules and reasoning about them. In such frameworks, object-level binding is encoded directly using meta-level binding. Although these systems seem superficially the same, they differ in a variety of ways; for example, in how they handle a context of assumptions and in what theorems about a given formal system can be expressed and proven. In this paper, we present several case studies which highlight a variety of different aspects of reasoning using HOAS, with the intention of providing a basis for qualitative comparison of different systems. We then carry out such a comparison among three systems: Twelf, Beluga, and Hybrid. We also develop a general set of criteria for comparing such systems. We hope that others will implement these challenge problems, apply these criteria, and further our understanding of the trade-offs involved in choosing one system over another for this kind of reasoning.

1 Introduction

In recent years, the POPLmark challenge [ABF⁺05] has stimulated considerable interest in mechanizing the meta-theory of programming languages, and the issued problems exercise many aspects that are known to be difficult to formalize. While several solutions have been submitted showing the diversity of possible approaches, it has been hard to compare them. Part of the reason is that while the proposed examples are typical for their domain, they do not highlight the differences between systems. We will bring a different view: As experts in designing and building logical frameworks, we propose a few challenge problems which highlight the differences between different meta-languages, and thereby hopefully provide a better understanding of what practitioners should be looking for.

Our focus in this paper is on encoding meta-theory of programming languages using higher-order abstract syntax (HOAS), where we encode object-level binders with meta-level binders. As a consequence, users can avoid implementing common and tricky routines dealing with variables, such as capture-avoiding substitution, renaming and fresh name generation. Because of this one can think

of HOAS encodings as the most advanced technology for specifying programming language meta-theory which leads to very concise and elegant encodings and provides the most support for such an endeavor. However concentrating on encoding binders neglects one important aspect: the support for hypothetical and parametric reasoning. Even in systems supporting HOAS, there is not a clear answer to this. On one side of the spectrum, we find the logical framework Twelf [PS99] or the dependently-typed functional language Beluga [Pie08,PD10]. Both systems provide direct support for contexts to keep track of hypotheses. In Twelf, contexts are implicit while in Beluga they are explicit. Supporting contexts directly has two advantages. First, it eliminates the need for building up a context and managing it explicitly via a first-order representation such as a list. More importantly, it eliminates the need to explicitly prove structural properties about contexts, such as weakening. Such built-in support for contexts allows for highly compact proofs. Second, using hypothetical and parametric reasoning provides us with direct meta-level support for applying substitution lemmas. Consequently, substitution lemmas come for free.

On the other side of the spectrum of systems supporting HOAS, we have, for instance, the two-level Hybrid system [MMF08,FM08] as implemented in Coq [BC04] and Isabelle/HOL [NPW02], Abella [Gac08], and the Tac prover [BSM10], where contexts are manually represented as lists. While the substitution lemma is still obtained for free because it is an application of the cut-rule, structural properties about contexts such as weakening must typically be proven separately as lemmas. These lemmas can be tedious and they may cause difficulties when automating induction proofs (see [BSM10]). On the other hand, since these systems do not rely on specific built-in procedures for dealing with contexts, there is more flexibility in how they are handled and the necessary reasoning is more transparent to the user. Consequently, proofs in these systems are often easier to understand and to trust.

This paper presents three case-studies which highlight the different treatments of hypothetical reasoning. Along the way, we develop a set of questions which allow a qualitative evaluation and comparison of different reasoning systems. These questions also provide guidance for users and developers in understanding better the differences and limitations. Due to space restrictions, we concentrate on the logical framework Twelf, the functional dependently-typed language Beluga, and the interactive theorem proving environment Hybrid. However, we hope that these problems will subsequently also be implemented using related approaches and serve as a starting point to understand commonalities and differences. Details about the challenge problems and their mechanization can be found in an electronic appendix which is available at <http://complogic.cs.mcgill.ca/beluga/benchmarks>.

2 Examples

In this section, we give an informal presentation and proofs of various properties of the lambda-calculus. We discuss in detail the first example which is

concerned with equality reasoning and then briefly sketch the other problems. Formal proofs will be discussed in later sections only for the first. All these examples are purposefully simple, so they can be easily understood and one can quickly appreciate the capabilities and trade-offs different systems offer. Yet we believe they are representative of the issues and problems arising when formalizing formal systems and proofs about them.

2.1 Equality reasoning for lambda-terms

We begin by defining the syntax of the (untyped) lambda-calculus together with a declarative definition of equality which includes reflexivity and transitivity in addition to the structural rules. We then define the algorithmic version of equality, which concentrates only on the structural rules. We model the declarative definition of equality by the judgment $\Psi \vdash \text{equal } M N$ and the algorithmic one by the judgment $\Phi \vdash \text{eq } M N$ and carefully define the contexts Ψ and Φ . The goal is to prove these two versions of equality to be equivalent.

$$\begin{array}{ll} \text{Term } M ::= y \mid \lambda x. M \mid \text{app } M_1 M_2 & \text{Context } \Phi ::= \cdot \mid \Phi, \text{equal } x x \\ & \text{Context } \Psi ::= \cdot \mid \Psi, \text{eq } x x \end{array}$$

Algorithmic Equality

$$\frac{\text{eq } x x \in \Psi}{\Psi \vdash \text{eq } x x} \quad \frac{\Psi, \text{eq } x x \vdash \text{eq } M N}{\Psi \vdash \text{eq } (\lambda x. M) (\lambda x. N)} \quad \frac{\Psi \vdash \text{eq } M_1 N_1 \quad \Psi \vdash \text{eq } M_2 N_2}{\Psi \vdash \text{eq } (\text{app } M_1 M_2) (\text{app } N_1 N_2)}$$

Declarative Equality

$$\frac{\text{equal } x x \in \Phi}{\Phi \vdash \text{equal } x x} \quad \frac{\Phi, \text{equal } x x \vdash \text{equal } M N}{\Phi \vdash \text{equal } (\lambda x. M) (\lambda x. N)} \quad \frac{}{\Phi \vdash \text{equal } M M}$$

$$\frac{\Phi \vdash \text{equal } M_1 N_1 \quad \Phi \vdash \text{equal } M_2 N_2 \quad \Phi \vdash \text{equal } M L \quad \Phi \vdash \text{equal } L N}{\Phi \vdash \text{equal } (\text{app } M_1 M_2) (\text{app } N_1 N_2) \quad \Phi \vdash \text{equal } M N}$$

It may be slightly unusual to keep the fact that a variable is equal to itself as a declaration in the context in both formulations. It is only strictly necessary in the first. There are two main reasons. 1) Explicitly introducing the appropriate assumption about each variable is a general methodology which scales to more expressive assumptions. For example, when we specify typing rules, we must introduce a typing context that keeps track of the fact that a given variable has a certain type. 2) Choosing this formulation will also make our proofs more elegant and compact, while at the same time highlight the issues which arise when working with two formal systems each using different assumptions.

We begin by proving that reflexivity and transitivity are indeed admissible from the algorithmic definition of equality.

Theorem 1 (Admissibility of Reflexivity and Transitivity).

1. If Ψ contains premises for all the free variables in M , then $\Psi \vdash \text{eq } M M$.
2. If $\Psi \vdash \text{eq } M L$ and $\Psi \vdash \text{eq } L N$ then $\Psi \vdash \text{eq } M N$.

The first theorem can be proven by induction on M . The second can be proven by induction on the first derivation. We now state that when we have a proof for $\text{equal } M N$ then we also have a proof using algorithmic equality.

Attempt 1 (Completeness). *If $\Phi \vdash \text{equal } M N$ then $\Psi \vdash \text{eq } M N$.*

However, we note that this statement does not contain enough information about how the two contexts Φ and Ψ are related. In the base case, where we have that $\Phi \vdash \text{equal } x x$, we must know that for every variable x in Φ there exists a corresponding assumption such that $\text{eq } x x$ in Ψ . There are two solutions to this problem. 1) We state how two contexts are related and then assume that if this relation holds the theorem holds. 2) We generalize the context used in the theorem such that it contains both assumptions as follows:

Generalized context $\Gamma ::= \cdot \mid \Gamma, \text{eq } x x, \text{equal } x x$

where we deliberately state that the assumption $\text{eq } x x$ always occurs together with the assumption $\text{equal } x x$, and then apply weakening and strengthening as needed to apply the equality inference rules. Both approaches can be mechanized and we discuss some of the trade-offs later. For now we will concentrate on the latter approach and state the revised generalized theorem.

Theorem 2 (Completeness). *If $\Gamma \vdash \text{equal } M N$ then $\Gamma \vdash \text{eq } M N$.*

Proof. Proof by induction on the first derivation. We show three cases which highlight the use of weakening and strengthening.

Case 1: Assumption from context

We know $\Gamma \vdash \text{equal } x x$ where $\text{equal } x x \in \Gamma$ by assumption. Because of the definition of Γ , we know that whenever we have an assumption $\text{equal } x x$, we also must have an assumption $\text{eq } x x$.

Case 2: Reflexivity rule

If the last step applied in the proof was the reflexivity rule $\Gamma \vdash \text{equal } M M$, then we must show that $\Gamma \vdash \text{eq } M M$. By the reflexivity lemma, we know that $\Psi \vdash \text{eq } M M$. By weakening the context Ψ , we obtain the proof for $\Gamma \vdash \text{eq } M M$.

Case 3: Equality rule for lambda-abstractions

| | |
|--|---|
| $\Gamma \vdash \text{equal } (\lambda x. M) (\lambda x. N)$ | by assumption |
| $\Gamma, \text{equal } x x \vdash \text{equal } M N$ | by decl. equality rule for lambda-abstraction |
| $\Gamma, \text{eq } x x, \text{equal } x x \vdash \text{equal } M N$ | by weakening |
| $\Gamma, \text{eq } x x, \text{equal } x x \vdash \text{eq } M N$ | by i.h. |
| $\Gamma, \text{eq } x x \vdash \text{eq } M N$ | by strengthening |
| $\Gamma \vdash \text{eq } (\lambda x. M) (\lambda x. N)$ | by alg. equality rule for lambda-abstraction |

This proof demonstrates many issues related to the treatment of bound variables and the treatment of contexts. First, we need to be able to apply a lemma which was proven in a context Ψ in a different context Γ . Second, we need to apply weakening and strengthening in the proof. Third, we need to be able to know the structure of the context and we need to be able to take advantage of it. We focus here on these structural properties of contexts, but of course many proofs also need the substitution lemma.

2.2 Reasoning about variable occurrences

In this example, we reason about the shape of terms instead of equality of terms. The idea is to compare terms up to variables. For example `lam x. lam y. app x y` would have the same shape as `lam x. lam y. app y x` but these two terms are obviously not equal. We use the judgment $\Phi \vdash \text{shape } M_1 M_2$ to describe that the term M_1 and the term M_2 have the same shape or structure. Thinking of the lambda-terms being described by a syntax tree, comparing the shape of two terms corresponds to comparing two syntax trees where we do not care about specific variable names which are at the leaves of it. The definition for `shape` $M_1 M_2$ can be found in the electronic appendix.

We now prove that if M_1 and M_2 have the same shape, then they must have the same number of variables using the judgment $\Phi \vdash \text{var-occ } M I$ where I describes the total number of variable occurrences in the term M . So for example, the total number of variable occurrences in the term `lam x. lam y. app (app y x) x` is 3. If we think of the lambda-term as a syntax tree, then I describes the number of leaves in the syntax tree described by the term M . We give three different variations, intended to show differences among systems.

Theorem 3.

1. *If $\Phi \vdash \text{shape } M_1 M_2$
then there exists an I such that $\Phi \vdash \text{var-occ } M_1 I$ and $\Phi \vdash \text{var-occ } M_2 I$.
Furthermore I is unique.*
2. *If $\Phi \vdash \text{shape } M_1 M_2$
then for all I . $\Phi \vdash \text{var-occ } M_1 I$ implies $\Phi \vdash \text{var-occ } M_2 I$.*
3. *If $\Phi \vdash \text{shape } M_1 M_2$ and $\Phi \vdash \text{var-occ } M_1 I$ then $\Phi \vdash \text{var-occ } M_2 I$.*

2.3 Reasoning about subterms in lambda-terms

For the next example, we define when a given lambda-term M is a subterm of another lambda-term N and hence we consider M to be structurally smaller than (or equal to) N using the following judgment: $\Psi \vdash \text{le } M N$. Rules for this judgment are given in the electronic appendix. We concentrate here on stating a very simple intuitive theorem that says that if for all terms N , if N is smaller than K implies that N is also smaller than L , then clearly K is smaller than L .

Theorem 4. *If for all N . $\Psi \vdash \text{le } N K$ implies $\Psi \vdash \text{le } N L$ then $\Psi \vdash \text{le } K L$.*

This theorem is interesting because in order to state it, we nest quantification and implications placing them outside the fragment of propositions directly expressible in systems such as Twelf.

3 Mechanization in Twelf and Beluga

In this section, we discuss how the previous examples are implemented in Twelf and Beluga. Both systems share an encoding of expressions and inference rules for declarative and algorithmic equality in the logical framework LF [HHP93]. There are several excellent tutorials on how to represent inference rules in the logical framework LF, and hence we keep this very short.

Formalization of lambda-terms and declarative and algorithmic equality Using HOAS, we represent binders in the object-language (see for example `lam x. M`) using binders in the meta-language, i.e., the logical framework LF. Hence the constructor `lam` takes in a function of type `exp → exp`. For example, the object-language term `lam x. lam y. app x y` will be represented in LF as `lam (λx. lam (λy. app x y))`. Bound variables found in the object language, are not explicitly represented in the meta-language.

| Object-language | Representation in LF |
|-----------------------|---------------------------------------|
| Term $M ::= y$ | <code>exp : type</code> |
| $\lambda x. M$ | <code>lam : (exp → exp) → exp.</code> |
| $\text{app } M_1 M_2$ | <code>app : exp → exp → exp.</code> |

We give the implementation of the declarative and algorithmic equality rules next using the two type families `eq` and `equal` respectively. Each inference rule is then represented as a type. Hypothetical derivations (as in the rule for lambda-abstraction) are represented as higher-order functions.

```

eq: exp → exp → type.
eq_lam : ((Πx : exp. eq x x) → eq (E x) (F x))
          → eq (lam (λx. E x)) (lam (λx. F x)).
eq_app : eq E1 F1 → eq E2 F2 → eq (app E1 E2) (app F1 F2).

equal: exp → exp → type.
e_l: ((Πx:exp. equal x x) → equal (T x) (T' x))
      → equal (lam (λx. T x)) (lam (λx. T' x)).
e_a: equal T2 S2 → equal T1 S1 → equal (app T1 T2) (app S1 S2).
e_r: equal T T.
e_t: equal T R → equal R S → equal T S.

```

Proofs as recursive functions Beluga is a functional language where (hypothetical) derivations are characterized by contextual objects and an inductive proof about derivations is written as a recursive function using pattern matching on them. Each case of the proof corresponds to one branch in the function. First, we define the context schema for the context Ψ which was used in defining algorithmic equality to track assumptions of the form `eq x x` (see page 3). Context schemas classify contexts just as types classify terms. It can be defined as follows: `schema eqCtx = block x:exp . eq x x;` This states that our context consists of blocks of assumptions, containing `x:exp` and `eq x x`. More formally, the `block`-construct introduces a Σ -type grouping the two declarations together.

The reflexivity theorem which stated that for all M there exists a proof for `eq M M` can then be implemented as a recursive function called `ref` which will have the following type: `rec ref : {ψ:(eqCtx)*} {M::exp[ψ]} (eq (M..)(M..))[ψ]`

This can be read as follows: for all contexts ψ which have schema `(eqCtx)*`, for all terms M , we have a proof that `(eq (M..)(M..))[ψ]`. Explicit quantification over the context variable ψ is written using curly brackets in `{ψ:(eqCtx)*}`. The schema is annotated with `*` to denote that declarations of the specified schema may be repeated and the context must be passed explicitly by the user. For universally quantifying over M , we use curly brackets in `{M::exp[ψ]}`. Central to Beluga is the idea of a contextual type. M for example has type `exp[ψ]` which describes an

object m which has type exp in the context ψ . m is hence an expression which may refer to variables in the context ψ . When we use m it is associated with a substitution which maps all the variables in ψ to the correct target context. In the example, we use m within the contextual type $(\text{eq } (M..) (M..))[\psi]$. Hence, M is declared in the context ψ and because it is also used in the context ψ , it is associated with the identity substitution, which is written as.. in our concrete syntax. Intuitively, it means m can depend on all the variables which occur in the context described by ψ . The derivation $\Psi \vdash \text{eq } M \ M$ is directly captured by the contextual type $(\text{eq } (M..) (M..))[\psi]$.

Before we represent the completeness theorem as a recursive function ceq , we define the schema of the generalized context, following our previous informal development as follows: `schema eCtx = block x:exp,u:eq x x.equal x x ;`

Finally, we state the type and implementation of the function ceq . We indicate that the context γ is implicit in the actual implementation of the proof and will be reconstructed by omitting the (...) when declaring the schema of γ .

```

rec ceq: {γ:eCtx} (equal (T..) (S..)) [γ] → (eq (T..) (S..)) [γ] =
fn e ⇒ case e of
| [γ] #p.3.. ⇒ [γ] #p.2..                                % Assumption from context
| [γ] e_r (T..) ⇒ ref [γ] <γ. _ >                      % Reflexivity
| [γ] e_t (D2..) (D1..) ⇒
  let [γ] F2.. = ceq ([γ] D2..) in
  let [γ] F1.. = ceq ([γ] D1..) in
    trans ([γ] F1..) ([γ] F2..)                            % Transitivity
| [γ] e_l (λx. λu. D.. x u) ⇒
  let [γ,b:block x:exp,u:eq x x . equal x x] F.. b.1 b.2 =
    ceq ([γ, b:block x:exp, u:eq x x . equal x x] D.. b.1 b.3)
  in
    [γ] eq_lam (λx. λv. F.. x v)                         % Abstraction
| [γ] e_a (D2..) (D1..) ⇒
  let [γ] F1.. = ceq ([γ] D1..) in
  let [γ] F2.. = ceq ([γ] D2..) in
    [γ] eq_app (F1..) (F2..) ;                               % Application

```

We explain the three cases shown also in the proof on page 4. First, let us consider the case where we used an assumption from the context. It is modelled using parameter variables $#p$ in Beluga. Operationally, $#p$ can be instantiated with any bound variable from the context γ . Since the context γ consists of blocks with the following structure: `block x:exp,u:eq x x . equal x x`, we in fact want to match on the third element of such a block. This is written as $#p.3..$. The type of $#p.3$ is $\text{equal } (\#p.1..) (\#p.1..)$. Since our context always contains a block and the parameter variable $#p..$ describes such a block, we know that there exists a proof for $\text{eq } (\#p.1..) (\#p.1..)$ which can be described by $#p.2..$.

Second, we consider the case where we applied the reflexivity rule e_r as a last step. In this case, we need to refer to the reflexivity lemma we proved about algorithmic equality. To use the function ref which implements the reflexivity lemma for algorithmic equality we however need a context of schema $eCtx$ but the context used in the proof for ceq is of schema $eCtx$. Since the schema $eCtx$ in fact contains at least as much information as the schema $eqCtx$, we should be allowed to pass a context of schema $eCtx$ when a context of schema $eqCtx$ is

required. This is achieved by incorporating context subsumption in Beluga (see [Sch00,HL07] for an introduction to context subsumption).

Third, we consider the case for `e_lam`. In this case, we extend the context with the new declarations about variables and pass to the recursive call `ceq` the derivation $[\gamma, b:\text{block } x:\text{exp}, u:\text{eq } x \cdot x.\text{equal } x \cdot x] D.. b.1 b.3$. Weakening is built-in. Although the derivation D only depends on the context $\psi, x:\text{exp}, u:\text{equal } x \cdot x$, we can use it in the context which also has the assumption $\text{eq } x \cdot x$. Applying the induction hypothesis corresponds to the recursive call. The result of recursive call is a derivation F , where F only depends on $x:\text{exp}$ and $u:\text{eq } x \cdot x$. In the on-paper proof we employed strengthening. Finally, we use F to assemble the final result `eq_lam` ($\lambda x. \lambda v. F .. x v$).

The cases where we applied the application rule `e_a` and the transitivity rule `e_t` as a last step are straightforward. In both cases, we simply appeal to the induction hypothesis on the subderivations $D1..$ and $D2..$. This is implemented as a recursive call to `ceq` using the derivation $[\gamma] D1..$ and the recursive call to `ceq` using the derivation $[\gamma] D2..$. Finally we assemble the result. In the case for applications we use the rule `eq_app` and in the case for transitivity we use the lemma `trans`.

Proofs as relations In Twelf, the proof is implemented as a relation between two derivations, and we separately check that it constitutes a total function. The mode declaration says how we must read the relation operationally. The theorem is represented as a type family, and each case of the proof is represented as one type (or clause). The proof is similar to the implementation in Beluga, with a few exceptions. In Twelf, the context in which we prove the theorem is implicit, and there is no generic variable case, but the variable case is folded into the case for lambda-abstraction. We begin by stating the reflexivity theorem as a relation in Twelf together with the corresponding world declaration. Similar to context schemas, world declarations allow us to describe the context in which the theorem is proven. However, unlike schemas, worlds also keep information about base cases. Since variable cases are handled implicitly, not explicitly, the context must not only list assumptions $x:\text{exp}$ and $u:\text{equal } x \cdot x$ but in addition a proof that reflexivity holds for x , i.e., $\text{ref } x u$.

```
ref:  $\prod T : \text{tp} . \text{equal } T T \rightarrow \text{type}$ . %mode ref +T -D.
%block r_block : block {x:term}{u:equal x x}{r_x: ref x u}.
%worlds (r_block) (ref T D).
```

We now inspect the implementation of the proof of the completeness proof from page 4. It will be very similar to our proof in Beluga, except for the treatment of base cases and contexts.

```
ceq: eq T S → equal T S → type. %mode ceq +E -D.
c_r: ref _ E
     → ceq eq_r E.

c_t: ceq D1 E1 → ceq D2 E2 → tr E1 E2 E
     → ceq (eq_t D2 D1) E.

c_l: ( $\prod x : \text{tm} . \prod u : \text{equal } x x . \prod t_x : \text{tr } u u u . \prod r_x : \text{ref } x u . \prod v : \text{eq } x x$ .
       ceq v u → ceq (E x v) (D x u))
     → ceq (eq_l E) (eq_l D).
```

```

c_a: ceq F1 D1 → ceq F2 D2
  → ceq (eq_a F2 F1) (equ_a D2 D1).
%block cl:block {x:term}{u:equal x x}{t_x:tr u u}{r_x:ref x u}{v: eq x x}
  {c_x: ceq v u}.
%worlds (cl) (ceq E D).
%total E (ceq E D).

```

We can read for example the case `c_a` for applications as follows: Given the relation `ceq F1 D1` (i.h. on the derivation `F1` and `D1`) and the relation `ceq F2 D2` (i.h. on the derivation `F2` and `D2`), we know `ceq (eq_a F2 F1) (equ_a D2 D1)`. This case is closely related to the case in our functional program. The differences arise in the case for lambda-abstractions. Since Twelf supports contexts only implicitly, we must introduce a variable `x` not only together with the assumption `equal x x` and `eq x x`, as we do in Beluga, but we also must assume that the reflexivity and transitivity lemma hold for this variable and that indeed there is a proof that guarantees that whenever we have `equal x x` we must have a proof for `eq x x`.

Because there is no explicit context and no explicit variable case when reasoning about formal systems, the base cases are scattered and pollute our context. Consequently, it now is harder to compose lemmas and reason about the relationship between different contexts. For example, the world described by blocks `r_block` is not a prefix of the world described by blocks `cl`. In Twelf, this will lead to world subsumption failure and the user needs to weaken manually the proof for reflexivity to include assumptions `t_x:trans u u`.³ Apart from the issues around contexts, the Twelf allows a very compact representation of the completeness proof. Weakening and strengthening is handled automatically. For a more detailed explanation regarding the formalization of proofs in the Twelf system and context subsumption, we refer the reader to [HL07].

4 Mechanization in Two-level Hybrid

The Hybrid approach [MMF08] exploits the advantages of HOAS within general theorem proving systems. We use a pretty-printed version of Coq concrete syntax in this paper. *Prop* is the type of meta-level formulas and the usual symbols (e.g., \rightarrow , \forall) represent the meta-level connectives and quantifiers. $\llbracket A_1; A_2; \dots; A_n \rrbracket \rightarrow A$ abbreviates $A_1 \rightarrow (A_2 \rightarrow \dots (A_n \rightarrow A) \dots)$, or equivalently $(A_1 \wedge A_2 \wedge \dots \wedge A_n) \rightarrow A$. The symbol $=$ denotes definitional equality. Free variables in inductive definitions and statements of theorems are implicitly universally quantified at the top-level of each clause or statement.

Hybrid provides a type *expr* and a set of operators on this type used to encode object-language syntax. It is built definitionally on the foundation of the meta-language of the underlying theorem prover; no axioms are introduced. The operators that are used in this paper, with their types are:

$$\text{CON} : \text{con} \rightarrow \text{expr} \mid \text{APP} : \text{expr} \rightarrow \text{expr} \rightarrow \text{expr} \mid \text{LAM} : (\text{expr} \rightarrow \text{expr}) \rightarrow \text{expr}.$$

³ Alternatively, we can also weaken the transitivity lemma and change the order of blocks.

We define the type con later to represent the constants of an object-language.

In the two-level approach used by Hybrid, a specification logic (SL) is defined inductively and used to encode inference rules of object-languages. Hypothetical and parametric judgments are encoded in the SL layer. In this paper, we use a simple SL, a sequent formulation of a fragment of second-order minimal logic with backchaining, adapted from [MM02] (and also used in [MMF08]). Its syntax and inference rules can be encoded directly as follows:

$$\begin{aligned}
\text{inductive } oo &:= \text{tt} : oo \mid \langle \cdot \rangle : atm \rightarrow oo \mid _ \text{and} _ : oo \rightarrow oo \rightarrow oo \\
&\quad \mid _ \text{imp} _ : atm \rightarrow oo \rightarrow oo \mid \text{all} : (expr \rightarrow oo) \rightarrow oo \\
\\
\text{inductive } _ \triangleright _ &: atm \text{ list} \rightarrow nat \rightarrow oo \rightarrow Prop := \\
s.\text{tt} : &\quad \llbracket \Gamma \triangleright_n G \rrbracket \rightarrow \Gamma \triangleright_n \text{tt} \\
s.\text{and} : &\quad \llbracket \Gamma \triangleright_n G_1; \Gamma \triangleright_n G_2 \rrbracket \rightarrow \Gamma \triangleright_{n+1} (G_1 \text{ and } G_2) \\
s.\text{all} : &\quad \llbracket (\forall x. \text{proper } x \rightarrow \Gamma \triangleright_n G x) \rrbracket \rightarrow \Gamma \triangleright_{n+1} (\text{all } x. G x) \\
s.\text{imp} : &\quad \llbracket A, \Gamma \triangleright_n G \rrbracket \rightarrow \Gamma \triangleright_{n+1} (A \text{ imp } G) \\
s.\text{init} : &\quad \llbracket A \in \Gamma \rrbracket \rightarrow \Gamma \triangleright_n \langle A \rangle \\
s.\text{bc} : &\quad \llbracket A \leftarrow G; \Gamma \triangleright_n G \rrbracket \rightarrow \Gamma \triangleright_{n+1} \langle A \rangle
\end{aligned}$$

In the inductive definition of oo , atm is a parameter used to represent atomic predicates of the object-language and $\langle \cdot \rangle$ coerces atoms into propositions of type oo . In the definition of the SL, we use the symbol \triangleright for the sequent arrow and decorate it with natural numbers to allow reasoning by (complete) induction on the *height* of a proof. For convenience we write $\Gamma \triangleright G$ if there exists an n such that $\Gamma \triangleright_n G$, and furthermore we simply write $\triangleright G$ when $\emptyset \triangleright G$. The first four clauses of the definition directly encode the introduction rules of a sequent calculus for this logic. Terms of type $expr$ are built on an underlying de Bruijn syntax. The use of the *proper* annotation rules out terms that have occurrences of bound variables that do not have a corresponding binder (*dangling indices*).⁴ In the last two rules, atoms are provable either by assumption or via *backchaining* over a set of Prolog-like rules, which encode the properties of the object-language. The notation $A \leftarrow G$ represents an instance of one of the clauses of this definition. The sequent calculus is parametric in those clauses.

A small set of structural rules of the SL is proved, and used to reason about object-languages. We prefix theorems formalized in Hybrid with “H-.”

H-Theorem 5 (Structural Properties).

- (a) *Height weakening*: $\llbracket \Gamma \triangleright_n G; n < m \rrbracket \rightarrow \Gamma \triangleright_m G$
- (b) *Context weakening*: $\llbracket \Gamma \triangleright_n G; \Gamma \subseteq \Gamma' \rrbracket \rightarrow \Gamma' \triangleright_n G$
- (c) *Atomic cut*: $\llbracket A, \Gamma \triangleright G; \Gamma \triangleright \langle A \rangle \rrbracket \rightarrow \Gamma \triangleright G$

Formalization of lambda-terms and declarative and algorithmic equality To represent the object-language, we fill in the definition of con , define new operators *app* and *lam* using the operators defined earlier for $expr$, and fill in the definition of atm , which includes the *is_tm* relation for well-formedness of terms as well as *eq* and *equal*. The inference rules are inductively defined using $(_ \leftarrow _)$.

⁴ Hybrid 0.2 described in [MMF08] includes an improvement that doesn’t require the *proper* predicate, but the proofs in this paper are not yet ported to the new version.

```

inductive con := cAPP : con | cLAM : con
app M1 M2 == (APP (APP (CON cAPP) M1) M2)
lam x. M x == (APP (CON cLAM) (LAM (λ x. M x)))

inductive atm := is_tm : expr → atm | eq, equal : expr → expr → atm
inductive _ ← _ : atm → oo → Prop :=
tm_lam : [[ abstr T ]] → is_tm (lam x. Tx) ← all x. (is_tm x) imp ⟨is_tm (Tx)⟩
tm_app : → is_tm (app T1 T2) ← ⟨is_tm T1⟩ and ⟨is_tm T2⟩
eq_lam : [[ abstr E; abstr F ]] → eq (lam x. Ex) (lam x. Fx) ←
all x. (eq x x) imp ⟨eq (Ex) (Fx)⟩
eq_app : → eq (app E1 E2) (app F1 F2) ←
⟨eq E1 F1⟩ and ⟨eq E2 F2⟩
e_l : [[ abstr T; abstr T' ]] → equal (lam x. Tx) (lam x. T'x) ←
all x. (is_tm x) imp (equal x x) imp ⟨equal (Tx) (T'x)⟩
e_a : → equal (app T1 T2) (app S1 S2) ←
⟨equal T1 S1⟩ and ⟨equal T2 S2⟩
e_r : → equal T T ← ⟨is_tm T⟩
e_t : → equal T S ← ⟨equal T R⟩ and ⟨equal R S⟩

```

The well-formedness clauses *tm_lam* and *tm_app* are required since Hybrid terms are untyped (all object-level terms have type *expr*). Each of the remaining clauses of the inductive definition is given the same name as the corresponding rule in the Twelf and Beluga encoding. Note that they are quite similar; the differences in the encodings include 1) the **abstr** conditions used to rule out meta-level functions that do not encode object-level syntax, and (2) the appearance of *is_tm* in the *e_l* and *e_r* clauses, which are required to prove *adequacy* of the encoding (see [FM08] for a fuller discussion of adequacy of Hybrid encodings). In particular, we prove:

$$\begin{aligned} &\triangleright \langle \text{eq } T \ S \rangle \rightarrow \triangleright \langle \text{is_tm } T \rangle \wedge \triangleright \langle \text{is_tm } S \rangle \\ &\triangleright \langle \text{equal } T \ S \rangle \rightarrow \triangleright \langle \text{is_tm } T \rangle \wedge \triangleright \langle \text{is_tm } S \rangle. \end{aligned}$$

Formalization of completeness for algorithmic equality In place of classifying contexts using context schemas or worlds declarations, we adopt the notion of a *context invariant*. This notion is informal; since we have an expressive logic at our disposal, we can define any predicate on contexts. We present one approach and briefly discuss a second one. In the first, we have three context invariants, one each for the proofs of reflexivity, transitivity, and completeness.

$$\begin{aligned} \text{ref_inv } \Phi \ \Psi &= (\forall x. \text{is_tm } x \in \Phi \rightarrow \text{eq } x \ x \in \Psi) \\ \text{tr_inv } \Psi &= (\forall x \ y. \text{eq } x \ y \in \Psi \rightarrow x = y) \\ \text{ceq_inv } \Phi \ \Psi &= \text{ref_inv } \Phi \ \Psi \wedge \text{tr_inv } \Psi \wedge (\forall x \ y. \text{equal } x \ y \in \Phi \rightarrow \text{eq } x \ y \in \Psi) \end{aligned}$$

Context invariants are used for two purposes here: 1) to represent how two contexts in different judgments are related (e.g., *ref_inv*), and 2) to represent information contained in the Σ -type groupings found in the **block** declarations in Beluga and Twelf (e.g., *tr_inv*). If we include enough information, as we do here, no weakening or strengthening is needed in the completeness proof. Instead, the following property is needed.

H-Lemma 6 (Context Extension).

$$\text{ceq_inv } \Phi \Psi \rightarrow \text{ceq_inv} (\text{equal } x \ x, \text{is_tm } x, \Phi) (\text{eq } x \ x, \Psi)$$

We now state the reflexivity and completeness theorems, and discuss the proof of the completeness theorem.

H-Theorem 7 (Reflexivity). $\llbracket \text{ref_inv } \Phi \Psi; \Phi \triangleright_n \langle \text{is_tm } T \rangle \rrbracket \rightarrow \Psi \triangleright_n \langle \text{eq } T \ T \rangle$

In addition to being necessary for adequacy, well-formedness definitions provide a convenient form of induction, which is used to prove the above theorem.

H-Theorem 8 (Completeness).

$$\llbracket \text{ceq_inv } \Phi \Psi; \Phi \triangleright_n \langle \text{equal } T \ S \rangle \rrbracket \rightarrow \Psi \triangleright_n \langle \text{eq } T \ S \rangle$$

The proof is by induction on n with induction hypothesis:

$$IH = \llbracket i < n; \text{ceq_inv } \Phi \Psi; \Phi \triangleright_i \langle \text{equal } T \ S \rangle \rrbracket \rightarrow \Psi \triangleright_i \langle \text{eq } T \ S \rangle.$$

A derivation of $\Phi \triangleright_n \langle \text{equal } T \ S \rangle$ must end in an application of the last two clauses of the definition of the SL (s_init or s_bc , page 10). In the s_init case (the assumption from context case), we know that $(\text{equal } T \ S) \in \Phi$. By the definition of ceq_inv , we know that $(\text{eq } T \ S) \in \Psi$. We use this fact and simply apply s_init to obtain $\Psi \triangleright_n \langle \text{eq } T \ S \rangle$, as desired.

When the derivation ends in s_bc , it must be the case that one of the four clauses defining declarative equality (page 11) was used. We consider reflexivity (e_r) and abstraction (e_l). In the former, we know that $T = S$ and $\Phi \triangleright_{n-1} \langle \text{is_tm } T \rangle$. By H-Theorem 7, we can conclude $\Psi \triangleright_{n-1} \langle \text{eq } T \ T \rangle$ and by H-Theorem 5(a), that $\Psi \triangleright_n \langle \text{eq } T \ T \rangle$.

In the abstraction case (e_l), we know that T and S have the form $(\text{lam } x. T'x)$ and $(\text{lam } x. S'x)$, respectively, and we must show:

$$\begin{aligned} &\llbracket IH; \text{ceq_inv } \Phi \Psi; \Phi \triangleright_n \langle \text{equal } (\text{lam } x. T'x) (\text{lam } x. S'x) \rangle \rrbracket \\ &\quad \rightarrow \Psi \triangleright_n \langle \text{eq } (\text{lam } x. T'x) (\text{lam } x. S'x) \rangle \end{aligned}$$

By repeated inversion of the SL rules on the last premise, and repeated backward application of these rules to the conclusion, we obtain:

$$\begin{aligned} &\llbracket IH; \text{ceq_inv } \Phi \Psi; \text{proper } x; (\text{equal } x \ x, \text{is_tm } x, \Phi) \triangleright_{n-4} \langle \text{equal } (T'x) (S'x) \rangle \rrbracket \\ &\quad \rightarrow (\text{eq } x \ x, \Psi) \triangleright_{n-3} \langle \text{eq } (T'x) (S'x) \rangle \end{aligned}$$

We can conclude $\text{ceq_inv} (\text{equal } x \ x, \text{is_tm } x, \Phi) (\text{eq } x \ x, \Psi)$ by H-Lemma 6 applied to the second premise, and then apply the induction hypothesis to obtain:

$$\begin{aligned} &\llbracket IH; \dots; (\text{eq } x \ x, \Psi) \triangleright_{n-4} \langle \text{eq } (T'x) (S'x) \rangle \rrbracket \\ &\quad \rightarrow (\text{eq } x \ x, \Psi) \triangleright_{n-3} \langle \text{eq } (T'x) (S'x) \rangle \end{aligned}$$

which is provable directly by an application of H-Theorem 5(a).

We can also prove this theorem using a generalized context as is done in Twelf and Beluga. Using this alternate approach, we have only one context, so

the context invariant no longer needs to express relationships between different contexts; now it only needs to contain the following information, which is also found in the `block` declarations in Beluga and Twelf.

$$\text{ceq_inv}' \Gamma = (\forall xy. \text{eq } x y \in \Gamma \rightarrow x = y) \wedge (\forall xy. \text{equal } x y \in \Gamma \rightarrow x = y).$$

Using this approach, we must also explicitly define weakening and strengthening functions on contexts, and lemmas about them. Such functions and lemmas depend on the object-language, but the lemmas are fairly easily proved using H-Theorem 5(b), which is the general weakening theorem of the SL. The reasoning required to prove this new version of H-Theorem 8 is similar to before, though slightly complicated by the need to explicitly apply the weakening and strengthening lemmas.

5 Criteria for Comparison

In this section we compare the approach taken in the three systems considered in this paper. More generally, we describe a list of questions which can be used to quantitatively compare systems and highlight their differences.

How do we represent contexts in proofs? Beluga supports explicit contexts when implementing proofs about LF objects. Context variables allow us to abstract over concrete contexts and the structure of contexts is defined by context schemas. Σ -types tie different declarations together. While Beluga shares the general ideas regarding representing and reasoning about contexts with the Twelf system, it makes the meta-theoretic reasoning about contexts which is hidden in Twelf explicit. In Twelf, the actual context of hypotheses remains implicit. In Hybrid, contexts are explicitly modelled using lists or sets in the SL, but do not appear in the specification of the inference rules of the object-language in the inductive definition of $(_ \leftarrow _)$.

How do we reason with contexts? Reasoning with contexts is particularly important when reasoning about the relationship between two formal systems (such as the equality example) and when we assemble larger proofs using lemmas. Systems like Twelf and Beluga also support structural reasoning about contexts; for example, weakening is supported by the underlying typing rules and context subsumption. This built-in support is sensitive to the ordering of elements in a context (or world) schema and may require explicit weakening as in the implementation of the completeness proof for equality in Twelf. In Hybrid, H-Theorem 5 supports simple reasoning about contexts. This theorem is carried out once and for all at the SL level, and reused for every object-language. More complicated reasoning about weakening and strengthening can be avoided in Hybrid by expressing relationships between contexts in different judgments. The trade-off is that we must define these relationships explicitly as context invariants and prove context extension lemmas. The meta-logic, however, provides considerable flexibility in expressing such invariants. In fact, as discussed earlier, we can express them so that they use generalized contexts and lead to

proofs that follow the corresponding Twelf and Beluga proofs quite closely. Doing so requires explicit weakening and strengthening lemmas that are specific to the object-language. Much of the reasoning that uses these kinds of lemmas is stereotyped and could be automated (although we have not yet done so).

How do we retrieve elements from a context? As the context is implicit in Twelf and the user has no access to it, the user cannot pattern match on elements from it directly. Instead of generic cases which pattern match on an element from the context, base cases in proofs are handled whenever an assumption is introduced, and in fact are treated as part of the context. This may lead to scattering of base cases and redundancy, and in addition complicates reasoning about contexts. In Beluga, retrieval is supported using parameter variables and projections. This is in fact crucial in the reflexivity `ref` and in the completeness proof `ceq`, where we use Σ -types to tie assumptions together and use projections on a parameter variable to retrieve different parts. Since the context is explicit in the SL level in Hybrid, when retrieval of elements is needed in the base case and other cases, it is done via simple list operations such as membership. The Coq libraries provide support for using such operations in proofs.

How easy is it to state a given theorem? The examples in sections 2.2 and 2.3 provide a wide range of statements. All discussed systems provide a two-level approach. However, the level which allows reasoning about formal systems is more expressive in Beluga and in Hybrid. These systems provide direct representations of the given theorems. Twelf’s meta-logic which is used to verify that a given relation constitutes a proof is not rich enough to handle nested quantification and implications directly. One solution is to implement an assertion logic which is then used to encode the actual theorem [SS08].

How do we apply a substitution lemma? In all known systems supporting HOAS encodings substitution lemmas come for “free.” While the examples in this paper do not make use of the substitution lemma, there are several well-known examples such as type preservation for MiniML. In the Twelf system and in Beluga, applying the substitution lemma is reduced to the substitution operation in the underlying logical framework. In Hybrid, the substitution lemma corresponds to the application of the SL cut-rule, expressed as H-Theorem 5(c).

How do we know we have implemented a proof? In a system such as Hybrid, we simply need to trust the underlying proof assistant and establish adequacy. In general, proofs proceed by induction on the definition of the SL with a sub-induction on the object-language. Coq provides extensive support for inductive reasoning, and the induction hypothesis is a premise that is applied explicitly when needed.

In systems such as Twelf or Beluga, we need to establish separately that the user implemented a total function. Twelf is a mature system and provides a coverage checker which in turn relies on the world declarations to ensure the base cases are covered. In addition, the termination checker verifies that a given relation is terminating according to a structural ordering specified by the user. This establishes that all appeals to the induction hypothesis are valid.

Beluga essentially adopts the same philosophy, although the current release does not include a coverage and termination checker. The theoretical foundation for coverage in Beluga is described in [DP09] and an implementation is planned for the future. Intuitively, pattern matching on a contextual object of type $A[\Psi]$ is exhaustive if we cover all constructors of type A plus the cases described by parameter variables, which cover the possibility that we have used an assumption from the context Ψ . For termination checking, we believe the ideas from [Pie05] can be easily adapted.

In general, writing cases using pattern matching by hand may result in a more compact proof since it provides a flexible way to write fall-through patterns or to simultaneously match on several objects. Hence, we may get away with writing fewer cases explicitly as compared to an interactive prover.

How easy is it to interface the system with, for example, support for natural numbers? In the example that counts variable occurrences, reasoning about natural numbers may be necessary and useful. Twelf and Beluga’s reasoning infrastructure does not support them and hence properties like addition and the totality of addition must be proven separately. This leads to some overhead in the actual proofs. Hybrid, on the other hand, relies heavily on the theorem prover’s built in data-type for natural numbers along with a large collection of lemmas and automated proof procedures (such as *omega* in Coq).

6 Conclusion

We presented several benchmark problems together with a general set of criteria for comparing reasoning systems which support the mechanization of formal systems. In addition, we discussed in detail the proofs of one of these problems in three systems (Beluga, Twelf, and Hybrid), and applied these criteria to compare them. This work is a starting point that will help users and developers to evaluate proof assistants which mechanize the reasoning about formal systems. It will also facilitate a better understanding of the differences between and limitations of these systems, as well as the impact of these design decisions in practice. This will provide guidance for users and stimulate discussion among developers.

We hope that these problems will subsequently also be implemented in systems using related approaches. In particular, the Delphin System [PS08] seems to lie between the three systems discussed in this paper. Similarly, it would be interesting to compare systems such as Abella as well as approaches not relying on HOAS encodings such as nominal encodings.

References

- [ABF⁺05] B. Aydemir et. al. Mechanized metatheory for the masses: The POPLmark challenge. In J. Hurd and T. F. Melham, editors, *18th International Conference on Theorem Proving in Higher Order Logics (TPHOLs), Lecture Notes in Computer Science* (LNCS 3603), pages 50–65. Springer, 2005.

- [BSM10] D. Baelde and Z. Snow and D. Miller. Focused inductive theorem proving. In J. Giesl and R. Hähnle, editors, *5th International Joint Conference on Automated Reasoning (IJCAR'10)*, Lecture Notes in Artificial Intelligence (LNAI), forthcoming. Springer, 2010.
- [BC04] Y. Bertot and P. Castéran. *Interactive Theorem Proving and Program Development. Coq'Art: The Calculus of Inductive Constructions*. Springer, 2004.
- [DP09] J. Dunfield and B. Pientka. Case analysis of higher-order data. In *LFMTP'08, Electr. Notes in Theor. Comput. Sci.*, 228:69–84, 2009.
- [FM08] A. P. Felty and A. Momigliano. Hybrid: A definitional two-level approach to reasoning with higher-order abstract syntax. *CoRR*, abs/0811.4367, 2008.
- [Gac08] A. Gacek. The Abella interactive theorem prover (system description). In *4th International Joint Conference on Automated Reasoning, Lecture Notes in Artificial Intelligence* (LNAI 5195), pages 154–161. Springer, 2008.
- [HHP93] R. Harper, F. Honsell, and G. Plotkin. A framework for defining logics. *Journal of the ACM*, 40(1):143–184, 1993.
- [HL07] R. Harper and D. R. Licata. Mechanizing metatheory in a logical framework. *Journal of Functional Programming*, 17(4-5):613–673, 2007.
- [MM02] R. C. McDowell and D. A. Miller. Reasoning with higher-order abstract syntax in a logical framework. *ACM Transactions on Computational Logic*, 3(1):80–136, 2002.
- [MMF08] A. Momigliano, A. J. Martin, and A. P. Felty. Two-level Hybrid: A system for reasoning using higher-order abstract syntax. In *LFMTP'07, Electr. Notes Theor. Comput. Sci.*, 196:85–93, 2008.
- [NPW02] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL: A Proof Assistant for Higher-Order Logic, Lecture Notes in Computer Science* (LNCS 2283). Springer, 2002.
- [Pie05] B. Pientka. Verifying termination and reduction properties about higher-order logic programs. *Journal of Automated Reasoning*, 34(2):179–207, 2005.
- [Pie08] B. Pientka. A type-theoretic foundation for programming with higher-order abstract syntax and first-class substitutions. In *35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'08)*, pages 371–382. ACM Press, 2008.
- [PD10] B. Pientka and J. Dunfield. Beluga:A Framework for Programming and Reasoning with Deductive Systems (System Description). In J. Giesl and R. Hähnle, editors, *5th International Joint Conference on Automated Reasoning (IJCAR'10)*, Lecture Notes in Artificial Intelligence (LNAI), forthcoming. Springer, 2010.
- [PS99] F. Pfenning and C. Schürmann. System description: Twelf — a meta-logical framework for deductive systems. In H. Ganzinger, editor, *16th International Conference on Automated Deduction (CADE-16), Lecture Notes in Artificial Intelligence* (LNAI 1632), pages 202–206. Springer, 1999.
- [PS08] A. B. Poswolsky and C. Schürmann. Practical programming with higher-order encodings and dependent types. In *17th European Symposium on Programming (ESOP'08), Lecture Notes in Computer Science* (LNCS 4960), pages 93–107. Springer, 2008.
- [Sch00] C. Schürmann. *Automating the Meta Theory of Deductive Systems*. PhD thesis, Department of Computer Science, Carnegie Mellon University, 2000. CMU-CS-00-146.
- [SS08] Carsten Schürmann and Jeffrey Sarnat. Structural logical relations. In *23rd Annual Symposium on Logic in Computer Science (LICS)*, pages 69–80. IEEE Computer Society, 2008.