

A Case-Study in Programming Coinductive Proofs: Howe’s Method

David Thibodeau,¹ Alberto Momigliano,² and Brigitte Pientka¹

¹ *School of Computer Science, McGill University, Montreal, Canada*

² *Dipartimento di Informatica, Università degli Studi di Milano, Italy*

Received Received: date / Accepted: date

Bisimulation proofs play a central role in programming languages in establishing rich properties such as contextual equivalence. They are also challenging to mechanize, since they require a combination of inductive and coinductive reasoning on open terms. In this paper we describe mechanizing the property that similarity in the call-by-name lambda calculus is a pre-congruence using Howe’s method in the **Beluga** proof environment. The formal development relies on three key ingredients: 1) we give a higher-order abstract syntax (HOAS) encoding of lambda-terms together with their operational semantics as intrinsically typed terms, thereby avoiding not only the need to deal with binders, renaming and substitutions, but keeping all typing invariants implicit; 2) we take advantage of **Beluga**’s support for representing open terms using first-class contexts and simultaneous substitutions: this allows us to directly state central definitions such as open simulation without resorting to the usual inductive closure operation and to encode very elegantly notoriously painful proofs such as the substitutivity of the Howe relation; 3) we exploit the possibility of reasoning by coinduction in **Beluga**’s reasoning logic. The end result is succinct and elegant, thanks to the high-level abstractions and primitives **Beluga** provides. We substantiate this claim by comparing this proof with a similar one, carried out in *Abella*, and less related developments in *Isabelle/HOL* and *Coq*. We hence believe that this mechanization is a text book example that illustrates **Beluga**’s strength at mechanizing challenging (co)inductive proofs using higher-order abstract syntax encodings.

1. Introduction

Logical frameworks such as *LF* (Harper et al., 1993) and *λProlog* (Miller and Nadathur, 2012) provide a meta-language for representing formal systems given via axioms and inference rules, factoring out common and recurring issues such as modelling variable bindings. They exploit an idea, dating back to Church, where we use a lambda calculus as the meta-language to uniformly model variable binding in our formal system. This technique is now commonly known as higher-order abstract syntax (HOAS) or (in its monomorphic version) “lambda”-tree syntax (Miller and Palamidessi, 1999). In particular, we can model uniformly variable binding by reusing its function space of the meta-language. As a consequence variables in the object language (OL) are represented by variables in

the meta-language and inherit thereby α -renaming and substitution from it. Moreover, this encoding technique scales to representing formal systems that use hypothetical and parametrical reasoning by providing generic support for managing hypotheses and the corresponding substitution lemmas. As users do not need to build up all this basic mathematical infrastructure, it is easier to prototype proof environments and mechanize formal systems. It also can have substantial benefits for proof checking and proof search.

While representing formal systems is a first step, the interesting part is *reasoning* about them, in particular reasoning *inductively* about HOAS representations within a proof environment. HOAS representations cannot be captured with the traditional set-theoretic view of inductive definitions as least fixed points, as they specifically allow and exploit negative occurrences that naturally arise when representing OL bindings.

One solution to this conundrum is the so-called “two-level” approach, as advocated by McDowell and Miller (1997), where we distinguish between a specification language and a reasoning logic above it, which supports at least some form of induction. The cited paper presented $FOLD^N$, which is basically a first order logic with *definitions* (fixed points) and natural number induction. Object logics are encoded in a specification language, which may vary and often is based on (possibly sub-structural) fragments of hereditary Harrop formulas. The method was tested on classical benchmarks such as subject reduction for PCF and its imperative variants.

Noteworthy, one of Dale Miller’s motivating examples has been the (meta)theory of process calculi, in particular the π -calculus. This brought to the forefront the issue of representing and reasoning about *infinite* behaviour. In fact, McDowell et al. (1996) were concerned with the representation of transition systems and their bisimulation: in agreement with Milner’s original presentation in *A Calculus of Communicating Systems*, bisimulation was captured *inductively* by computing the greatest fixed point starting from the universal relation and closing downwards by intersection. This is doable, but notoriously awkward to work with and in fact Milner swiftly adopted the notion of *coinduction* in his subsequent *Communication and Concurrency*.

In the late 1990, coinduction was available in general proof assistants such as Isabelle/HOL and Coq, the first by encoding the standard Tarski’s fixed point theory in higher order logic, the latter by guarded induction; several reasonably large case studies were carried out, not without some difficulties (Ambler and Crole, 1999; Honsell et al., 2001; Hirschhoff, 1997). These case studies further demonstrated the challenges in modelling variable bindings and building up the required infrastructure, as lambda-tree syntax is fundamentally incompatible with the foundations of these proof systems. It turns out instead that it is quite natural to step from $FOLD^N$ to support (co)inductive reasoning; Momigliano and Tiu (2003) simply viewed definitions as least and greatest fixed points adding rules for fixed point induction. With the orthogonal ingredient of ∇ -quantifier to abstract over variable names (Gacek et al., 2008), this line of research culminated in the Abella proof assistant (Abella, 2012), which until recently was, in fact, the only proof assistant which supported natively HOAS and coinduction, as exemplified in some non insignificant case studies (Tiu and Miller, 2010; Momigliano, 2012).

The other main player in HOAS logical frameworks is LF (Harper et al., 1993): Pfen-

ning advocated using it as a *meta*-logical framework by representing inductive proofs as relations. To ensure that a relation describes a valid inductive proof, external checks guarantee that the implemented relation constitutes a total function, i.e. covers all cases and all appeals to the induction hypothesis are well-founded. This led to the proof environment Twelf (Pfenning and Schürmann, 1999), which has been used widely, see for a significant case study (Lee et al., 2007). However, Twelf did not seem to lend itself to coinductive reasoning.

To address these and other shortcomings, Pientka (2008) designed a reasoning logic on top of LF that allows us to directly analyze and manipulate LF objects. **Beluga** (Pientka and Dunfield, 2010) implements this idea. To model derivation trees that depend on assumptions, LF objects are paired with their surrounding context (Nanevski et al., 2008; Pientka, 2008; Pientka and Dunfield, 2008). Inductive proofs are then implemented as *recursive* functions that directly pattern match on contextual LF objects. **Beluga** provides an explicit proof language that makes explicit context reasoning via first-class contexts and first-class simultaneous substitutions together with their equational theory. Moreover, it supports inductive and stratified definitions in addition to higher-order functions (Cave and Pientka, 2012; Pientka and Cave, 2015; Jacob-Rao and Pientka, 2017), thereby going substantially beyond the expressive power of Twelf.

One might say that the proof and the type-theoretic approaches are converging towards a core reasoning logic that supports least and greatest fixed points and equality within first-order logic, as pioneered in Baelde (2011). This might be more obvious in the inductive case where we are more familiar with the computational interpretation of proofs. We readily interpret pattern matching in a program as case analysis in a proof and accept that recursive calls on structurally smaller objects correspond to well-founded appeals to the induction hypothesis. Coinductive reasoning in type theory is less well understood. In fact, guarded corecursion in Coq, for example, does not preserve types (Giménez, 1996; Oury, 2008). To overcome these difficulties, Pientka and collaborators (Abel et al., 2013; Abel and Pientka, 2016; Thibodeau et al., 2016; Jacob-Rao and Pientka, 2017) proposed in prior work a novel computational interpretation of coinductive proofs. While finite (inductive) data is defined using *constructors* and analyzed via *pattern matching*, we define infinite (coinductive) data by *observations* we can make about it. We can reason about such observations using *copattern matching*. A function about finite data represents an inductive proof, if we cover all cases and all recursive calls are on structurally smaller objects. This guarantees that the function is *total*, that is, defined on all inputs and terminating. Dually, a total function about infinite data then corresponds to a coinductive proof, if we cover all possible observations on the output and all recursive calls are guarded by an observation. This guarantees that the function is defined on all possible outputs and remains productive, as we only proceed to evaluate the corecursive function when we apply it to an observation.

As a contribution to a better understanding of the relationship between the logical and computational interpretation of coinductive proofs, the present paper reappraises the proof that similarity in the call-by-name lambda calculus with lists is a pre-congruence using Howe’s method (Howe, 1996). This is a challenging proof since it requires a com-

bination of inductive and coinductive reasoning on open terms. We mechanize this proof in **Beluga**, relying on three key ingredients:

1. we give a HOAS encoding of lambda-terms together with their operational semantics as intrinsically typed terms, thereby avoiding not only the need to deal with binders, renaming and substitutions, but keeping all typing invariants implicit;
2. we take advantage of **Beluga**'s support for representing open terms using first-class contexts and simultaneous substitutions: this allows us to directly state a notion such as open simulation without resorting to the usual inductive closure operation and to encode very elegantly notoriously painful proofs such as the substitutivity of the Howe relation;
3. we exploit the possibility of reasoning by coinduction in **Beluga**'s reasoning logic.

The end result is succinct and elegant, thanks to the high-level abstractions and primitives **Beluga** provides. We substantiate this claim by comparing this proof with the aforementioned proof in **Abella** (Momigliano, 2012), and less related developments in **Isabelle/HOL** (Ambler and Crole, 1999) and **Coq** (Honsell et al., 2001).

The paper starts in Section 2 with a summary description of Howe's method and discusses the challenges it poses to its mechanization. The latter is detailed in Section 3, together with a proof of adequacy of our encoding of similarity (Section 3.4) and an example derivation of two terms being similar (Section 3.10). We review related work in Section 4 and conclude in Section 5. Appendix A contains a brief overview of **Beluga**'s syntax.

2. A Summary of Howe's Method

First let us fix our programming languages as the simply-typed λ -calculus with recursion over (lazy) lists, which we call PCFL following Pitts (1997). Its types consist of the unit type (written as \top), function types, and lists (written as $[\tau]$).

Types	τ	::=	$\top \mid \tau \rightarrow \tau \mid [\tau]$
Terms	m, p, q	::=	$x \mid \mathbf{lam} \ x. p \mid m_1 \ m_2 \mid \mathbf{fix} \ x. m \mid \langle \rangle$ $\mid \mathbf{nil} \mid m_1 :: m_2 \mid \mathbf{lcase} \ m \ \mathbf{of} \ \{\mathbf{nil} \Rightarrow n \mid h :: t \Rightarrow p\}$
Values	v	::=	$\langle \rangle \mid \mathbf{lam} \ x. p \mid \mathbf{nil} \mid m :: n$

The typing rules for PCFL and the big step lazy operational semantics denoted by $m \Downarrow v$ are standard and we omit them here. In particular, lists are only evaluated lazily, as the definition of values shows. The interested reader can skip ahead to their encoding in LF in the Section 3.1 or consult Pitts (1997).

2.1. Proving Bisimilarity a Congruence Using Howe's Method

Suppose we want to say when two programs (two closed terms) have the same behavior. A well known characterization is Morris-style *contextual equivalence*: m and n are equivalent when, if inserted in *any* larger program fragment (context), both larger programs evaluate to the same value, or equivalently both terminate. While this notion of

program equivalence is intuitive, it is indeed difficult to reason about it, mainly due to the quantification on every possible context.[†]

Bisimilarity has emerged as a more manageable, yet, in this setting, equivalent idea. Roughly, m and n are *bisimilar* if whenever m evaluates to a value, so does n , and all the subprograms of the resulting values are also bisimilar, and vice versa. To simplify the presentation, we will concentrate on the notion of *similarity*, from which bisimilarity can be obtained by symmetry, that is taking the conjunction of similarity and its inverse; this is possible thanks to determinism of evaluation.

Definition 1 (Applicative simulation). An *applicative simulation* is a family of typed binary relations R_τ on programs satisfying the following conditions:

- if $m R_\tau n$ then $m \Downarrow \langle \rangle$ entails $n \Downarrow \langle \rangle$.
- if $m R_{[\tau]} n$ then $m \Downarrow \text{nil}$ entails $n \Downarrow \text{nil}$.
- if $m R_{[\tau]} n$ then $m \Downarrow h :: t$ entails that there are h' and t' such that $n \Downarrow h' :: t'$ for which $h R_\tau h'$ and $t R_{[\tau]} t'$.
- if $m R_{\tau \rightarrow \tau'} n$ then $m \Downarrow \text{lam } x. m'$ for any $x:\tau \vdash m':\tau'$ entails that there exists a $y:\tau \vdash n':\tau'$ such that $n \Downarrow \text{lam } y. n'$ and for every $r:\tau$, $m'[r/x] R_{\tau'} n'[r/y]$;

We can make sense of the non-wellfoundedness nature of the last two conditions by noting that the *union* of two applicative simulations is still a simulation and so there exists the *largest* one, which we call applicative *similarity*. This relation can also be characterized using the Knaster-Tarski fixed point theorem, as the greatest fixed point of an appropriate endofunction Φ on families of typed relations. The definition of the function follows the simulation relation, and has, for example at $\tau \rightarrow \tau'$, $m \Phi(R_{\tau \rightarrow \tau'}) n$ just in case whenever $m \Downarrow \text{lam } x. p$ for any p , there exists a n' such that $n \Downarrow \text{lam } y. n'$ and for every $r:\tau$, $p[r/x]$ is $R_{\tau'}$ -related to $n'[r/y]$; hence, similarity is the set *coinductively* defined by Φ , a relation we write as $m \preceq_\tau n$. This yields a co-induction principle that we describe first in its generality and below we show it instantiated to applicative similarity.

$$\frac{\exists S \text{ s.t. } a \in S \quad S \subseteq \Phi(S)}{a \in \text{gfp}(\Phi)} \text{ CI}$$

$$\frac{\exists S_\tau \text{ s.t. } m S_\tau n \quad S_\tau \text{ is an applicative simulation}}{m \preceq_\tau n} \text{ CI-} \preceq$$

It is not difficult to show that similarity is a pre-order and we detail the proof of reflexivity using rule $\text{CI-} \preceq$ to highlight the similarities with the type-theoretic one based instead on the notion of observation, on which our mechanization relies.

Theorem 1 (Reflexivity of applicative similarity). $\forall m \tau, m \preceq_\tau m$.

[†] This notion can and has been simplified, starting from Milner's context lemma (Milner, 1977) and going through the CIU theorem (Mason and Talcott, 1991). Some mechanizations are also available (Ambler and Crole, 1999; Ford and Mason, 2003; McLaughlin et al., 2017), as we discuss further in Section 4.

Proof. To show the result we need to provide an appropriate simulation S and check the simulation conditions. Just choose S_τ to be the family $\{(m, m) \mid \cdot \vdash m : \tau\}$.

We then consider each case in the applicative simulation definition.

- if $m S_\top m$, then $m \Downarrow \langle \rangle$ entails $m \Downarrow \langle \rangle$: immediate;
- if $m S_{[\tau]} m$, then $m \Downarrow \text{nil}$ entails $m \Downarrow \text{nil}$: immediate;
- assume $m S_{[\tau]} m$ and $m \Downarrow h :: t$; pick h', t' to be h, t and by the definition of the simulation, it holds $h S_\tau h$ and $t S_{[\tau]} t$;
- assume $m S_{\tau \rightarrow \tau'} m$ and $m \Downarrow \text{lam } x. m'$; again by picking m' for n' , again by the definition of the simulation it is obvious that for every $r : \tau$, $m'[r/x] S_{\tau'} m'[r/x]$.

□

In many cases, we do not have to look much further than the statement of the theorem to come up with an appropriate simulation, i.e. we can read off the definition of simulation from it — and this is indeed the case for all the coinductive proofs in the following development. However, to show the equivalence of specific programs we may have to come up with a complex bisimulation, possibly defined inductively and/or “up to”. This phenomenon is well-known in inductive theorem proving, where sometimes the induction hypothesis coincides with the statement of the theorem, but in other cases it needs to be generalized in an appropriate lemma. The fixed point rules conflate those two aspects, generalization and lemma application, in one go. With an abuse of language, we will say that we prove a statement by coinduction and say that we appeal to the use of the “coinductive hypothesis” when the simulation corresponds to the statement of the theorem.

When dealing with program equivalence, *equational* in addition to coinductive reasoning would be helpful and this is why it is crucial to establish bisimilarity to be a *congruence*, i.e. a relation respecting the way terms are constructed. Since in this paper we restrict ourselves to similarity, we target pre-congruence. Given the presence of variable-binding operators, we need to consider relations over *open* terms, that is families of binary relations of terms indexed by a typing context Γ in addition to a type τ , which we write as $\Gamma \vdash m \mathcal{R}_\tau n$.

Definition 2 (Compatible Relation). A relation $\Gamma \vdash m \mathcal{R}_\tau n$ is *compatible* when:

- (C0) $\Gamma \vdash \langle \rangle \mathcal{R}_\top \langle \rangle$;
- (C1) $\Gamma, x : \tau \vdash x \mathcal{R}_\tau x$;
- (C2) $\Gamma, x : \tau \vdash m \mathcal{R}_{\tau'} n$ entails $\Gamma \vdash (\text{lam } x. m) \mathcal{R}_{\tau \rightarrow \tau'} (\text{lam } x. n)$;
- (C3) $\Gamma \vdash m_1 \mathcal{R}_{\tau \rightarrow \tau'} n_1$ and $\Gamma \vdash m_2 \mathcal{R}_\tau n_2$ entails $\Gamma \vdash (m_1 m_2) \mathcal{R}_{\tau'} (n_1 n_2)$;
- (C4) $\Gamma, x : \tau \vdash m \mathcal{R}_\tau n$ entails $\Gamma \vdash (\text{fix } x. m) \mathcal{R}_\tau (\text{fix } x. n)$;
- (C5) $\Gamma \vdash m_1 \mathcal{R}_\tau n_1$ and $\Gamma \vdash m_2 \mathcal{R}_{[\tau]} n_2$ entails $\Gamma \vdash (m_1 :: m_2) \mathcal{R}_{[\tau]} (n_1 :: n_2)$;
- (C6a) $\Gamma \vdash \text{nil} \mathcal{R}_{[\tau]} \text{nil}$;
- (C6b) $\Gamma \vdash m_1 \mathcal{R}_{[\tau]} m_2$, $\Gamma \vdash n_1 \mathcal{R}_{\tau'} n_2$ and $\Gamma, h : \tau, t : [\tau] \vdash p_1 \mathcal{R}_{\tau'} p_2$ entails $\Gamma \vdash (\text{lcase } m_1 \text{ of } \{\text{nil} \Rightarrow n_1 \mid h :: t \Rightarrow p_1\}) \mathcal{R}_{\tau'} (\text{lcase } m_2 \text{ of } \{\text{nil} \Rightarrow n_2 \mid h :: t \Rightarrow p_2\})$.

Definition 3 (Pre-congruence). A *pre-congruence* is a compatible transitive relation.

By the very definition of simulation at arrow type it is clear that a key property for our development is for a relation to be preserved by pairwise substitution:

$$\Gamma, y:\tau \vdash m_1 \mathcal{R}_{\tau'} m_2 \text{ and } \Gamma \vdash n \mathcal{R}_{\tau} n' \text{ entails } \Gamma \vdash [n/y]m_1 \mathcal{R}_{\tau'} [n'/y]m_2.$$

We are generalizing this property here using *simultaneous* substitutions. This streamlines in particular our formal development, see Section 3.6. Fig. 1 gives rules for well-typed simultaneous substitutions $\Gamma' \vdash \sigma : \Gamma$, where σ replaces variables in Γ with terms typable in Γ' ; then, it states when two such substitutions are \mathcal{R} -related:

$$\begin{array}{c} \text{Well-Typed Simultaneous Substitutions: } \Psi \vdash \sigma : \Gamma \\ \hline \Psi \vdash \cdot : \cdot \quad \frac{\Psi \vdash \sigma : \Gamma \quad \Psi \vdash m : \tau}{\Psi \vdash \sigma, m/x : \Gamma, x : \tau} \\ \\ \text{Related Simultaneous Substitutions: } \Gamma' \vdash \sigma_1 \mathcal{R}_{\Gamma} \sigma_2 \\ \hline \Psi \vdash \cdot \mathcal{R} \cdot \quad \frac{\Psi \vdash \sigma_1 \mathcal{R}_{\Gamma} \sigma_2 \quad \Psi \vdash m \mathcal{R}_{\tau} n}{\Psi \vdash (\sigma_1, m/x) \mathcal{R}_{\Gamma, x:\tau} (\sigma_2, n/x)} \end{array}$$

Fig. 1. (Related) simultaneous substitutions

Definition 4 (Substitutive relation). A relation is *substitutive* (Sub) iff $\Gamma \vdash m_1 \mathcal{R}_{\tau} m_2$ and $\Gamma' \vdash \sigma_1 \mathcal{R}_{\Gamma} \sigma_2$ entails $\Gamma' \vdash [\sigma_1]m_1 \mathcal{R}_{\tau} [\sigma_2]m_2$.

Some other properties are admissible:

Lemma 2 (Elementary Admissible Properties).

(Ref) If a relation is compatible, then it is *reflexive*;

(Wkn) If $\Gamma \vdash m \mathcal{R}_{\tau} n$ and $\Gamma \subseteq \Gamma'$ then $\Gamma' \vdash m \mathcal{R}_{\tau} n$;

(Cus) If \mathcal{R}_{τ} is substitutive and reflexive, then it is also *closed under substitution*:

$$\Gamma \vdash m_1 \mathcal{R}_{\tau} m_2 \text{ and } \Psi \vdash \sigma : \Gamma \text{ entails } \Psi \vdash [\sigma]m_1 \mathcal{R}_{\tau} [\sigma]m_2.$$

The definition of similarity applies only to closed terms. It is therefore customary to *extend* similarity to *open terms* via instantiation. We do this using *grounding* substitutions:

Definition 5 (Open similarity). $\Gamma \vdash m \preceq_{\tau}^{\circ} m'$ iff $[\sigma]m \preceq_{\tau} [\sigma]m'$, for any $\cdot \vdash \sigma : \Gamma$.

Now, it is immediate that open similarity is a pre-order and hence (C1) and transitivity hold. Further, (C2) also holds, since similarity satisfies

$$\text{lam } x. m \preceq_{\tau \rightarrow \tau'} \text{lam } x. n \text{ iff for all } p:\tau, [p/x]m \preceq_{\tau'} [p/x]n$$

However, a direct attempt to prove pre-congruence of open similarity breaks down when dealing with (C3) and proving that open similarity is substitutive.

Howe's idea (Howe, 1996) was to introduce a *candidate* relation $\preceq_{\tau}^{\mathcal{H}}$ (see Fig. 2), which — contains (open) similarity,

$$\begin{array}{c}
\frac{\Gamma \vdash \langle \rangle \preceq_{\top}^{\circ} n}{\Gamma \vdash \langle \rangle \preceq_{\top}^{\mathcal{H}} n} ep \quad \frac{\Gamma, x:\tau \vdash x \preceq_{\tau}^{\circ} n}{\Gamma, x:\tau \vdash x \preceq_{\tau}^{\mathcal{H}} n} var \quad \frac{\Gamma, x:\tau \vdash m \preceq_{\tau'}^{\mathcal{H}} m' \quad \Gamma \vdash \text{lam } x. m' \preceq_{\tau \rightarrow \tau'}^{\circ} n}{\Gamma \vdash \text{lam } x. m \preceq_{\tau \rightarrow \tau'}^{\mathcal{H}} n} fun \\
\\
\frac{\Gamma, x:\tau \vdash m \preceq_{\tau}^{\mathcal{H}} m' \quad \Gamma \vdash \text{fix } x. m' \preceq_{\tau}^{\circ} n}{\Gamma \vdash \text{fix } x. m \preceq_{\tau}^{\mathcal{H}} n} fix \\
\\
\frac{\Gamma \vdash m_1 \preceq_{\tau \rightarrow \tau'}^{\mathcal{H}} m'_1 \quad \Gamma \vdash m_2 \preceq_{\tau}^{\mathcal{H}} m'_2 \quad \Gamma \vdash m'_1 m'_2 \preceq_{\tau'}^{\circ} n}{\Gamma \vdash m_1 m_2 \preceq_{\tau'}^{\mathcal{H}} n} app \\
\\
\frac{\Gamma \vdash \text{nil} \preceq_{\top}^{\circ} n}{\Gamma \vdash \text{nil} \preceq_{\top}^{\mathcal{H}} n} nil \quad \frac{\Gamma \vdash m_1 \preceq_{\tau}^{\mathcal{H}} m'_1 \quad \Gamma \vdash m_2 \preceq_{[\tau]}^{\mathcal{H}} m'_2 \quad \Gamma \vdash m'_1 :: m'_2 \preceq_{[\tau]}^{\circ} n}{\Gamma \vdash m_1 :: m_2 \preceq_{[\tau]}^{\mathcal{H}} n} cons \\
\\
\frac{\Gamma \vdash m \preceq_{[\tau]}^{\mathcal{H}} m' \quad \Gamma \vdash m_1 \preceq_{\tau'}^{\mathcal{H}} m'_1 \quad \Gamma, h:\tau, t:[\tau] \vdash m_2 \preceq_{\tau'}^{\mathcal{H}} m'_2 \quad \Gamma \vdash \text{lcase } m \text{ of } \{\text{nil} \Rightarrow m'_1 \mid h :: t \Rightarrow m'_2\} \preceq_{\tau'}^{\circ} n}{\Gamma \vdash \text{lcase } m \text{ of } \{\text{nil} \Rightarrow m_1 \mid h :: t \Rightarrow m_2\} \preceq_{\tau'}^{\mathcal{H}} n} lcase
\end{array}$$

Fig. 2. Definition of the Howe relation

— can be shown to be almost a substitutive pre-congruence, where the “almost” refers to being semi-transitive,

and then to prove that it does coincide with similarity.

The informal proof consists of several lemmata:

- 1 Semi-transitivity: the composition of the Howe relation with open similarity is contained in the former. The proof goes by case analysis using transitivity of open similarity.
- 2 The Howe relation is reflexive. Induction on typing, using reflexivity of open similarity.
- 3 Compatibility: (C0)–(C6) hold, an easy consequence of (2).
- 4 Open similarity is contained in Howe, which follows immediately from (1) and (2).
- 5 The Howe relation is substitutive, see Lemma 6.
- 6 The Howe relation “mimics” the simulation conditions:
 - If $\langle \rangle \preceq_{\top}^{\mathcal{H}} n$, then $n \Downarrow \langle \rangle$.
 - If $\text{nil} \preceq_{[\tau]}^{\mathcal{H}} n$, then $n \Downarrow \text{nil}$.
 - If $\lambda x. m \preceq_{\tau \rightarrow \tau'}^{\mathcal{H}} n$, then $n \Downarrow \lambda x. m'$ and for every $q:\tau$ we have $[q/x]m \preceq_{\tau'}^{\mathcal{H}} [q/x]m'$.
 - If $m :: m' \preceq_{[\tau]}^{\mathcal{H}} n$, then $n \Downarrow p :: p'$, with $m \preceq_{\tau}^{\mathcal{H}} p$ and $p \preceq_{[\tau]}^{\mathcal{H}} p'$.

By inversion on the Howe relation and definition of similarity, using semi-transitivity and, in the lambda-case, substitutivity of the Howe relation.

- 7 Downward closure: if $p \preceq_{\tau}^{\mathcal{H}} q$ and $p \Downarrow v$, then $v \preceq_{\tau}^{\mathcal{H}} q$. Induction on evaluation, and inversion on Howe and similarity, with an additional case analysis on v .
- 8 $p \preceq_{\tau}^{\mathcal{H}} q$ entails $p \preceq_{\tau} q$. By coinduction, using the coinductive hypothesis, point (6) and (7).

Once all of these properties have been proved, we are ready for the main result:

Theorem 3. $\Gamma \vdash p \preceq_{\tau}^{\mathcal{H}} q$ iff $\Gamma \vdash p \preceq_{\tau}^{\circ} q$

Equational Theory of Simultaneous Substitution	
$[\sigma]x$	$= \sigma(x)$
$[\sigma](\text{lam } x. m)$	$= \text{lam } x. [\sigma, x/x]m$
$[\sigma](m \ n)$	$= [\sigma]m \ [\sigma]n$
$(\langle \rangle)[\sigma]$	$= \langle \rangle$
$[\sigma](\text{nil})$	$= \text{nil}$
$[\sigma](m :: n)$	$= [\sigma]m :: [\sigma]n$
$[\sigma](\text{fix } x. m)$	$= \text{fix } x. [\sigma, x/x]m$
$[\sigma](\text{lcase } m \text{ of } \{ \text{nil} \Rightarrow n \mid h :: t \Rightarrow p \})$	$= \text{lcase } [\sigma]m \text{ of } \{ \text{nil} \Rightarrow [\sigma]n \mid h :: t \Rightarrow [\sigma, h/h, t/t]p \}$
$[\sigma_2](\cdot)$	$= \cdot$
$[\sigma_2](\sigma_1, m/x)$	$= [\sigma_2]\sigma_1, [\sigma_2]m/x$

Lemma 4 (Substitution Lemma and Weakening Property).

- 1 If $\Gamma' \vdash \sigma : \Gamma$ and $\Gamma \vdash m : \tau$ then $\Gamma' \vdash [\sigma]m : \tau$.
- 2 If $\Gamma' \vdash \sigma : \Gamma$ then $\Gamma', y : \tau \vdash \sigma : \Gamma$.

Lemma 5 (Substitution Properties).

- 1 $[\sigma, n/x]m = [n/x]([\sigma, x/x]m)$
- 2 $[\sigma', n/x]\sigma = [n/x]([\sigma', x/x]\sigma)$
- 3 $[\sigma_2](\sigma_1)m_1 = [[\sigma_2]\sigma_1]m_1$
- 4 $[\sigma_2](\sigma_1)\sigma = [[\sigma_2]\sigma_1]\sigma$
- 5 Let $\text{id} = x_1/x_1, \dots, x_n/x_n$ be the identity substitutions for $\Gamma = x_1:\tau_1, \dots, x_n:\tau_n$, then $[\text{id}]m = m$ and $[\text{id}]\sigma = \sigma$. Moreover, $[\sigma]\text{id} = \sigma$. A special case is when $\Gamma = \cdot$. In this case we have $\text{id} = \cdot$. Moreover, $[\cdot]m = m$, $[\cdot]\sigma = \sigma$, and $[\sigma]\cdot = \cdot$.

Fig. 3. Properties of Simultaneous Substitutions

Proof. Right to left is point (4) above. Conversely, proceed by induction on Γ using (8) for the base case and closure under substitution for the step. \square

Corollary 1. Open similarity is a pre-congruence.

2.2. On the role of substitutions in Howe's method

Substitutions play a central role in the overall proof that similarity is a pre-congruence. In the on paper proof, we silently exploit equational laws about substitution; however they can cause significant pain during mechanization. We summarize the definition of substitution for our term language together with its mostly straightforward equational theory in Fig. 3. To illustrate how we rely on these substitution properties in proofs, we show here in more detail the proof of substitutivity and pay particular attention to the properties in Fig. 3. Recall that the definition of $\Psi \vdash \sigma_1 \lesssim_{\Gamma}^{\mathcal{H}} \sigma_2$ is just an instance of the definition of related simultaneous substitutions.

Lemma 6 (Substitutivity of the Howe relation). Suppose we have $\Gamma \vdash m_1 \lesssim_{\tau}^{\mathcal{H}} m_2$ and $\Psi \vdash \sigma_1 \lesssim_{\Gamma}^{\mathcal{H}} \sigma_2$; then $\Psi \vdash [\sigma_1]m_1 \lesssim_{\tau}^{\mathcal{H}} [\sigma_2]m_2$.

Proof. By induction on the derivations of $\Gamma \vdash m_1 \lesssim_{\tau}^{\mathcal{H}} m_2$.

$$\begin{array}{c}
\text{Case } \frac{\Gamma \vdash m \lesssim_{\tau \rightarrow \tau'}^{\mathcal{H}} m' \quad \Gamma \vdash n \lesssim_{\tau}^{\mathcal{H}} n' \quad \Gamma \vdash m' n' \lesssim_{\tau'}^{\circ} r}{\Gamma \vdash m n \lesssim_{\tau'}^{\mathcal{H}} r} \text{ app} \\
\Psi \vdash [\sigma_1]m \lesssim_{\tau \rightarrow \tau'}^{\mathcal{H}} [\sigma_2]m' \quad \text{by IH} \\
\Psi \vdash [\sigma_1]n \lesssim_{\tau}^{\mathcal{H}} [\sigma_2]n' \quad \text{by IH} \\
\Psi \vdash [\sigma_2](m' n') \lesssim_{\tau'}^{\circ} [\sigma_2]r \quad \text{by closure under substitution (Cus)} \\
[\sigma_2](m' n') = [\sigma_2]m' [\sigma_2]n' \quad \text{by subst. prop.} \\
\Psi \vdash [\sigma_1]m [\sigma_1]n \lesssim_{\tau'}^{\mathcal{H}} [\sigma_2]r \quad \text{by def. of Howe relation} \\
\Psi \vdash [\sigma_1](m n) \lesssim_{\tau'}^{\mathcal{H}} [\sigma_2]r \quad \text{by subst. prop.} \\
\\
\text{Case } \frac{\Gamma, x:\tau \vdash m \lesssim_{\tau'}^{\mathcal{H}} m' \quad \Gamma \vdash \text{lam } x. m' \lesssim_{\tau \rightarrow \tau'}^{\circ} r}{\Gamma \vdash \text{lam } x. m \lesssim_{\tau \rightarrow \tau'}^{\mathcal{H}} r} \text{ fun} \\
\Psi \vdash \sigma_1 \lesssim_{\Gamma}^{\mathcal{H}} \sigma_2 \quad \text{by assumption} \\
\Psi, x : \tau \vdash \sigma_1 \lesssim_{\Gamma}^{\mathcal{H}} \sigma_2 \quad \text{by weakening (Lemma 4.2)} \\
\forall \sigma \text{ where } \cdot \vdash \sigma : \Psi, x:\tau \text{ we have } [\sigma]x \lesssim_{\tau} [\sigma]x \quad \text{by reflexivity of similarity} \\
\Psi, x:\tau \vdash x \lesssim_{\tau}^{\circ} x \quad \text{by def. of open similarity} \\
\Psi, x:\tau \vdash x \lesssim_{\tau}^{\mathcal{H}} x \quad \text{by def. of Howe relation} \\
\Psi, x:\tau \vdash \sigma_1, x/x \lesssim_{\Gamma, x:\tau}^{\mathcal{H}} \sigma_2, x/x \quad \text{by def. of Howe relation for substitutions} \\
\Psi, x:\tau \vdash [\sigma_1, x/x]m \lesssim_{\tau'}^{\mathcal{H}} [\sigma_2, x/x]m' \quad \text{by IH} \\
\Psi \vdash [\sigma_2](\text{lam } x. m') \lesssim_{\tau \rightarrow \tau'}^{\circ} [\sigma_2]r \quad \text{by closure under substitution(Cus)} \\
[\sigma_2](\text{lam } x. m') = \text{lam } x. [\sigma_2, x/x]m' \quad \text{by subst. prop.} \\
\Psi \vdash (\text{lam } x. [\sigma_1, x/x]m) \lesssim_{\tau \rightarrow \tau'}^{\mathcal{H}} [\sigma_2]r \quad \text{by def. of Howe relation} \\
[\sigma_1](\text{lam } x. m) = \text{lam } x. [\sigma_1, x/x]m \quad \text{by subst. prop.} \\
\Psi \vdash [\sigma_1](\text{lam } x. m) \lesssim_{\tau \rightarrow \tau'}^{\mathcal{H}} [\sigma_2]r \\
\text{The other cases are analogous.} \quad \square
\end{array}$$

3. Mechanizing Howe's Method in Beluga

We discuss in this Section the proof that similarity in PCFL is a pre-congruence using Howe's method in the proof environment Beluga.

Beluga is a proof environment that supports both specifying formal systems and reasoning about them. To specify formal systems such as PCFL we use the logical framework LF. This allows us to take advantage of higher-order abstract syntax. A key challenge when reasoning about LF objects is that we must consider potentially open objects. In Beluga, this dilemma is resolved by viewing all LF objects together with the context in which they are meaningful (Nanevski et al., 2008) as *contextual* LF objects and by abstracting not only over LF objects but also over contexts. We then view contextual objects and contexts as a particular index domain about which we can reason using a first-order logic with (co)induction principles on our index domain and first-class equality on index objects. Under the Curry-Howard isomorphism this logic corresponds to a functional language with indexed (co)inductive types that supports (co)pattern matching.

Meta-theoretic proofs about formal systems are implemented as (co)recursive functions in `Beluga`. We summarize in Appendix A the source level syntax of `Beluga`, where we concentrate on the parts that are relevant for our development. It is by no means a complete description, but it may serve as a useful high-level introduction to understanding `Beluga` programs. For a more formal introduction to the theoretical foundations, we refer the reader to (Cave and Pientka, 2012; Thibodeau et al., 2016).

3.1. Encoding Syntax in LF

We adopt the usual HOAS-encoding for binding operators in our OL such as OL functions and fixed points, and make essential use of LF's dependent types (see Fig. 4). In particular the type family `term` encodes *intrinsically*-typed terms. This will make our overall mechanization more elegant and compact, as we do not need to reason about well-typed terms separately.

```

LF tp : type =
| top : tp
| arr : tp → tp → tp
| list: tp → tp;

LF term : tp → type =
| app  : term (arr S T) → term S → term T
| lam  : (term S → term T) → term (arr S T)
| fix  : (term T → term T) → term T
| unit : term top
| nil  : term (list T)
| cons : term T → term (list T) → term (list T)
| lcase: term (list S) → term T → (term S → term (list S) → term T)
      → term T;

```

Fig. 4. LF definition of intrinsically typed terms

Variables such as `T` and `S` that are used in declaring the type of the LF constants are abstracted over at the outside and we rely on type reconstruction to infer the type of these variables (Pientka, 2013). These variables are treated as implicit variables and we subsequently omit passing them when forming `term` objects.

3.2. Encoding the Operational Semantics with Indexed Inductive Types

To illustrate how we can use inductive types in `Beluga`, we encode the *value* and *evaluation* judgment as computation-level type families that are indexed by closed well-typed terms. This is demonstrably equivalent to encode the same judgments at the LF level.

How do we enforce that an LF object is closed? This is accomplished by a contextual

type $[\vdash \text{term } T]$, where the context that appears to the left hand side of the turnstile is empty; to improve readability we simply write $[\text{term } T]$. Note that we can embed contextual types into *Beluga* types, but not vice-versa. There is a strict separation between LF definitions that form our index objects and *Beluga* types that talk about LF definitions.

The inductive type family *Value* encodes closed well-typed expressions. Similarly, the inductive type family *Eval* relates two closed expressions of the same type, where the first big-step evaluates to the second (see Fig. 5).

```

inductive Value : [term T] → type =
  | Val-lam   : Value [lam λx.N]
  | Val-unit  : Value [unit]
  | Val-nil   : Value [nil]
  | Val-cons  : Value [cons M1 M2];

inductive Eval : [term T] → [term T] → type =
  | Ev-app    : Eval [M1] [lam λx.N] → Eval [N[M2]] [V]
              → Eval [app M1 M2] [V]
  | Ev-val    : Value [V]
              → Eval [V] [V]
  | Ev-fix    : Eval [M[fix λx.M]] [V]
              → Eval [fix λx.M] [V]
  | Ev-case-nil : Eval [M] [nil] → Eval [M1] [V]
              → Eval [lcase M M1 (λh.λt.M2)] [V]
  | Ev-case-cons: Eval [M] [cons H L] → Eval [E2[H, L]] [V]
              → Eval [lcase M M1 (λh.λt.M2)] [V];

```

Fig. 5. Inductive definition of values and evaluation

Beluga has a sophisticated notion of first-class simultaneous substitution. Consider the rule *Ev-app*, where to build the evaluation derivation for *Eval* [app M1 M2] [V] we have to supply an evaluation derivation for *Eval* [M1] [lam λx.N] and *Eval* [N[M2]] [V] where *N* stands for a term of type *S* that may refer to $x:\text{term } T$. The substitution that in standard LF would be represented as meta-level application $N\ M2$, here consists of the singleton simultaneous substitution $N[M2]$ that keeps its domain, namely x , implicit. In general, all capitalized variables denote LF objects that may depend on LF declarations. Given an LF term *N* that depends on a context γ , we can use *N* in a context ψ by associating *N* with a simultaneous substitution σ with domain ψ and codomain γ , with type $[\psi \vdash \gamma]$. This *closure* is written in post-fix notation as $N[\sigma]$. If ψ is equivalent to γ , then σ becomes the identity substitution which we may drop. Often we need to weaken an LF object that is closed, in order to use it in a context γ . For example, the type *T* is closed in $[\text{term } T]$. To use *T* in a context γ , we need to associate it with a *weakening substitution* which is simply written as $[\]$. Hence, $[\gamma \vdash \text{term } T][\]$ describes a *term* object of type *T* in a context γ . More generally, we can weaken an object *M* that is defined in a context γ

to an object that is defined in an extended context $\gamma, \mathbf{x}:\mathbf{term} \mathbf{T}[]$ by associating \mathbf{M} with a general weakening substitution, written as $[\dots]$. Finally, we note that the η -expanded form of $\mathbf{Eval} \ [\mathbf{M}] \ [\mathbf{lam} \ \lambda x.N]$ is $\mathbf{Eval} \ [\mathbf{M}] \ [\mathbf{lam} \ \lambda x.N[\mathbf{x}]]$. The substitution that maps \mathbf{x} to itself is simply written as $[\mathbf{x}]$.

Inductive types in **Beluga** correspond to least fixed points over an index domain type-theoretically defined using labelled sums and Σ -type (Cave and Pientka, 2012; Thibodeau et al., 2016). Such inductive types must in general satisfy the positivity condition. The surface definition of **Value** is translated in the following notation where sums are represented as list of labels together with their computation-level types wrapped with $\langle \ \rangle$.

$\mu\mathbf{Value}.\lambda T, M.$

$\langle \mathbf{Val-lam} : \Sigma S_1, S_2:[\mathbf{tp}].T = (\mathbf{arr} \ S_1 \ S_2) \times \Sigma N:[x:\mathbf{term} \ S_1 \vdash \ \mathbf{term} \ S_2[]].M = \mathbf{lam} \ \lambda x.N,$

$\mathbf{Val-unit} : T = \mathbf{top} \times M = \mathbf{unit},$

$\mathbf{Val-nil} : \Sigma S:[\mathbf{tp}].T = \mathbf{list} \ S \times M = \mathbf{nil}$

$\mathbf{Val-cons} : \Sigma S:[\mathbf{tp}].T = \mathbf{list} \ S \times \Sigma M_1:[\mathbf{term} \ S].\Sigma M_2:[\mathbf{term} \ (\mathbf{list} \ S)].M = \mathbf{cons} \ M_1 \ M_2 \ \rangle$

We view **Value** still as a least fixed point definition, although there is no recursive reference to **Value**. The latter happens in the fragment of the inductive type for **Eval**:

$\mu\mathbf{Eval}.\lambda T, M, W.$

$\langle \mathbf{Ev-app} : \Sigma S:[\mathbf{tp}], M_1:[\mathbf{term} \ (\mathbf{arr} \ S \ T)], M_2:[\mathbf{term} \ S], N:[x:\mathbf{term} \ T \vdash \ \mathbf{term} \ S[]].$

$M = (\mathbf{app} \ M_1 \ M_2) \times \mathbf{Eval} \ [M_1] \ [\mathbf{lam} \ \lambda x.N] \times \mathbf{Eval} \ [N[M_2]] \ [W]$

$\mathbf{E-fix} : \Sigma N:[x:\mathbf{term} \ T \vdash \ \mathbf{term} \ T[]].M = \mathbf{fix} \ \lambda x.N \times \mathbf{Eval} \ [N[\mathbf{fix} \ \lambda x.N]] \ [W]$

$\mathbf{E-val} : M = W \times \mathbf{Value} \ [W] \ \dots \ \rangle$

3.3. Encoding Similarity Using Indexed Coinductive Types

In **Beluga**, we also can state coinductive type families and in particular similarity as a coinductive definition that relates closed well-typed terms.

While inductive types are defined by constructors, we define *coinductive* types by the *observations* we can make (Abel et al., 2013; Thibodeau et al., 2016). To define the coinductive type $\mathbf{Sim} \ [\mathbf{T}] \ [\mathbf{M}] \ [\mathbf{N}]$, we declare observations $\mathbf{Sim_unit}$, $\mathbf{Sim_nil}$, $\mathbf{Sim_cons}$, and $\mathbf{Sim_lam}$; each one corresponds to a case in our definition of applicative simulation — compare Def. 1 to Fig. 6.

When we define an indexed inductive type, the indices impose obligations that must be satisfied in order to construct an object. When we define a coinductive one, indices guard what observations we can make. If the guard is true, then we can make the observation and proceed. We write the observation together with its type on the left side of $::$ and on the right side we give the result type of the observation that describes our proof obligation. For example, we can make the observation $\mathbf{Sim_unit}:\mathbf{Sim} \ [\mathbf{top}] \ [\mathbf{M}] \ [\mathbf{N}]$, if we can show that $\mathbf{Eval} \ [\mathbf{M}] \ [\mathbf{unit}] \rightarrow \mathbf{Eval} \ [\mathbf{N}] \ [\mathbf{unit}]$. It corresponds directly to “ $m \Downarrow \langle \rangle$ entails $n \Downarrow \langle \rangle$ ” in Def. 1. Note that \mathbf{M} and \mathbf{N} are implicitly quantified at the outside. The definition of the observation $\mathbf{Sim_nil}$ follows a similar schema.

The result of the observation $\mathbf{Sim_cons}$ on $\mathbf{Sim} \ [\mathbf{list} \ \mathbf{T}] \ [\mathbf{M}] \ [\mathbf{N}]$ requires that if $m \Downarrow h :: t$ then there are h' and t' such that $n \Downarrow h' :: t'$ for which $h R_\tau h'$ and $t R_{[\tau]} t'$. We hence need a way to encode an *existential* property. Although existentials (i.e. Σ -types) exist

```

coinductive Sim : {T:[tp]} [term T] → [term T] → type =
| (Sim_unit : Sim [top] [M] [N]) ::
    Eval [M] [unit] → Eval [N] [unit]
| (Sim_nil : Sim [list T] [M] [N]) ::
    Eval [M] [nil] → Eval [N] [nil]
| (Sim_cons : Sim [list T] [M] [N]) ::
    Eval [M] [cons H L] → ExSimCons [H] [L] [N]
| (Sim_lam : Sim [arr S T] [M] [N]) ::
    Eval [M] [lam λx.M'] → ExSimLam [x:term S ⊢ M'] [N]

and inductive ExSimCons:[term T]→[term (list T)]→[term (list T)]→ type =
| ExSimcons: Eval [N] [cons H' L']
    → Sim [T] [H] [H'] → Sim [list T] [L] [L']
    → ExSimCons [H] [L] [N]

and inductive ExSimLam: [x:term S ⊢ term T[]] → [term (arr S T)] → type =
| ExSimlam: Eval [N] [lam λx.N']
    → ({R:[term S]} Sim [T] [ N' [R] ] [ N' [R] ])
    → ExSimLam [x:term S ⊢ N'] [N]

```

Fig. 6. Coinductive definition of applicative similarity

in our theoretical foundation, the implementation of Beluga does not support them at the top level, as they always can be realized using indexed inductive types. We therefore define an indexed inductive type `ExSimCons` that relates h , t and n .

Last, we need to represent the result of observing `Sim_lam` that encodes the corresponding part from the definition:

$m \Downarrow \text{lam } x. m'$ for any $x:\sigma \vdash m':\tau$ entails
that there exists a $y:\sigma \vdash n':\tau$ such that $n \Downarrow \text{lam } y. n'$ and for every $r:\sigma$, $m'[r/x] R_\tau n'[r/y]$

We again resort to defining an inductive type `ExSimLam` that relates the term `P` with type `[x:term S ⊢ term T[]]`, i.e. `M'` has type `term T[]` under the assumption of the variable `x` having type `term S`. Hence we can simply write `[x:term S ⊢ M']`, as we interpret `M'` within the context `x:term S`. As T denotes a closed type, we associate it with a weakening substitution, when it is used in a non-empty context. The relation `ExSimLam` exists if `Eval [N] [lam λx.N']` and for all $R:[\text{term } S]$ we know `Sim [T] [M' [R]] [N' [R]]`. Finally, we remark that the coinductive type `Sim` and inductive types `ExSimCons` and `ExSimLam` are defined mutually.

In the type-theoretic foundation that underlies Beluga, the coinductive type family `Sim` is encoded using a greatest fixed point that is defined using records, universals (written using Π), and implications.

```

νSim.λT.λM.λN.
{ S.unit : T = top → Eval [M] [unit] → Eval [N] [unit]
  S.lam  : ΠS1:[tp].ΠS2:[tp].ΠM':[x:term S1 ⊢ term S2]]. T = arr S1 S2
    → Eval [M] [lam λx.M']
    → ΣN':[x:term S1 ⊢ term S2]]. Eval [N] [lam λx.N'] × ΠR:[term S1].Sim [S2] [M'[R]] [N'[R]]
}

```

We only show the encoding for lambda-expressions and omit the observations we can make on lists to keep it readable. In this internal representation the guards such as $T = \text{top}$ or $T = \text{arr } S_1 S_2$ are made explicit. We further inlined the definition of `ExSimlam` to keep the definition compact.

3.4. Adequacy of Coinductive Encoding

We next sketch the adequacy of the encoding of similarity; a full proof, such as those in the electronic appendix of Tiu and Miller (2010) would fill a dozen pages and require to spell out the static and dynamic semantics of (co)inductive Beluga (Thibodeau et al., 2016). Instead, we rely on our intuitive understanding of (co)inductive types; to get started, we assume the adequacy of LF encodings, whereby we denote the mapping of terms m and types τ to their encodings as $\ulcorner m \urcorner$ and $\ulcorner \tau \urcorner$ respectively. Conversely, the decoding of an LF object \mathbb{M} and \mathbb{T} into terms and types is written $\llcorner \mathbb{M} \llcorner$ and $\llcorner \mathbb{T} \llcorner$ respectively.

Lemma 7. For any term m and type τ , we have $\llcorner \ulcorner m \urcorner \llcorner = m$ and $\llcorner \ulcorner \tau \urcorner \llcorner = \tau$.

Proof. Standard, following for example Pfenning (1997). □

We further build on the adequacy of the encoding of substitutions. In particular, the translation of $[\sigma]m$ is equivalent to first translating σ and the term m to their corresponding representations in LF and then relying on the built-in simultaneous LF substitution operation of applying $\ulcorner \sigma \urcorner$ to $\ulcorner m \urcorner$. The encoding $\ulcorner \sigma \urcorner$ is defined inductively on the substitution σ as expected: $\ulcorner \cdot \urcorner = \hat{}$ and $\ulcorner \sigma, m/x \urcorner = \ulcorner \sigma \urcorner, \ulcorner m \urcorner$. Further, recall that we write the application of a simultaneous LF substitution in prefix form, while we write the closure of an LF object together with an LF substitution in post-fix.

Lemma 8 (Compositionality). $\ulcorner [\sigma]m \urcorner = \ulcorner \sigma \urcorner \ulcorner m \urcorner$.

Proof. Generalization of the compositionality lemma for LF. □

Lemma 9 (Soundness). If $\text{Sim } \ulcorner \tau \urcorner \ulcorner m \urcorner \ulcorner n \urcorner$ then $m \preceq_\tau n$.

Proof. (Sketch) We apply rule $CI-\preceq$ to unfold $m \preceq_\tau n$ selecting the family S_τ to be

$$\{(m, n) \mid \text{Sim } \ulcorner \tau \urcorner \ulcorner m \urcorner \ulcorner n \urcorner\}$$

We then show that S_τ satisfies the simulation conditions unfolding the definition of S_τ . □

Before addressing the other direction, we briefly contrast the more familiar inductive reasoning with the coinductive reasoning we will use. To prove a conjecture *inductively*

on an object, we consider all possible ways such an object can be constructed and we reason inductively about some notion of *size* of an object or a derivation. In an inductive proof, we may assume that the conjecture holds for objects of size k and show that the conjecture holds for objects of size m where $k < m$.

For example, to prove that for all term M and V , if $\mathcal{D} : \text{Eval } [M] [V]$ then $\mathcal{F} : \text{Value } [V]$, we proceed by induction on the height n of \mathcal{D} , the derivation of $\text{Eval } [M] [V]$. We therefore prove the following by considering all possible constructors that we can use to build such a derivation \mathcal{D} .

IH For all $k < n$, for all term M and V ,
 if $\mathcal{D} : \text{Eval } [M] [V]$ and $\text{size}(\mathcal{D}) = k$ then $\mathcal{F} : \text{Value } [V]$

To show For all term M and V , if $\mathcal{D} : \text{Eval } [M] [V]$ and $\text{size}(\mathcal{D}) = n$ then $\mathcal{F} : \text{Value } [V]$

Dually, to prove a conjecture *coinductively*, we consider all possible observations we can make about an object and we *reason inductively on the number of observations*, which we refer to as “depth”. In a coinductive proof, we assume that the conjecture holds when we can make k observations about the object, and we show that the conjecture also holds when we make n observations about it where $k < n$. In essence, to prove a statement by coinduction we reason by complete induction on the number of observations. For example, if we want to prove reflexivity of simulation, i.e. for all terms M and types T , $\mathcal{D} : \text{Sim } [T] [M] [M]$, then we proceed by induction on the number of observation on \mathcal{D} and consider all possible observations we can make about \mathcal{D} .

IH For all $k < n$, for all term M and types T ,
 $\mathcal{D} : \text{Sim } [T] [M] [M]$ and $\text{depth}(\mathcal{D}) = k$

To show For all term M and types T , $\mathcal{D} : \text{Sim } [T] [M] [M]$ and $\text{depth}(\mathcal{D}) = n$

We are now ready to address the other direction of the adequacy statement. For a more formal justification of reasoning about inductive data via sizes and coinductive data via observations we refer the reader to (Abel and Pientka, 2016, 2013).

Lemma 10 (Completeness). If $m \preceq_\tau n$, then $\text{Sim } [\ulcorner \tau \urcorner] [\ulcorner m \urcorner] [\ulcorner n \urcorner]$.

Proof. We proceed by complete induction on the number of observations we can make on $\text{Sim } [\ulcorner \tau \urcorner] [\ulcorner m \urcorner] [\ulcorner n \urcorner]$.

IH For all $k < j$, for all term m and types τ ,
 If $\mathcal{S} : m \preceq_\tau n$, then $\mathcal{D} : \text{Sim } [\ulcorner \tau \urcorner] [\ulcorner m \urcorner] [\ulcorner n \urcorner]$ and $\text{depth}(\mathcal{D}) = k$

To show for all term m and types τ ,
 If $\mathcal{S} : m \preceq_\tau n$, then $\mathcal{D} : \text{Sim } [\ulcorner \tau \urcorner] [\ulcorner m \urcorner] [\ulcorner n \urcorner]$ and $\text{depth}(\mathcal{D}) = j$

Observation **S_unit**.

To show: If $m \preceq_\tau n$ then $\ulcorner \tau \urcorner = \text{top} \rightarrow \text{Eval } [\ulcorner m \urcorner] [\text{unit}] \rightarrow \text{Eval } [\ulcorner n \urcorner] [\text{unit}]$.

Assume $m \preceq_\tau n$, $\ulcorner \tau \urcorner = \text{top}$, and $\text{Eval } [\ulcorner m \urcorner] [\text{unit}]$

$\tau = \top$

by definition of $\ulcorner \tau \urcorner$

$m \Downarrow \langle \rangle$ entails $n \Downarrow \langle \rangle$	by Def. 1 using the assumption $m \preceq_\tau n$
$\perp \text{Eval } [\ulcorner m \urcorner] [\text{unit}] \Downarrow = m \Downarrow \langle \rangle$	by decoding of <code>Eval</code>
$n \Downarrow \langle \rangle$	by previous lines
$\ulcorner n \Downarrow \langle \rangle \urcorner = \text{Eval } [\ulcorner n \urcorner] [\text{unit}]$	by encoding of <code>Eval</code>

Observation `S_lam`.

IH	For all $k < j$, if $m \preceq_\tau n$ then $\mathcal{D} : \text{Sim } [\ulcorner \tau \urcorner] [\ulcorner m \urcorner] [\ulcorner n \urcorner]$ and $\text{depth}(\mathcal{D}) = k$
To show	If $m \preceq_\tau n$ then $\mathcal{D} .\text{S_lam} : \text{Sim } [\ulcorner \tau \urcorner] [\ulcorner m \urcorner] [\ulcorner n \urcorner]$ and $\text{depth}(\mathcal{D} .\text{S_lam}) = j$

Making an observation corresponds to projecting with the dot notation the field `S_lam` of the record. We further note that $\text{depth}(\mathcal{D} .\text{S_lam}) = \text{depth}(\mathcal{D}) + 1$. Since we are making the observation `S_lam`, we can unfold the definitions. Hence, it suffices to show

if $m \preceq_\tau n$ then
for all $\mathbf{S}_1, \mathbf{S}_2, \mathbf{M}'$. $\ulcorner \tau \urcorner = \text{arr } \mathbf{S}_1 \ \mathbf{S}_2 \rightarrow \text{Eval } [\ulcorner m \urcorner] [\text{lam } \lambda x. \mathbf{M}']$
 \rightarrow there exists \mathbf{N}' s.t. $\text{Eval } [\mathbf{N}] [\text{lam } \lambda x. \mathbf{N}']$
and (for all \mathbf{R} . $\mathcal{D} : \text{Sim } [\mathbf{S}_2] [\mathbf{M}'[\mathbf{R}]] [\mathbf{N}'[\mathbf{R}]]$)

Moreover, $\text{depth}(\mathcal{D})$ is clearly less than $\text{depth}(\mathcal{D} .\text{S_lam})$ and we may appeal to the induction hypothesis, which can be specialized to the following statement:

if $[r/x]m' \preceq_{s_2} [r/y]n'$ then $\text{Sim } [\mathbf{S}_2] [\mathbf{M}'[\mathbf{R}]] [\mathbf{N}'[\mathbf{R}]]$ where $\ulcorner m' \urcorner = \mathbf{M}'$, $\ulcorner r \urcorner = \mathbf{R}$, $\ulcorner n' \urcorner = \mathbf{N}'$.

Assume $m \preceq_\tau n$, $\ulcorner \tau \urcorner = (\text{arr } \mathbf{S}_1 \ \mathbf{S}_2)$, and $\text{Eval } [\ulcorner m \urcorner] [\text{lam } \lambda x. \mathbf{M}']$
 $\tau = s_1 \rightarrow s_2$ since $\ulcorner s_1 \rightarrow s_2 \urcorner = \text{arr } \mathbf{S}_1 \ \mathbf{S}_2$ where $s_1 = \perp \mathbf{S}_1 \perp$ and $s_2 = \perp \mathbf{S}_2 \perp$.
for any $x:s_1 \vdash m':s_2$. $m \Downarrow \text{lam } x. m'$ entails that
there exists a $y:s_1 \vdash n':s_2$ such that $n \Downarrow \text{lam } y. n'$
and for every $r:s_1$, $[r/x]m' \preceq_{s_2} [r/y]n'$; by definition of $m \preceq_\tau n$
 $\text{Eval } [\ulcorner m \urcorner] [\text{lam } \lambda x. \mathbf{M}'] = \ulcorner m \Downarrow \perp \text{lam } \lambda x. \mathbf{M}' \perp \urcorner$ by encoding of `Eval`
 $\text{lam } \lambda x. \mathbf{M}' = \ulcorner \text{lam } x. m' \urcorner$ by encoding of terms
there exists a $y:s_1 \vdash n':s_2$ such that $n \Downarrow \text{lam } y. n'$
and for every $r:s_1$, $[r/x]m' \preceq_{s_2} [r/y]n'$ by previous lines
 $\ulcorner n \Downarrow \text{lam } y. n' \urcorner = \text{Eval } \ulcorner n \urcorner (\text{lam } \lambda y. \ulcorner n' \urcorner)$ by encoding of `Eval`

Assume $\mathbf{R}:\text{term } \mathbf{S}_1$.
 $\ulcorner [r/x]m' \urcorner = \mathbf{M}'[\mathbf{R}]$ and $\ulcorner [r/y]n' \urcorner = \ulcorner n' \urcorner[\mathbf{R}]$ by Theorem 8 (Compositionality)
 $\text{Sim } [\mathbf{S}_2] [\mathbf{M}'[\mathbf{R}]] [\ulcorner n' \urcorner[\mathbf{R}]]$ by the specialized induction hypothesis
using $[r/x]m' \preceq_{s_2} [r/y]n'$ from the previous line using $[r/x]m' \preceq_{s_2} [r/y]n'$ from the previous line

Therefore, there exists \mathbf{N}' , namely $\ulcorner n' \urcorner$, and $\text{Eval } [\ulcorner n \urcorner] [\text{lam } \lambda y. \ulcorner n' \urcorner]$ and for all $\mathbf{R}:\text{term } \mathbf{S}_1$, we have $\text{Sim } [\mathbf{S}_2] [\mathbf{M}'[\mathbf{R}]] [\ulcorner n' \urcorner[\mathbf{R}]]$. Hence, $\text{Sim } [\ulcorner s_1 \rightarrow s_2 \urcorner] [\ulcorner m \urcorner] [\ulcorner n \urcorner]$. This concludes this case. \square

Theorem 11 (Adequacy of Encoding of Simulation as Coinductive Type).
 $m \lesssim_{\tau} n$ iff $\text{Sim } [\ulcorner \tau \urcorner] [\ulcorner m \urcorner] [\ulcorner n \urcorner]$.

We remark that while soundness follows the same structure of (Honsell et al., 2001) and (Tiu and Miller, 2010), the possibility to induct on the number of observation allows one to establish completeness in a novel and easier way w.r.t. an analogous result in (Tiu and Miller, 2010), which had to resort to a complex induction on the structure of the arguments of the coinductively defined predicate/type.

Subsequently, we will not make the number of observations explicit in coinductive arguments, but simply allow corecursive calls when they are guarded by an observation.

3.5. Writing Coinductive Proofs Using Copattern Matching

In *Beluga*, we implement (co)inductive proofs as (co)recursive functions using (co)pattern matching (Abel et al., 2013). Let us reconsider first the proof that similarity is reflexive (Fig. 7): for all $T, M, \text{Sim } [T] [M] [M]$. The type of the function `sim_refl` encodes this statement directly. We leave T and M implicit, as these arguments can be reconstructed by *Beluga*.

```

rec sim_refl : Sim [T] [M] [M] =
fun .Sim_unit d => d
  | .Sim_nil d => d
  | .Sim_cons d => ExSimcons d sim_refl sim_refl
  | .Sim_lam d => ExSimlam d (fun [R] => sim_refl)

```

Fig. 7. Corecursive program showing that similarity is reflexive

To prove our statement, we consider each case by writing the observation on the left hand side of our corecursive function and provide a witness of the appropriate type on the right hand side. We write observations as *projections* prefixing them with a dot.

We recall that each observation is implicitly guarded by a constraint (for example `Sim_unit` is guarded by $T = \text{top}$) and we can only make the observation if the constraint is satisfied. *Beluga* reconstructs proofs for such equality guards and associates them implicitly with the observation. If the guard is not satisfied, i.e. no proof that $T = \text{top}$ for example exists, then the case is trivially satisfied, and we omit it from our function definition.

In general, the arguments of a function definition in *Beluga* may consist of index objects, patterns describing inductive objects, and copatterns (i.e. observations) defining coinductive objects. In fact, the function `sim_refl` takes first the two arguments that we left implicit, namely T and M , before it expects each observation.

Observation `Sim_unit` and $T = \text{top}$. In this case we need to construct a witness for `Eval [M] [unit] → Eval [M] [unit]`. This is simply the identity function that maps

$d:\text{Eval } [M] \text{ [unit]}$ to itself. Note that instead of returning a function, we write d on the left hand side of our function definition.

Observation Sim_nil and $T = \text{list } T'$. Similar to the previous case.

Observation Sim_cons and $T = \text{list } T'$. In this case, $d:\text{Eval } [M] \text{ [cons H T]}$ and we need to supply a witness for $\text{ExSimCons } [H] \text{ [L]} \text{ [M]}$: we choose d for $\text{Eval } [M] \text{ [cons H T]}$ and make two corecursive calls on T and H respectively. We note that the appeal to the corecursive calls is justified since observation .Sim_cons about $\text{Sim } [\text{list } T'] \text{ [M]} \text{ [M]}$ increases the number of observations made.

Hence, a corecursive call is only valid if it is guarded by at least an observation on the left hand side of a function definition, as we comment on next.

Observation Sim_lam and $T = (\text{arr } S \ S')$. We have $d:\text{Eval } [M] \text{ [lam } \lambda x.M']$ and we need to provide a witness for $\text{ExSimLam } [x:\text{term } S \vdash M'] \text{ [M]}$: we again choose d for $\text{Eval } [M] \text{ [lam } \lambda x.M']$ and then build a function of type $\{R:[\text{term } S]\} \text{ Sim } [S'] \text{ [M']} [R] \text{ [M']} [R]$ that makes a corecursive call to sim_refl . We again observe that the corecursive call is justified, as it is guarded by the observation .Sim_lam .

As we mentioned above, in order for a term like sim_refl to be considered a proof, it needs to be *covering* and *productive*. To be covering, a term of codata type needs to have a branch for each possible observation. In this case, since sim_refl does not require a particular shape for the type of M and N , we need to provide all the possible observations. However, if we were to prove reflexivity only for terms of type $\text{list } T$, i.e. $\text{Sim } [\text{list } T] \text{ [M]} \text{ [M]}$, then we are justified to only consider the branches for Sim_cons and Sim_nil . Coverage is discussed in details in (Thibodeau et al., 2016).

Similarly to other languages with support for coinduction, we achieve productivity through a guardedness check. However, this check is different from the one implemented in languages such as Coq or Agda. In Coq or Agda, we define coinductive types through lazy constructors. Thus, recursive calls are restricted, or guarded, as to be performed only under such constructors. With copatterns, coinductive terms are defined via observations which serve to delay computation. As such, guardedness is obtained by restricting recursive occurrences to be performed under observations; hence each step of unfolding requires at least one observation to be applied.

Next, we consider the proof that applicative similarity is transitive: if $\text{Sim } [T] \text{ [M]} \text{ [N]}$ and $\text{Sim } [T] \text{ [N]} \text{ [R]}$, then $\text{Sim } [T] \text{ [M]} \text{ [R]}$ (Fig. 8). We comment a couple of cases:

Observation Sim_unit and $T = \text{top}$. We have the following assumptions: $s1:\text{Sim } [\text{top}] \text{ [M]} \text{ [N]}$, $s2:\text{Sim } [\text{top}] \text{ [N]} \text{ [R]}$, and $d:\text{Eval } [M] \text{ [unit]}$ and we need to provide a witness for $\text{Eval } [R] \text{ [unit]}$. By the observation Sim_unit on $s2$ (written as the projection $s2.\text{Sim_unit}$), we obtain a function $\text{Eval } [N] \text{ [unit]} \rightarrow \text{Eval } [R] \text{ [unit]}$. We pass to it the result of $s1.\text{Sim_unit } d$.

Observation Sim_lam and $T = (\text{arr } S \ S')$. We have the assumptions: $s1:\text{Sim } [\text{Arr } S \ T] \text{ [M]} \text{ [N]}$ and $s2:\text{Sim } [\text{Arr } S \ T] \text{ [N]} \text{ [R]}$ and $d:[\text{eval } M \ (\text{lam } \lambda x.M')]$, and we need to build a witness for $\text{ExSim } [x:\text{term } S \vdash M'] \text{ [R]}$. We first observe $s1.\text{Sim_lam}$ and pass $d:\text{Eval } [M] \text{ [lam } \lambda x.M']$. This hence gives us $\text{ExSimLam } [x:\text{term } S \vdash M'] \text{ [N]}$, which provides us with $d1:\text{Eval } [N] \text{ [lam } \lambda x.N']$ and $s2:\{V:[\text{term } S]\} \text{ Sim } [T] \text{ [M']} [V] \text{ [N']} [V]$. Similarly, we observe $s2.\text{Sim_lam}$ passing $d1:\text{Eval } [N] \text{ [lam } \lambda x.N']$. Hence we

```

rec sim_trans : Sim [T] [M] [N] → Sim [T] [N] [R] → Sim [T] [M] [R] =
fun s1 s2 .Sim_unit d ⇒ s2.Sim_unit (s1.Sim_unit d)
| s1 s2 .Sim_nil d ⇒ s2.Sim_nil (s1.Sim_nil d)
| s1 s2 .Sim_cons d ⇒
  let ExSimcons d1' s1' s1'' = s1.Sim_cons d in
  let ExSimcons d2' s2' s2'' = s2.Sim_cons d1' in
  ExSimcons d2' (sim_trans s1' s2'') (sim_trans s1'' s2'')
| s1 s2 .Sim_lam d ⇒
  let ExSimlam d1 s3 = s1.Sim_lam d in
  let ExSimlam d2 s4 = s2.Sim_lam d1 in
  ExSimlam d2 (fun [P] ⇒ sim_trans (s3 [P]) (s4 [P]))

```

Fig. 8. Corecursive program showing that similarity is transitive

obtain $\text{ExSimLam } [x:\text{term } S \vdash N'] [R]$, which provides us with $\text{d2:Eval } [R] [\text{lam } \lambda x .R']$ and $\text{s4:}\{V:\{\text{term } S\}\} \text{Sim } [T] [N'[V]] [R'[V]]$. Building a witness for $\text{ExSim } [x:\text{term } S \vdash M'] [R]$ requires us to supply a derivation $\text{d2:Eval } [R] [\text{lam } \lambda x .R']$ and a function of type $\{W:\{\text{term } S\}\} \text{Sim } [T] [M'[W]] [R'[W]]$. We build that function making a corecursive call.

3.6. Defining Open Similarity Using First-Class Contexts and Substitutions

Similarity only relates closed terms. However, in general, we want to be able to reason about similarity of open terms, i.e. terms that depend on a context γ . In *Beluga*, we can declare schemas of contexts and work with contexts first-class. Context schemas classify contexts in the same way types classify terms and kinds classify types, describing the shape of each declaration in a context. Moreover, we can take advantage of first-class substitutions to *relate* two contexts. In particular, we can describe *grounding* substitutions with the type $[\vdash \gamma]$ where the range of the substitution is empty.

We begin by defining the schema of contexts that can occur in our development:

```
schema ctx = term T;
```

Here we declare the schema `ctx` that states that each declaration of a context γ of schema `ctx` can only contain variable declarations of type `term T` for some type `T`. For example, the context $x:\text{term top}, y:\text{term (list top)}$ is a valid context of schema `ctx`. On the other hand, a context $x:\text{term unit}, a:\text{tp}$ is not.

We can now state open similarity as an inductive type relating well-typed terms in the context γ . In the kind of the inductive type `OSim`, we make the type `T` explicit, but leave γ implicit. This distinction is reflected in *Beluga*'s source syntax. We use curly braces $\{ \}$ to describe explicit index arguments and round ones $()$ to give type annotations implicitly.

We can now define open similarity. Two terms $[\gamma \vdash M]$ and $[\gamma \vdash N]$ are openly similar if for all grounding substitutions σ , they are similar.

```

inductive Howe : (γ:ctx){T:[tp]} [γ ⊢ term T[] ] → [γ ⊢ term T[] ] → type =
  | Howe_unit : OSim [top] [γ ⊢ unit] [γ ⊢ M]
    → Howe [top] [γ ⊢ unit] [γ ⊢ M]
  | Howe_var : {#p:[γ ⊢ term T[]]} OSim [T] [γ ⊢ #p] [γ ⊢ M]
    → Howe [T] [γ ⊢ #p] [γ ⊢ M]
  | Howe_lam : Howe [T] [γ,x:term S[] ⊢ M] [γ,x:term S[] ⊢ N]
    → OSim [arr S T] [γ ⊢ lam λx.N] [γ ⊢ R]
    → Howe [arr S T] [γ ⊢ lam λx.M] [γ ⊢ R]
  | Howe_app : Howe [arr S T] [γ ⊢ M] [γ ⊢ M']
    → Howe [S] [γ ⊢ N] [γ ⊢ N']
    → OSim [T] [γ ⊢ app M' N'] [γ ⊢ R]
    → Howe [T] [γ ⊢ app M N] [γ ⊢ R]
  ...

```

Fig. 9. The Howe relation

```

inductive OSim : (γ:ctx){T:[tp]} [γ ⊢ term T[]] → [γ ⊢ term T[]] → type =
  | OSimC : ({σ:[ ⊢ γ]} Sim [T] [M[σ]] [N[σ]]) → OSim [T] [γ ⊢ M] [γ ⊢ N]

```

We can easily show that open similarity is closed under substitutions by simply composing the input substitution σ with the closing substitution σ' .

```

rec osim_cus : (γ:ctx) (ψ:ctx) {σ:[ψ ⊢ γ]} OSim [T] [γ ⊢ M] [γ ⊢ N]
  → OSim [T] [ψ ⊢ M[σ]] [ψ ⊢ N[σ]] =
fun [ψ ⊢ σ] s ⇒ let OSimC f = s in OSimC (fun [σ'] ⇒ f [σ[σ']])

```

3.7. Defining the Howe Relation on Open Terms

The encoding of the Howe relation (see Fig. 9) is, in our possibly biased view, one of the high point of the formalization: it follows very closely its mathematical formulation, while retaining all the powerful abstractions that *Beluga* offers. This is apparent in the the variable case, where *Beluga*'s *parameter* variables, ranging over elements from the context γ (written as $[\gamma \vdash \#p]$), permit us to precisely characterize when a variable is Howe related to a term M in the given context, while looking remarkably similar to the on-paper version. The same applies to lambda-abstractions case, where one notes the correct scoping of M , N and R w.r.t. γ . The cases for `fix` and `lcase` follow the same principle and we omit them for space reasons.

This is notably different from the only other comparable formalization (Momigliano, 2012), which, lacking the possibility of abstracting over contexts had to rely on a very concrete notion of context. This in turn made the rest of the development, in particular the proof of the substitutivity of the Howe relation fairly painful.

Using reflexivity and transitivity of open similarity, respectively, we can show reflexivity and semi transitivity of the candidate relation. We only show the types.

```

rec howe_refl : {γ:ctx} {M:[γ ⊢ term T[] ]} Howe [T] [γ ⊢ M] [γ ⊢ M]

```

```

inductive HoweSubst :
  { $\gamma$ :ctx} ( $\psi$ :ctx) { $\sigma_1$  : [ $\psi \vdash \gamma$ ]} { $\sigma_2$  : [ $\psi \vdash \gamma$ ]} type =
  | HNil : HoweSubst [] [ $\psi \vdash$ ] [ $\psi \vdash$ ]
  | HCons : HoweSubst [ $\gamma$ ] [ $\psi \vdash \sigma_1$ ] [ $\psi \vdash \sigma_2$ ]
    → Howe [T] [ $\psi \vdash M$ ] [ $\psi \vdash N$ ]
    → HoweSubst [ $\gamma, x$ :term T[]] [ $\psi \vdash \sigma_1, M$ ] [ $\psi \vdash \sigma_2, N$ ]

rec howeSubst_wkn : HoweSubst [ $\gamma$ ] [ $\psi \vdash \sigma_1$ ] [ $\psi \vdash \sigma_2$ ]
  → HoweSubst [ $\gamma$ ] [ $\psi, x$ :term S[]  $\vdash \sigma_1$ [...]] [ $\psi, x$ :term S[]  $\vdash \sigma_2$ [...]]

```

Fig. 10. The Howe relation on substitutions

```

rec howe_osim_trans : ( $\gamma$ :ctx) Howe [T] [ $\gamma \vdash M1$ ] [ $\gamma \vdash M2$ ]
  → OSim [T] [ $\gamma \vdash M2$ ] [ $\gamma \vdash M3$ ]
  → Howe [T] [ $\gamma \vdash M1$ ] [ $\gamma \vdash M3$ ]

```

From this it immediately follows that open similarity is a Howe relation.

```

rec osim_howe : ( $\gamma$ :ctx) OSim [T] [ $\gamma \vdash M$ ] [ $\gamma \vdash N$ ]
  → Howe [T] [ $\gamma \vdash M$ ] [ $\gamma \vdash N$ ] =
fun (s : OSim [T] [ $\gamma \vdash M$ ] [ $\gamma \vdash N$ ]) ⇒
  howe_osim_trans (howe_refl [ $\gamma$ ] [ $\gamma \vdash M$ ]) s

```

3.8. Substitutivity of the Howe Relation

As remarked in Section 2.2, a crucial point of the proof is showing that the Howe relation is substitutive. Traditionally, substitution properties tend to be tedious to prove in proof assistants due to the necessity to reason manually about contextual meta-theory. Here, Beluga’s contextual abstractions significantly reduces the amount of boilerplate work needed for that proof.

We first encode (Fig. 10) $\Gamma' \vdash \sigma_1 \lesssim_{\Gamma}^H \sigma_2$ using an inductive type that relates two simultaneous substitutions. The base case relates empty substitutions, written as [$\psi \vdash$]. In the inductive case, the substitution [$\psi \vdash \sigma_1, M$] and [$\psi \vdash \sigma_2, N$] are related, if so are [$\psi \vdash \sigma_1$] and [$\psi \vdash \sigma_2$] and [$\psi \vdash M$] is Howe related to [$\psi \vdash N$]. Compare it with the mathematical definition 4.

In the subsequent proofs, we rely on the weakening property of simultaneous substitutions: namely, that weakening preserves Howe-relatedness, see function `howeSubst_wkn` in Fig. 10. In Beluga, weakening a substitution is simply achieved by composing it with the weakening substitution `[...]`, which has here domain ψ and range ψ, x :term S[]. This is supported in Beluga’s theory of simultaneous substitutions Cave and Pientka (2013), which internalizes the notions in Fig. 3. We also need the following reflexivity property of `HoweSubst`, which holds by a simple induction on substitutions:

```

rec howeSubst_refl : ( $\gamma$ :ctx) ( $\psi$ :ctx) { $\sigma$ : [ $\psi \vdash \gamma$ ]} HoweSubst [ $\gamma$ ] [ $\psi \vdash \sigma$ ] [ $\psi \vdash \sigma$ ]

```

The proof of substitutivity in Beluga appears in Fig. 11. Making use of the lemmas described above, we see that it straightforwardly represents the proof of Lemma 6. We only show here the same two cases we described in the on paper proof, but the remaining cases follow a similar pattern. What is remarkable in this program with respect to the on paper proof is that there are no explicit references to the substitution properties outside of the weakening of the Howe relation on substitutions. The encoding is very concise and captures the essential steps in the proof.

```

rec howe_subst : Howe [T] [ $\gamma \vdash M$ ] [ $\gamma \vdash N$ ]
  → HoweSubst [ $\gamma$ ] [ $\psi \vdash \sigma_1$ ] [ $\psi \vdash \sigma_2$ ]
  → Howe [T] [ $\psi \vdash M[\sigma_1]$ ] [ $\psi \vdash N[\sigma_2]$ ] =
fun h (hs:HoweSubst [ $\gamma$ ] [ $\psi \vdash \sigma_1$ ] [ $\psi \vdash \sigma_2$ ]) ⇒
case h of
...
| Howe_lam h' s ⇒
  Howe_lam (howe_subst h' (HCons (howeSubst_wkn hs)
    (howe_refl [ $\psi, x:\text{term } \_$ ] [ $\psi, x:\text{term } \_ \vdash x$ ])))
    (osim_cus [ $\psi \vdash \sigma_2$ ] s)
| Howe_app h1' h2' s ⇒
  Howe_app (howe_subst h1' hs) (howe_subst h2' hs) (osim_cus [ $\psi \vdash \sigma_2$ ] s);

```

Fig. 11. Substitutivity property of the Howe relation

3.9. Main Theorem

The key lemma in our main theorem is the proof that the Howe relation is downward closed:

```

rec down_closed : Eval [P] [V] → Howe [T] [P] [Q] → Howe [T] [V] [Q]

```

The proof of this lemma (point 7 at page 8) relies on several previous lemmas such as transitivity of closed and open similarity, semi-transitivity and substitutivity of the Howe relation, together with the unfolding of similarity using the observations. The proof is otherwise straightforward but long. We leave it to the online documentation.

We are now almost ready to prove that the Howe relation is an open similarity. For this proof, we first establish lemmas that mimic the similarity conditions (previous point 6). For example: If $\lambda x. m \preceq_{\tau \rightarrow \tau'}^H n$, then $n \Downarrow \lambda x. m'$ and for every $q:\tau$ we have $[q/x]m \preceq_{\tau'}^H [q/x]m'$. Again, as we do not have existential types, we encode the existence of a term N' using the inductive types `HoweAbs`. A fragment of the type signature is as follows:

```

inductive HoweAbs : [ $x:\text{term } T \vdash \text{term } T' []$ ] → [ $\text{term } (\text{arr } T T')$ ] → type =
  | HoweAbsC : Eval [N] [ $\text{lam } \lambda x. N'$ ]
    → ( $\{Q:[\text{term } T]\}$  Howe [T'] [M'] [Q]] [N' [Q]])
    → HoweAbs [ $x:\text{term } T \vdash M'$ ] [N];

```

```

rec howe_sim : Howe [T] [M] [N] → Sim [T] [M] [N] =
fun h .Sim_unit e ⇒ howe_ev_unit (down_closed e h)
  | h .Sim_nil e ⇒ howe_ev_nil (down_closed e h)
  | h .Sim_cons e ⇒
    let HoweConsC e' h1 h2 = howe_ev_cons (down_closed e h) in
      ExSimcons e' (howe_sim h1) (howe_sim h2)
  | h .Sim_lam e ⇒
    let HoweAbsC e' f = howe_ev_abs (down_closed e h) in
      ExSimlam e' (fun [P] ⇒ howe_sim (f [P]))

rec howe_osim : {γ:ctx} Howe [T] [γ ⊢ M] [γ ⊢ N]
  → OSim [T] [γ ⊢ M] [γ ⊢ N] =
fun [γ] h ⇒ OSimC (fun [σ] ⇒ howe_sim (howe_subst (howeSubst_refl [σ])));

```

Fig. 12. The Howe relation is an open simulation

```

rec howe_ev_abs : Howe [arr T T'] [lam λx.M'] [N]
  → HoweAbs [x:term T ⊢ M'] [N]

```

We are now ready to prove that the Howe relation is an open simulation. We do this by first proving that, in a closed context, the Howe relation is a similarity, then we embed the open version into an open similarity. To do so, we construct out of the input substitution σ for open similarity a derivation $\cdot \vdash \sigma \preceq_{\Gamma}^H \sigma$ by reflexivity. The proofs appear in Fig. 12.

3.10. A Concrete Example of Similarity

In this section we show how we can interactively build an actual simulation between two terms, namely that `two` is simulated by `suc one`, following the example in (Pitts, 2011). We represent the numbers via Church encodings, by which `one` $\equiv \lambda f.\lambda x.f x$, `two` $\equiv \lambda f.\lambda x.f (f x)$, and `suc` $\equiv \lambda n.\lambda x.\lambda y.x (n x y)$. We thus want to prove the following theorem:

```

rec sim_two_succ_one :
  Sim [_] [lam λf.lam λx.app f (app f x)]
    [app (lam λn. lam λx. lam λy. app x (app (app n x) y))
      (lam λf.lam λx.app f x)]

```

Since our operational semantics does not reduce under a lambda, the evaluation of `suc one` results in $\lambda x.\lambda y.x ((\lambda u.\lambda w.u w) x y)$ rather than evaluating directly to $\lambda f.\lambda x.f (f x)$. Thus, trying to build the simulation directly by relying on the evaluation sequence would leave us stuck trying to relate $e_1 ((\lambda u.\lambda w.u w) e_1 e_2)$ with $e_1 (e_1 e_2)$, for some e_1 and e_2 . Since e_1 is abstract, we cannot β -reduce and progress in the evaluation.

We will build the proof incrementally, by inserting *holes*, denoted by `?` and refining

them, analogously to Agda or Epigram’s methodology. We start with the following program body:

```
fun .Sim_lam (Ev-val Val-lam) ⇒
  ExSimlam (Ev-app (Ev-val Val-lam) (Ev-val Val-lam)) sim_lemma1
```

where `sim_lemma1` is used to abstract over the nested copattern matching:

```
rec sim_lemma1:{E:[exp (arr T T)]}
  Sim [arr T T]
    [lam (λx. app E[] (app E[] x))]
    [lam (λy. app E[] (app (app (lam (λf. lam (λw. app f w))) E[] y))] =
  fun [E1] .Sim_lam (Ev-val Val-lam) ⇒
    ExSimlam (Ev-val Val-lam)
      (fun [E2] ⇒ ?)
```

By design, it is easy to show that Howe’s relation is a pre-congruence (`howe_cong_app` below), and since we proved the equivalence between Howe’s relation and similarity, it follows that similarity is pre-congruence as well:

```
rec howe_cong_app : Howe [arr S T] [M1] [M2] → Howe [S] [N1] [N2]
  → Howe [T] [app M1 N1] [app M2 N2] =
  fun h1 h2 ⇒ Howe_app h1 h2 (osim_refl []);

rec sim_cong_app : Sim [arr S T] [M1] [M2] → Sim [S] [N1] [N2]
  → Sim [T] [app M1 N1] [app M2 N2] =
  fun s1 s2 ⇒ howe_sim (howe_cong_app (sim_howe s1) (sim_howe s2));
```

Using this result and reflexivity of similarity, we can thus refine the body of `sim_lemma1`:

```
fun [E1] .Sim_lam (Ev-val Val-lam) ⇒
  ExSimlam (Ev-val Val-lam)
    (fun [E2] ⇒ sim_cong_app (sim_refl [E1]) ?)
```

where the current hole has type:

```
Sim [T] [app E1 E2] [app (app (lam (λu. lam (λw. app u w))) E1) E2]
```

Now, we can easily use a derivation of the evaluation of the left-hand side to derive the evaluation of the right-hand side of this similarity as follows:

```
rec ev1 : Eval [app E1 E2] [V] →
  Eval [app (app (lam (λu. lam (λw. app u w))) E1) E2] [V] =
  fun d ⇒ Ev-app (Ev-app (Ev-val Val-lam) (Ev-val Val-lam)) d;
```

Constructing the above simulation requires us to match on the possible values that `app E1 E2` can take through the possible observations, in fact all of them as the type `T` is abstract. We then use `ev1` on the derivations of `Eval [app E1 E2] [V]` for the given `V` and reflexivity when needed.

```
rec sim_lemma2 : Sim [T] [app E1 E2]
  [app (app (lam (λu. lam (λw. app u w))) E1) E2] =
  fun .Sim_lam (d : Eval [app E1 E2] [lam (λx.E)]) ⇒
```

```

    ExSimlam (ev1 d) (fun [V] => sim_refl [E[V]])
  | .Sim_top d => ev1 d
  | .Sim_nil d => ev1 d
  | .Sim_cons d => ExSimcons (ev1 d) (sim_refl [_]) (sim_refl [_])

```

Here, we use a type annotation on `d` to expose `E` in order for unification to accept the call to `sim_refl`. Using this lemma, we can complete the body of `sim_lemma1`:

```

fun [E1] .Sim_lam (Ev-val Val-lam) =>
  ExSimlam (Ev-val Val-lam)
    (fun [E2] => sim_cong_app (sim_refl [E1]) sim_lemma2)

```

This concludes the proof.

4. Related Work

In this section we will *not* review the various approaches to co-induction in proof assistants or even theorem provers (Leino and Moskal, 2014), let alone the literature on program equivalence, but will make some comments about the former nevertheless as we go along.

The first HOAS-like formal verification of the congruence of a notion of bisimilarity concerned the π -calculus (Honsell et al., 2001) and was carried out in Coq using the Weak HOAS approach and instantiating the *Theory of Contexts* to axiomatizing properties of names. As common in many coinductive developments in Coq, the authors soon ran afoul of the guardedness checker in *Cofix*-style proofs and had to resort to an explicit greatest-fixed point encoding for Strong Late Bisimilarity. Abella’s take to the same issue (Tiu and Miller, 2010) seems preferable; that paper details, among so much more, a proof that similarity is a pre-congruence for the finite π -calculus. The encoding is most elegant, where all issues involving bindings, names, and substitutions are handled declaratively without explicit side-conditions, thanks to the ∇ -quantifier.

Encoding bisimilarity in the λ -calculus, in particular via Howe’s method, brings in additional challenges, as we have seen. We are aware of three formalization through the years:

- 1 In (Ambler and Crole, 1999) the authors verify in Isabelle/HOL 98 the same result of the present paper and a bit more (they also show that similarity coincides with contextual pre-order) for PCFL using De Bruijn indexes as an encoding techniques for binders. The development, for the congruence part, consists of around 160 lemmas/theorems, and it confirms a common belief about (standard) concrete syntax approaches: doable, but very hard-going;
- 2 A partial improvement was presented in (Momigliano et al., 2002), which was based on the HOAS approach implemented in an early version of the *Hybrid* tool (Feltz and Momigliano, 2012), but one crucial lemma was left unproven, tellingly: Howe’s substitutivity, related to the difficulty of lifting in one-level Hybrid term substitution as β -conversion to substitution on judgments;
- 3 Momigliano (2012) fixed the situation, giving a complete Abella proof for the simply typed calculus with unit. The proof consists of circa 45 theorems, 1/4 of which

devoted to maintaining typing invariants in (open)similarity and the in candidate relation, 1/7 of which instead used to make sure that some ∇ -quantified variable cannot occur in certain predicates. The main source of difficulty was again in the proof of substitutivity of the Howe relation, which is formalized at the meta-level. Since Abella does not have a first class notion of substitution and has good supports for *implicit* but not for *explicit* contexts, such as the one occurring in open similarity and the Howe relation, significant effort was required in handling weakening and exchange. In fact, it may not be a coincidence that the author never summoned the energy to extend this proof to PCFL, as promised in (Momigliano, 2012).

In a very recent paper McLaughlin et al. (2017) give a formalization of the coincidence of observational and applicative approximation not going through the candidate relation, but triangulating with a notion of *logical* (as in logical relations) approximation. This is then extended to CIU approximation. The encoding uses first-order syntax for terms, but a form of weak HOAS for judgments following Allais et al. (2017), and it is therefore compatible with a standard proof assistant as Agda. Similarly to us, it leverages the use of intrinsically well-typed and scoped terms and simultaneous substitutions, although the latter are not supported natively by the framework. Interestingly, it offers an elegant notion of concrete context (and thus of Morris approximation) that seems much easier to reason with than previous efforts (Ford and Mason, 2003).

5. Conclusion

We outlined how to use the *Beluga* proof environment to encode a text book example of reasoning about program equivalence using Howe’s method for PCFL. This reinforced several observations that we have been making in other case studies, viz. (Cave and Pientka, 2015):

- Using intrinsically typed terms instead of working with explicit typing invariants makes our encoding more compact and easier to deal with, whereas HOAS maintained our terms well-scoped. These observations are by now gaining traction also in non-HOAS approach as Allais et al. (2017).
- The support for first-class simultaneous substitutions and contexts lead us to formulate for example substitutivity more generally; but this paid off in our mechanization, as many required lemmas became simpler to prove thanks to *Beluga*’s built-in handling of the theory of substitutions.
- Thanks to catering for both indexed inductive and coinductive data-types in *Beluga*, the encoding of similarity and of the candidate relation was compact and concise.

Future work. The attentive reader will have noticed that we have not shown that bisimilarity coincides with contextual equivalence, as e.g., in (Ambler and Crole, 1999). The difficulty lies in the encoding of the latter notion: using concrete non-capture-avoiding context seems *prima facie* incompatible with HOAS, while Lassen’s and Pitt’s context-less approach, which defines contextual equivalence as the largest *adequate and compatible* relation (Lassen, 1998; Pitts, 2011), requires extending *Beluga* to at least second-order

quantification. This is work in progress. Program equivalence can and has been tackled with *step-indexed logical relations* (Ahmed, 2006): **Beluga** has proven to be, arguably, the most advanced environment to carry such style of proofs (Cave and Pientka, 2015) and it would be interesting to compare the two developments and see how they scale to more complex programming languages such as in (Pitts, 2005).

In the future, we plan to extend the implementation of our coverage checker to also handle coinductive definitions and copattern matching following Thibodeau et al. (2016). Further, we plan to extend the totality checker to check productivity of our functions defined by copattern matching. This will aim to exploit the simple structural criteria of being guarded by an observation as outlined in Abel and Pientka (2013).

While we hope that we have succeeded in showing that **Beluga** is an excellent environment for the meta-theory of program equivalence, it is not well suited (yet) to verifying that concrete pieces of code may be equivalent. To approach this, we need not only to improve the interactive proof construction mode similarly to what Agda offers, but to investigate proof search and refutation, possibly following Bedwyr (Baelde et al., 2007).

References

- Abel, A. and Pientka, B. (2013). Well-founded recursion with copatterns: a unified approach to termination and productivity. In *Proceedings of the 18th ACM SIGPLAN International Conference on Functional Programming (ICFP '13)*, pages 185–196.
- Abel, A. and Pientka, B. (2016). Well-founded recursion with copatterns and sized types. *Journal of Functional Programming*, 26:e2 (61 pages).
- Abel, A., Pientka, B., Thibodeau, D., and Setzer, A. (2013). Copatterns: Programming infinite structures by observations. In *40th ACM Symp. on Principles of Programming Languages (POPL'13)*, pages 27–38. ACM Press.
- Abella (2012). The Abella prover. Available at <http://abella-prover.org/>.
- Ahmed, A. (2006). Step-indexed syntactic logical relations for recursive and quantified types. In Sestoft, P., editor, *15th European Symposium on Programming (ESOP'06)*, pages 69–83. Springer.
- Allais, G., Chapman, J., McBride, C., and McKinna, J. (2017). Type-and-scope safe programs and their proofs. In Bertot, Y. and Vafeiadis, V., editors, *6th ACM SIGPLAN Conference on Certified Programs and Proofs (CPP'17)*, pages 195–207. ACM.
- Ambler, S. and Crole, R. L. (1999). Mechanized operational semantics via (co)induction. In Bertot, Y., Dowek, G., Hirschowitz, A., Paulin, C., and Théry, L., editors, *12th International Conference on Theorem Proving in Higher Order Logics (TPHOLs'99)*, Lecture Notes in Computer Science (LNCS 1690), pages 221–238. Springer.
- Baelde, D. (2011). Least and greatest fixed points in linear logic. *ACM Transactions on Computational Logic*, 13(1).
- Baelde, D., Gacek, A., Miller, D., Nadathur, G., and Tiu, A. (2007). The Bedwyr system for model checking over syntactic expressions. In *21st Conference on Automated Deduction (CADE)*, Lecture Notes in Computer Science (LNCS 4603), pages 391–397. Springer.
- Cave, A. and Pientka, B. (2012). Programming with binders and indexed data-types. In

- 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'12), pages 413–424. ACM Press.
- Cave, A. and Pientka, B. (2013). First-class substitutions in contextual type theory. In *8th ACM SIGPLAN International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP'13)*, pages 15–24. ACM Press.
- Cave, A. and Pientka, B. (2015). A case study on logical relations using contextual types. In Cervesato, I. and Chaudhuri, K., editors, *10th International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP'15)*, pages 18–33. Electronic Proceedings in Theoretical Computer Science (EPTCS).
- Felty, A. P. and Momigliano, A. (2012). Hybrid: A definitional two-level approach to reasoning with higher-order abstract syntax. *Journal of Automated Reasoning*, 48(1):43–105.
- Ford, J. and Mason, I. A. (2003). Formal foundations of operational semantics. *Higher-Order and Symbolic Computation*, 16(3):161–202.
- Gacek, A., Miller, D., and Nadathur, G. (2008). Combining generic judgments with recursive definitions. In Pfenning, F., editor, *23rd Symposium on Logic in Computer Science*. IEEE Computer Society Press.
- Giménez, E. (1996). *Un Calcul de Constructions Infinies et son application à la vérification de systèmes communicants*. PhD thesis, Ecole Normale Supérieure de Lyon. Thèse d'université.
- Harper, R., Honsell, F., and Plotkin, G. (1993). A framework for defining logics. *Journal of the ACM*, 40(1):143–184.
- Hirschhoff, D. (1997). A full formalisation of pi-calculus theory in the calculus of constructions. In Gunter, E. L. and Felty, A. P., editors, *10th International Conference on Theorem Proving in Higher Order Logics (TPHOLs'97)*, Lecture Notes in Computer Science (LNCS 1275), pages 153–169. Springer.
- Honsell, F., Miculan, M., and Scagnetto, I. (2001). Π -calculus in (co)inductive type theories. *Theoretical Computer Science*, 2(253):239–285.
- Howe, D. J. (1996). Proving congruence of bisimulation in functional programming languages. *Information and Computation*, 124(2):103–112.
- Jacob-Rao, R. and Pientka, B. (2017). Index-stratified types.
- Lassen, S. B. (1998). *Relational Reasoning about Functions and Nondeterminism*. PhD thesis, Dept of Computer Science, Univ of Aarhus.
- Lee, D. K., Crary, K., and Harper, R. (2007). Towards a mechanized metatheory of standard ml. In *34th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'07)*, pages 173–184. ACM Press.
- Leino, K. R. M. and Moskal, M. (2014). Co-induction simply. In Jones, C., Pihlajasaari, P., and Sun, J., editors, *19th International Symposium on Formal Methods (FM'14)*, pages 382–398. Springer.
- Mason, I. and Talcott, C. (1991). Equivalence in functional languages with effects. *Journal of Functional Programming*, 1(03):287–327.
- McDowell, R. and Miller, D. (1997). A logic for reasoning with higher-order abstract syntax. In Winskel, G., editor, *12th Symp. on Logic in Computer Science*, pages 434–445. IEEE Computer Society Press.

- McDowell, R., Miller, D., and Palamidessi, C. (1996). Encoding transition systems in sequent calculus: Preliminary report. *Electronic Notes in Theoretical Computer Science (ENTCS)*, 3.
- McLaughlin, C., McKinna, J., and Stark, I. (2017). Triangulating context lemmas. In *Submitted*.
- Miller, D. and Nadathur, G. (2012). *Programming with Higher-Order Logic*. Cambridge University Press, New York, NY, USA, 1st edition.
- Miller, D. and Palamidessi, C. (1999). Foundational aspects of syntax. *ACM Comput. Surv.*, 31(3es).
- Milner, R. (1977). Fully abstract models of typed λ -calculi. *Theoretical Computer Science*, 4(1):1 – 22.
- Momigliano, A. (2012). A supposedly fun thing I may have to do again: A HOAS encoding of Howe’s method. In *7th International Workshop on Logical Frameworks and Meta-languages: Theory and Practice (LFMTP’12)*, pages 33–42. ACM.
- Momigliano, A., Ambler, S., and Crole, R. L. (2002). A Hybrid encoding of Howe’s method for establishing congruence of bisimilarity. *Electr. Notes Theor. Comput. Sci.*, 70(2).
- Momigliano, A. and Tiu, A. (2003). Induction and co-induction in sequent calculus. In Coppo, M., Berardi, S., and Damiani, F., editors, *Post-proceedings of TYPES 2003*, Lecture Notes in Computer Science (LNCS 3085), pages 293–308.
- Nanevski, A., Pfenning, F., and Pientka, B. (2008). Contextual modal type theory. *ACM Transactions on Computational Logic*, 9(3):1–49.
- Oury, N. (2008). Coinductive types and type preservation. Message on the coq-club mailing list.
- Pfenning, F. (1997). Computation and deduction.
- Pfenning, F. and Schürmann, C. (1999). System description: Twelf — a meta-logical framework for deductive systems. In Ganzinger, H., editor, *16th International Conference on Automated Deduction (CADE-16)*, Lecture Notes in Artificial Intelligence (LNAI 1632), pages 202–206. Springer.
- Pientka, B. (2008). A type-theoretic foundation for programming with higher-order abstract syntax and first-class substitutions. In *35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL’08)*, pages 371–382. ACM Press.
- Pientka, B. (2013). An insider’s look at LF type reconstruction: Everything you never wanted to know. *J. Funct. Program.*, 23(1):1–37.
- Pientka, B. and Cave, A. (2015). Inductive Beluga: Programming Proofs (System Description). In Felty, A. P. and Middeldorp, A., editors, *25th International Conference on Automated Deduction (CADE-25)*, Lecture Notes in Computer Science (LNCS 9195), pages 272–281. Springer.
- Pientka, B. and Dunfield, J. (2008). Programming with proofs and explicit contexts. In *ACM SIGPLAN Symposium on Principles and Practice of Declarative Programming (PPDP’08)*, pages 163–173. ACM Press.
- Pientka, B. and Dunfield, J. (2010). Beluga: a framework for programming and reasoning with deductive systems (System Description). In Giesl, J. and Haehnle, R., editors, *5th*

- International Joint Conference on Automated Reasoning (IJCAR'10)*, Lecture Notes in Artificial Intelligence (LNAI 6173), pages 15–21. Springer.
- Pitts, A. M. (1997). Operationally Based Theories of Program Equivalence. In Dybjer, P. and Pitts, A. M., editors, *Semantics and Logics of Computation*.
- Pitts, A. M. (2005). Typed operational reasoning. In Pierce, B. C., editor, *Advanced Topics in Types and Programming Languages*, chapter 7, pages 245–289. The MIT Press.
- Pitts, A. M. (2011). Howe’s method for higher-order languages. In Sangiorgi, D. and Rutten, J., editors, *Advanced Topics in Bisimulation and Coinduction*, volume 52, chapter 5, pages 197–232. Cambridge University Press.
- Thibodeau, D., Cave, A., and Pientka, B. (2016). Indexed codata. In Garrigue, J., Keller, G., and Sumii, E., editors, *21st ACM SIGPLAN International Conference on Functional Programming (ICFP'16)*, pages 351–363. ACM.
- Tiu, A. and Miller, D. (2010). Proof search specifications of bisimulation and modal logics for the π -calculus. *ACM Trans. Comput. Logic*, 11(2):1–35.

Appendix A. Overview of Beluga Source Language

A Beluga signature consists of LF declarations, inductive and coinductive definitions, and programs. For each LF type family \mathbf{a} we declare its LF kind together with the constants that allow us to form objects that inhabit the type family. We also support mutual LF definitions that do not capture in our grammar below not to complicate matters.

Signature Decl. $\mathcal{D} ::= \mathbf{LF} \ \mathbf{a} : \mathcal{K}_{\text{LF}} = \mathbf{c}_1 : \mathcal{A}_1 \mid \dots \mid \mathbf{c}_k : \mathcal{A}_k;$
 $\mathbf{inductive} \ \mathbf{a} : \mathcal{K} = \mathbf{c}_1 : \mathcal{T}_1 \mid \dots \mid \mathbf{c}_n : \mathcal{T}_n;$
 $\mathbf{coinductive} \ \mathbf{a} : \mathcal{K} = (\mathbf{c}_1 : \mathcal{T}_1) :: \mathcal{T}'_1 \mid \dots \mid (\mathbf{c}_n : \mathcal{T}_n) :: \mathcal{T}'_n;$
 $\mathbf{rec} \ \mathbf{c} : \mathcal{T} = \mathcal{E};$

(Co)inductive definitions correspond to (co)indexed recursive types and semantically are interpreted as (greatest) least fixed points. We define an indexed recursive type family by defining constructors \mathbf{c}_i that can be used to construct elements of a given indexed recursive type. We define a corecursive type by the observations we can make about it using indexed records, where we write the field \mathbf{c}_i together with the type \mathcal{T}_i from which we can project the result \mathcal{T}'_i . Given a Beluga term of \mathcal{E} type \mathcal{T}_i we may use the projection \mathbf{c}_i and the result of $\mathcal{E}.\mathbf{c}_i$ is then of type \mathcal{T}'_i .

We describe Beluga's type and terms more precisely in Fig. 13. The syntax for LF kinds, types and terms is close to the syntax \mathbf{x} in the Twelf system. We write curly braces $\{ \ }$ for the dependent LF function space and allow users to write simply \rightarrow , if there is no dependency. LF kinds, types, and LF terms used in a signature declaration must be pure, i.e. they cannot refer to closures highlighted in blue and written as $\mathbf{x}[\sigma]$ and $\#\mathbf{p}[\sigma]$. Here \mathbf{x} and $\#\mathbf{p}$ are meta-variables that are bound and introduced in Beluga types and patterns.

LF Kinds	\mathcal{K}_{LF}	$::=$	type $\mid \{ \mathbf{x} : \mathcal{A} \} \mathcal{K}_{\text{LF}} \mid \mathcal{A} \rightarrow \mathcal{K}_{\text{LF}}$
LF Types	\mathcal{A}	$::=$	a $\mathcal{M}_1 \dots \mathcal{M}_n \mid \{ \mathbf{x} : \mathcal{A} \} \mathcal{A}' \mid \mathcal{A} \rightarrow \mathcal{A}'$
LF Terms	\mathcal{M}	$::=$	$\mathbf{x} \mid \lambda \mathbf{x} . \mathcal{M} \mid \mathbf{c} \ \mathcal{M}_1 \dots \mathcal{M}_n \mid \mathbf{x}[\sigma] \mid \#\mathbf{p}[\sigma]$
LF Subst.	σ	$::=$	$_ \mid \dots \mid \sigma, \mathcal{M}$
LF Context	Ψ, Φ	$::=$	$_ \mid \psi \mid \Psi, \mathbf{x} : \mathcal{A}$
Contextual Type	\mathcal{U}	$::=$	$[\Psi \vdash \mathbf{a} \ \mathcal{M}_1 \dots \mathcal{M}_n] \mid [\text{ctx}] \mid [\Psi \vdash \Phi] \mid \dots$
Contextual Object	\mathcal{C}	$::=$	$[\Psi \vdash \mathcal{M}] \mid [\Psi] \mid \dots$

Beluga Kinds	\mathcal{K}	$::=$	type $\mid \{ \mathbf{x} : \mathcal{U} \} \mathcal{K} \mid \mathcal{U} \rightarrow \mathcal{K}$
Beluga Types	\mathcal{T}	$::=$	$\{ \mathbf{x} : \mathcal{U} \} \mathcal{T} \mid (\mathbf{x} : \text{ctx}) \mathcal{T} \mid \mathcal{T} \rightarrow \mathcal{T} \mid \mathcal{U} \mid \mathbf{a} \ \mathcal{C}_1 \dots \mathcal{C}_n$
Beluga Terms	\mathcal{E}	$::=$	fun $\mathcal{B} \mid \mathbf{let} \ \mathcal{P} = \mathcal{E} \ \mathbf{in} \ \mathcal{E} \mid \mathcal{E} \ \mathcal{E} \mid \mathcal{C} \mid \mathbf{x} \mid \mathbf{c} \mid \mathcal{E} . \mathbf{c}$
Beluga Branches	\mathcal{B}	$::=$	$_ \mid (\mathcal{B} \mid \mathcal{R}_1 \dots \mathcal{R}_n \Rightarrow \mathcal{E})$
Beluga Pattern	\mathcal{P}	$::=$	$\mathbf{x} \mid \mathcal{C} \mid \mathbf{c} \ \mathcal{P}_1 \dots \mathcal{P}_n$
Beluga Copatterns	\mathcal{R}	$::=$	$\cdot \mid \mathcal{P} \ \mathcal{R} \mid \cdot \mathbf{c} \ \mathcal{R}$

Fig. 13. Grammar of Beluga

Substitutions in closures are either empty, written as $_$ here, or a weakening substitution \dots , which we use in practice to transition from a context to a possible extension. Substitutions can also be built by extending a substitution σ with a LF term \mathcal{M} .

LF contexts may be empty, consists of a context variable, or are built by concatenating to a LF context a LF variable declaration.

Contextual types and terms pair a LF context together with a LF type or LF term respectively. As LF contexts are first-class in **Beluga**, they form valid contextual objects. LF contexts are in general classified by a schema that allows us to state an invariant the LF context satisfies; here we only add the one schema we have used in this development (Section 3.6), namely `ctx`.

Finally, we come to **Beluga**'s computation language. It allows us to make statements about contextual types and objects and we highlight them in the color blue. It is in many ways similar to standard functional programming languages with two exceptions: first, contextual types and objects are the special domain and we support not only pattern matching, but also copattern matching, i.e., our patterns may include projections describing the observations we can make about a given (output) type. We only describe here the part of **Beluga**'s term language that is relevant for our example. It consists of function definitions that use (co)pattern matching, let-expressions, and applications. Further, we include computation-level variables, constructors and constants described using `c`, and projections.

In addition to indexed dependent function space, $\{\mathbf{x}:\mathcal{U}\}\mathcal{T}$, **Beluga**'s type language supports simple types, embedding contextual types, and (co)inductive definitions, described as $\mathbf{a} \ C_1 \dots C_n$. Note that in $\{\mathbf{x}:\mathcal{U}\}\mathcal{T}$ we make \mathbf{x} explicit and hence any computation-level expression of that type expects first an object of type \mathcal{U} . We also allow a limited form of implicit type annotation for context variables; by writing $(\mathbf{x}:\text{ctx})\mathcal{T}$ we can abstract over the context variable \mathbf{x} and declare its context schema, while keeping \mathbf{x} implicit.