

Formal Program Optimization in Nuprl Using Computational Equivalence and Partial Types

Vincent Rahli, Mark Bickford, and Abhishek Anand

Cornell University, Ithaca, NY, USA

Abstract. This paper extends the proof methods used by the Nuprl proof assistant to reason about the computational behavior of its untyped programs. We have implemented new methods to prove non-trivial bisimulations between programs and have successfully applied these methods to formally optimize distributed programs such as our synthesized and verified version of Paxos, a widely used protocol to achieve software based replication. We prove new results about the basic computational equality relation on terms, and we extend the theory of partial types as the basis for stating internal results about the computation system that were previously treated only in the meta theory of Nuprl. All the lemmas presented in this paper have been formally proved in Nuprl.

1 Introduction

This paper presents proof techniques implemented in the Nuprl proof assistant [16,27,4] to reason about its own computation system and programming language, an applied lazy (call-by-name) λ -calculus. Since the computation system is universal (Turing complete), we need to reason using *partial types* introduced by Constable and Smith [35,17] and extended by Crary [18].¹ The *bisimulation* relation defined by Howe turned out to form a contextual equivalence relation [23,24], and is therefore the basic computational equality on Nuprl terms. Internally it becomes the equality on the partial type **Base** of all untyped Nuprl terms, both programs and data. The canonical values of this type are the terminating terms, the values of the type system.

Nuprl's logic is defined on top of this computation system. It is an extensional Constructive Type Theory (CTT) [16] which relies on ternary partial equivalence relations that express when two terms are equal in a type. For example, the type $1 + 1 =_{\mathbb{N}} 2$ expresses that $1 + 1$ and 2 are equal natural numbers (we write $x \in T$, for $x =_T x$). Each type is defined by such a relation.

Over the past two decades much progress has been made to enrich Nuprl and make it a practical programming language as well as a logical system in which one can verify properties of Nuprl programs [35,17,18,21,25,26]. During that period, Nuprl's theory was extended with, e.g., intersection types, union types, partial types, a call-by-value operator, rules to reason about computation, and in particular rules about the fixpoint operator. Recently, we have extended Nuprl

¹ Crary gave a denotational semantics for an ML dialect using partial types.

with new operators called *canonical form tests* (similar to Lisp and Scheme’s type predicates) so that programs can distinguish between primitive canonical operators such as the pair or lambda constructors, and we have developed new ways to reason about these new constructs. This gives us more tools to program in Nuprl and reason about these programs.

Nuprl’s intersection and partial types add expressive power. They allow us to reason about a larger class of practical programs and express more program properties. However, using typed equivalences to transform programs can be unnecessarily complex because programs are not annotated with types and both type checking and type inference are undecidable in Nuprl. Instead, we can reason about untyped program equivalences (e.g., between partial functions), which are easier to use because they only require trivial type reasoning. Such equivalences are highly useful for program transformation such as program optimization.

Using untyped reasoning, we have proved many bisimulations involving data structures such as lists. We have also used these techniques in our work on process synthesis [11,33], where processes are defined as recursive functions of a co-recursive type. Our synthesized processes were initially too slow to be used in industrial strength systems. In response to that issue, we have developed a proof methodology to simplify and optimize them. We have applied that methodology to various synthesized consensus protocols such as 2/3-consensus [14] or Multi-Paxos [28], and observed a significant speed-up. These synthesized consensus protocols have successfully been used in a replicated database [34]. This paper illustrates these proof techniques using a simple running example: appending the empty list to a term. It then illustrates their use to optimize distributed processes synthesized from protocol specifications.

Finally, being able to reason about Nuprl’s programming language directly in Nuprl is another step towards a longstanding goal of building a correct-by-construction, workable logical programming environment [22]. An obvious question is then, could we build a verified compiler for Nuprl in Nuprl that generates reasonably fast code? Modern proof assistants that implement constructive type theories such as Coq [9,1], Isabelle [8,7], or Nuprl rely on unverified compilers. Even though the programs they generate, e.g., by extraction from proofs, are *correct-by-construction*, one could argue whether the machine code obtained after compilation is still correct. Thus, we would like these proof assistants to be expressive enough to program and verify optimized compilers for their underlying programming languages, and to program these proof assistants in themselves.

The contributions of this paper are as follows: (1) we introduce new formal untyped reasoning techniques for proving bisimulations, which expose more of the computation system to formal reasoning; and (2) we apply these techniques to optimize distributed processes.

2 Nuprl’s Programming Language

2.1 Syntax

Nuprl is defined on top of an applied lazy untyped λ -calculus. Fig. 1 introduces a subset of this language, where \underline{n} ranges over integers. Because this language

$v ::= \underline{n}$	(integer)	$\lambda x.t$	(lambda)
$\langle t_1, t_2 \rangle$	(pair)	$\text{inl}(t)$	(left injection)
$\text{inr}(t)$	(right injection)	Ax	(axiom)
$t ::= x$	(variable)	$\text{isaxiom}(\boxed{t_1}, t_2, t_3)$	(isaxiom)
v	(value)	$\text{ispair}(\boxed{t_1}, t_2, t_3)$	(ispair)
$\boxed{t_1} t_2$	(application)	$\text{islambd}(\boxed{t_1}, t_2, t_3)$	(islambd)
$\text{fix}(\boxed{t})$	(fixpoint)	$\text{isinl}(\boxed{t_1}, t_2, t_3)$	(isinl)
$\text{let } x ::= \boxed{t_1} \text{ in } t_2$	(call-by-value)	$\text{isinr}(\boxed{t_1}, t_2, t_3)$	(isinr)
$\text{let } x ::= \boxed{t_1} \text{ in } t_2$	(call-by-valueall)	$\text{isint}(\boxed{t_1}, t_2, t_3)$	(isint)
$\text{let } x, y = \boxed{t_1} \text{ in } t_2$	(spread)		
$\text{case } \boxed{t_1} \text{ of } \text{inl}(x) \Rightarrow t_2 \mid \text{inr}(y) \Rightarrow t_3$	(decide)		
$\text{if } \boxed{t_1} = \boxed{t_2} \text{ then } t_3 \text{ else } t_4$	(integer equality)		

Fig. 1: Syntax of Nuprl's programming language

is lazy, its values² (or canonical forms) are either integers, lambda abstractions, pairs, injections, or Ax . The canonical form Ax (sometimes written as \star) is the unique canonical inhabitant of true propositions that do not have any nontrivial computational meaning in CTT, such as $0 =_{\mathbb{N}} 0$ which is an axiom of the logic. Non-canonical terms (non-values) have arguments that are said to be *principal*. These principal arguments indicate which subterms of a non-canonical term have to be evaluated before checking whether the term itself is a redex or not. Principal arguments of terms are marked with boxes in the above table. In the rest of this paper, variables will be obvious from the context (we often use x and y such as in Fig. 1), we use v for values, and the other letters can be any term. When it is more readable we write $t_1(t_2)$ instead of $t_1 t_2$.

As mentioned above, we have recently added new primitive operators to Nuprl: the canonical form tests such as ispair . Adding these primitive forms was a design decision we made to distinguish between canonical forms (e.g., see list_ind 's definition below) and therefore exploit Howe's bisimulation even further. Our experiments with them have proven to be very fruitful.

Let us now define a few useful abstractions: let \perp (bottom) be $\text{fix}(\lambda x.x)$, let $\pi_1(t)$ be $(\text{let } x, y = t \text{ in } x)$, and let $\pi_2(t)$ be $(\text{let } x, y = t \text{ in } y)$.

Free and bound variables are defined as usual. We write $t[x \setminus u]$ (and more generally $t[x_1 \setminus u_1; \dots; x_n \setminus u_n]$) for the term t in which all the free occurrences of x have been replaced by u . Terms are identified up to alpha-equivalence.

Let Top be the following type: for all closed terms t_1 and t_2 , $t_1 =_{\text{Top}} t_2$. Top 's equality is trivial because it identifies all elements. This type is especially useful to assign types to terms in contexts where their structure or behavior is irrelevant. When discussing types it is important to remember that a type is an equivalence relation on a set of terms and not simply a set of terms. Type A is a subtype of type B (written $A \sqsubseteq B$) if $x =_A y$ implies $x =_B y$. This means not only that every term in A is also in B , but that equality in A refines equality in B . Hence, $T \sqsubseteq \text{Top}$ for every type T . Sec. 3.1 discusses the type Base , which

² The only other values currently in Nuprl are tokens, atoms, and types, but more values can be added because the system is open-ended.

Core calculus:			
$(\lambda x.F) a$			$\rightarrow F[x \setminus a]$
$\mathbf{let} x, y = \langle t_1, t_2 \rangle \mathbf{in} F$			$\rightarrow F[x \setminus t_1; y \setminus t_2]$
$\mathbf{case} \mathbf{inl}(t) \mathbf{of} \mathbf{inl}(x) \Rightarrow F \mid \mathbf{inr}(y) \Rightarrow G$			$\rightarrow F[x \setminus t]$
$\mathbf{case} \mathbf{inr}(t) \mathbf{of} \mathbf{inl}(x) \Rightarrow F \mid \mathbf{inr}(y) \Rightarrow G$			$\rightarrow G[y \setminus t]$
$\mathbf{if} \underline{n}_1 = \underline{n}_2 \mathbf{then} t_1 \mathbf{else} t_2$			$\rightarrow t_1, \text{ if } \underline{n}_1 = \underline{n}_2$
$\mathbf{if} \underline{n}_1 \neq \underline{n}_2 \mathbf{then} t_1 \mathbf{else} t_2$			$\rightarrow t_2, \text{ if } \underline{n}_1 \neq \underline{n}_2$
$\mathbf{fix}(t)$			$\rightarrow t \mathbf{fix}(t)$
$\mathbf{let} x := t_1 \mathbf{in} t_2$			$\rightarrow t_2[x \setminus t_1], \text{ if } t_1 \text{ is a value}$
Canonical form tests:			
$\mathbf{ispair}(\langle t, t' \rangle, t_1, t_2) \rightarrow t_1$	$\mathbf{ispair}(v, t_1, t_2) \rightarrow t_2, \text{ if } v \text{ is not a pair}$		
$\mathbf{isaxiom}(\mathbf{Ax}, t_1, t_2) \rightarrow t_1$	$\mathbf{isaxiom}(v, t_1, t_2) \rightarrow t_2, \text{ if } v \text{ is not axiom}$		
$\mathbf{islambd}(\lambda x.t, t_1, t_2) \rightarrow t_1$	$\mathbf{islambd}(v, t_1, t_2) \rightarrow t_2, \text{ if } v \text{ is not a lambda}$		
$\mathbf{isinl}(\mathbf{inl}(t), t_1, t_2) \rightarrow t_1$	$\mathbf{isinl}(v, t_1, t_2) \rightarrow t_2, \text{ if } v \text{ is not a left injection}$		
$\mathbf{isinr}(\mathbf{inr}(t), t_1, t_2) \rightarrow t_1$	$\mathbf{isinr}(v, t_1, t_2) \rightarrow t_2, \text{ if } v \text{ is not a right injection}$		
$\mathbf{isint}(\underline{n}, t_1, t_2) \rightarrow t_1$	$\mathbf{isint}(v, t_1, t_2) \rightarrow t_2, \text{ if } v \text{ is not an integer}$		

Fig. 2: Nuprl's operational semantics

contains all Nuprl terms, but does not have this property (i.e. not every type T is a subtype of \mathbf{Base}^3) because equality on \mathbf{Base} is Howe's bisimulation relation.

2.2 Operational Semantics

Fig. 2 presents some of Nuprl's reduction rules. This figure does not show the reduction rule for the call-by-valueall operator because it is slightly more complicated. This operator is like call-by-value but continues recursively evaluating subterms of pairs and injections.⁴

At any point in a computation, either a value is produced, or the computation is stuck, or we can take another step. For example, $(\mathbf{let} x, y = \mathbf{Ax} \mathbf{in} F)$ is a meaningless term that cannot evaluate further. It is stuck on the wrong kind of principal argument: \mathbf{Ax} instead of a pair. Using the proof techniques presented below, in Sec. 4.3 we prove that this term is computationally equivalent to \perp . We can prove such results using \mathbf{ispair} and $\mathbf{isaxiom}$, and do not know of any other way discussed in the literature to accomplish this. Intuitively, we prove this lemma using the fact that $\mathbf{isaxiom}$ can compute to different values depending on whether its first argument computes to \mathbf{Ax} or not. For example, $\mathbf{isaxiom}(t, 0, 1)$ reduces to 0 if t is \mathbf{Ax} and to 1 if t is, e.g., a pair. Note that even though they are computationally equal, $(\mathbf{let} x, y = \mathbf{Ax} \mathbf{in} F)$ and \perp are fundamentally different in the sense that one could potentially detect whether a term is stuck (by slightly modifying our destructors such as *spread* or *decide*), but one cannot detect whether a term diverges or not.

³ Being extentional, function types are in general not subtypes of \mathbf{Base} .

⁴ The call-by-valueall operator is similar to a restricted form of Haskell's *deepseq* operator. It can be defined using the other primitive operators (see the expanded version of this article at <http://www.nuprl.org/Publications/>), but for simplicity reasons we introduce it as a primitive in this paper.

2.3 Datatypes

Booleans As usual, we define Booleans using the disjoint union type as follows: $\mathbb{B} = \mathbf{Unit} + \mathbf{Unit}$. The \mathbf{Unit} type is defined as $0 =_{\mathbb{Z}} 0$ and therefore \mathbf{Ax} is its only inhabitant (up to computation). We define the Boolean true \mathbf{tt} as $\mathbf{inl}(\mathbf{Ax})$, and the Boolean false \mathbf{ff} as $\mathbf{inr}(\mathbf{Ax})$. Using the *decide* operator we define a conditional operator as follows: $\mathbf{if } t_1 \mathbf{ then } t_2 \mathbf{ else } t_3 = \mathbf{case } t_1 \mathbf{ of } \mathbf{inl}(x) \Rightarrow t_2 \mid \mathbf{inr}(x) \Rightarrow t_3$.

Lists We define lists as follows using *Nuprl*'s union type [26] and recursive type [32] that allows one to build inductive types:⁵ $\mathbf{List}(T) = \mathbf{rec}(L.\mathbf{Unit} \cup T \times L)$. The type constructor \cup creates the union of two types, not the disjoint union. The members of $A \cup B$ are members of A or B , not injections of them. A list is either a member of \mathbf{Unit} , i.e., \mathbf{Ax} , or a pair. The empty list \mathbf{nil} is defined as \mathbf{Ax} , and the cons operation, denoted by \bullet , as the pair constructor. We can distinguish an empty list and a non empty list because \mathbf{Unit} and the product type are disjoint. Using \mathbf{fix} , we define the following “list induction” operator:

$$\begin{aligned} \mathbf{list_ind}(L, b, f) = \\ \mathbf{fix}(\lambda F.\lambda L.\mathbf{ispair}(L, \mathbf{let } h, t = L \mathbf{ in } f \ h \ t \ (F \ t), \mathbf{isaxiom}(L, b, \perp))) \ L \end{aligned}$$

To define such a function that takes a list as input, we need to be able to test whether it is a pair or \mathbf{Ax} . If we were to use the *spread* operator, we could destruct pairs, but computations would get stuck on \mathbf{Ax} which we use to represent the empty list. Therefore, we need an operator such as the \mathbf{ispair} canonical form test which allows us to perform two different computations depending on whether its first argument computes to a pair or not. Note that if $\mathbf{list_ind}$'s first argument does not compute to a pair or to \mathbf{Ax} , then the term diverges as opposed to returning an arbitrary value. This is necessary to prove untyped equivalences between list operations. We define the append and map operations as follows:

$$\begin{aligned} t_1 \ @ \ t_2 &= \mathbf{list_ind}(t_1, t_2, \lambda h.\lambda t.\lambda r.h \bullet r) \\ \mathbf{map}(f, t) &= \mathbf{list_ind}(t, \mathbf{nil}, \lambda h.\lambda t.\lambda r.(f \ h) \bullet r) \end{aligned}$$

3 Computational Equivalence

3.1 Simulations and Bisimulations

Howe [23,24] defined the simulation or approximation relation \leq using the following co-inductive rule: $t_1 \leq t_2$ if and only if (if t_1 computes to a canonical form $\Theta(u_1, \dots, u_n)$ of the language defined in Sec. 2.1, then t_2 computes to a canonical form $\Theta(u'_1, \dots, u'_n)$ such that for all $i \in \{1, \dots, n\}$, $u_i \leq u'_i$). We say that t_1 approximates t_2 or that t_2 simulates t_1 . This relation is reflexive (w.r.t. the terms defined in Sec. 2.1) and transitive. Howe then defined the bisimulation relation \sim as the symmetric closure of \leq (i.e., $t_1 \sim t_2$ iff $t_1 \leq t_2$ and $t_2 \leq t_1$), and proved that \leq and \sim are congruences w.r.t. *Nuprl*'s computation system.⁶

⁵ This new definition of lists replaces the one from *Nuprl*'s book [16] where lists are considered as primitive objects. Using *Nuprl*'s replay functionality, we were able to successfully replay the entire *Nuprl* library using this new definition of lists.

⁶ Howe proved that \sim is a congruence w.r.t. a lazy computation system by proving that all the operators of the system satisfy a property called *extensionality*. The expanded version of this article proves that the new operators introduced in this paper satisfy that property.

The following *context property* follows from the fact that \sim is a congruence:

$$\forall i : \{1, \dots, n\}. t_i \leq u_i \Rightarrow G[x_1 \setminus t_1; \dots; x_n \setminus t_n] \leq G[x_1 \setminus u_1; \dots; x_n \setminus u_n]$$

Howe's bisimulation relation respects computation, i.e., if $t_1 \sim t_2$ then (1) t_1 computes to a value iff t_2 computes to a value, and (2) if t_1 computes to a value v_1 then t_2 computes to a value v_2 with same outer operator such that $v_1 \sim v_2$.

Because \perp does not compute to a canonical form, by definition $\perp \leq t$ is true for any term t ; hence for example $\langle \mathbf{Ax}, \perp \rangle \leq \langle \mathbf{Ax}, \mathbf{Ax} \rangle$. Similarly, because \mathbf{Ax} is not a pair, $(\mathbf{let } x, y = \mathbf{Ax} \mathbf{ in } x)$ does not compute to a canonical form, and by definition, $\mathbf{let } x, y = \mathbf{Ax} \mathbf{ in } x \leq t$ is true for any term t (we prove $\mathbf{let } x, y = \mathbf{Ax} \mathbf{ in } F \sim \perp$ in Sec. 4.3). However, $\mathbf{Ax} \leq \perp$ is not true because \perp diverges while \mathbf{Ax} is a value; hence $\langle \mathbf{Ax}, \mathbf{Ax} \rangle \leq \langle \mathbf{Ax}, \perp \rangle$ is not true either.

Let us write $\mathbf{halts}(t)$ if t reduces to a value—we say that t converges. We can define convergence using call-by-value because the call-by-value operator $(\mathbf{let } x := t_1 \mathbf{ in } t_2)$ first evaluates t_1 . The term t_1 converges if and only if the term $(\mathbf{let } x := t_1 \mathbf{ in } \mathbf{Ax})$ evaluates to \mathbf{Ax} . So we simply define $\mathbf{halts}(t)$ to be the simulation $\mathbf{Ax} \leq (\mathbf{let } x := t \mathbf{ in } \mathbf{Ax})$. Because \mathbf{Ax} is a canonical value then $\mathbf{Ax} \leq (\mathbf{let } x := t \mathbf{ in } \mathbf{Ax})$ is true if and only if $(\mathbf{let } x := t \mathbf{ in } \mathbf{Ax})$ computes to \mathbf{Ax} , i.e., if and only if t computes to a value.

Constable and Smith [35,17] introduced partial types to reason about computations that might not halt. For any type T , the partial type \overline{T} contains all members of T as well as all divergent terms, and has the following equality: two terms are equal in \overline{T} if they have the same convergence behavior (i.e., either neither computes to a value or both compute to a value), and when they converge, they are equal in T . An important partial type is $\mathbf{Base} = \overline{\mathbf{Value}}$ where \mathbf{Value} is the type of all closed canonical terms of the computation system with \sim as its equality. Because \mathbf{Base} is a partial type, it contains converging as well as diverging terms, and equal terms have the same convergence behavior.

3.2 Simple Facts About Lists

Sec. 4 proves that for all terms f and t in \mathbf{Top} , $\mathbf{map}(f, t) @ \mathbf{nil} \sim \mathbf{map}(f, t)$.

If t is a list, the first expression $(\mathbf{map}(f, t) @ \mathbf{nil})$ requires two passes over the list t while the second expression $(\mathbf{map}(f, t))$ requires only one. This simple bisimulation will be our running example to illustrate the techniques we use to optimize our distributed processes (discussed in Sec. 5).

Note that this lemma would be easy to prove by induction on the list t if we were using the list type instead of \mathbf{Top} . However, we might need to instantiate t with a term for which it would be non-trivial to prove that it is a list because \mathbf{Nuprl} is based on an extension of the *untyped* λ -calculus and type inference and type checking are *undecidable*. In addition, if we were to use a typed equality (instead of \sim) for substitution in some context, then we would also have to prove that the context is functional over the type of the equality. That is, to rewrite in the term $C[t]$ of type B using $t =_A u$, we have to prove that $\lambda z. C[z]$ is of type $A \rightarrow B$. Moreover, the above equivalence is indeed true for any term t , e.g., it is true when t is a stream.

Note that it is not true that for all terms t , $t @ \text{nil} \sim t$. For example, by definition of $@$, $(\lambda x.x) @ \text{nil} \sim \perp$. However, the bisimulation $\lambda x.x \sim \perp$ is not true because the simulation $\lambda x.x \leq \perp$ is not true. This shows that there are some terms t for which $t \leq t @ \text{nil}$ does not hold.⁷ However, Sec. 4 proves that for all terms t , $t @ \text{nil} \leq t$. A corollary of that lemma is that $\text{map}(f, t) @ \text{nil} \leq \text{map}(f, t)$.

4 Proof Techniques

This section presents three proof techniques we use to prove bisimulations: Cray's least upper bound property [18], patterns of reasoning regarding our new canonical form tests, and patterns of reasoning regarding our `halts` operator. It also presents three derived proof techniques called lifting, normalization, and strictness. Using these techniques, we prove $\text{map}(f, t) @ \text{nil} \sim \text{map}(f, t)$, and in Sec. 5, we optimize distributed processes.

4.1 Least Upper Bound Property

Using the properties of \leq and that $\text{fix}(f) = f \text{ fix}(f)$, it is easy to prove by induction on n that $\forall n : \mathbb{N}. f^n(\perp) \leq \text{fix}(f)$ [18]. So $\text{fix}(f)$ is an upper bound of its approximations. The least upper bound property [18, Theorem 5.9] is:

Rule [least-upper-bound]. $\forall n : \mathbb{N}. G(f^n(\perp)) \leq t \Rightarrow G(\text{fix}(f)) \leq t$.

4.2 Canonical Form Tests

In order to reason about its programs, we gave Nuprl the ability to reason about the *canonical form tests* such as `ispair`, `isaxiom`, etc.⁸ These effective operations on `Base` allow us to reason in the programming language, where in the past we resorted to reflection in the logic [6].

Membership Rules

Rule [ispair-member]. *To prove that `ispair`(t_1, t_2, t_3) $\in T$, it is enough to prove `halts`(t_1), and that both t_2 and t_3 are members of T .*

We introduce similar rules for the other canonical form tests. Using this rule we can trivially prove the following fact:

Lemma 1. *For all terms t in `Base`, if `halts`(t) then `ispair`(t, tt, ff) $\in \mathbb{B}$.*

The same is true for the other tests. Using these facts, we can, e.g., decide whether a converging term is a pair or not.

Semi-decision Rules Depending on how `ispair` computes we can deduce various pieces of information. If we know that `ispair`(t_1, t_2, t_3) always computes to t_2 and cannot compute to t_3 then we know that t_1 is a pair. If we know that `ispair`(t_1, t_2, t_3) always computes to t_3 and cannot compute to t_2 then we know that t_1 is not a pair. These properties are captured by the following two rules:

⁷ The expanded version of this article provides a characterization of the terms that satisfy that property.

⁸ The proofs that the rules introduced in this section are valid w.r.t. Allen's PER (Partial Equivalence Relations) semantics [2,3] are presented in the expanded version of this article.

$\vdash \forall F:\text{Top}. (\text{let } x, y = Ax \text{ in } F[x; y] \sim \text{bottom}())$
\mid
BY (SqReasoning
THEN Assert $\lceil (\text{if } Ax \text{ is a pair then } 0 \text{ otherwise } 1) = 1 \rceil$.
THEN (Reduce 0 THEN Auto THEN AutoPairEta [2;1] (-1))

Fig. 3: Computational equivalence between \perp and a stuck term

Rule [ispair]. *To prove $t \in \text{Top} \times \text{Top}$ (i.e., t is a pair), it is enough to prove $\text{ispair}(t, \text{inl}(a), \text{inr}(b)) \sim \text{inl}(a)$ for some terms a and b .*

Rule [not-ispair]. *To prove $\text{ispair}(t_1, t_2, t_3) \sim t_3$, it is enough to prove that $\text{ispair}(t_1, \text{inl}(a), \text{inr}(b)) \sim \text{inr}(b)$ for some terms a and b .*

We introduce similar rules for the other canonical form tests. Using these rules we can prove such results as (similar results are true for the other tests):

Lemma 2. *For all terms t, a, b in Base, if $\text{halts}(t)$ then $t \sim \langle \pi_1(t), \pi_2(t) \rangle \vee \text{ispair}(t, a, b) \sim b$.*

Proof. By Lemma 1, $\text{ispair}(t, \text{tt}, \text{ff}) \in \mathbb{B}$. Therefore, either $\text{ispair}(t, \text{tt}, \text{ff}) \sim \text{tt}$ or $\text{ispair}(t, \text{tt}, \text{ff}) \sim \text{ff}$ (this is true for any Boolean). If $\text{ispair}(t, \text{tt}, \text{ff}) \sim \text{tt}$ then using rule [ispair] we obtain that t is a pair and therefore $t \sim \langle \pi_1(t), \pi_2(t) \rangle$. If $\text{ispair}(t, \text{tt}, \text{ff}) \sim \text{ff}$ then using rule [not-ispair] we obtain that $\text{ispair}(t, a, b) \sim b$. \square

4.3 Convergence

Rule [convergence]. *To prove $t_1 \leq t_2$, one can assume $\text{halts}(t_1)$.*

This rule follows directly from \leq 's definition. For example, to prove $\text{let } x, y = p \text{ in } F \leq \text{let } x, y = q \text{ in } G$, one can assume that $\text{halts}(\text{let } x, y = p \text{ in } F)$.

Nuprl also has rules to reason about $\text{halts}(t)$. If a non-canonical term converges, then its principal arguments have to converge to the appropriate canonical forms as presented in Fig 2. For example the following two rules follow from the operational semantics of *spread* and *ispair* (we have similar rules for the other non-canonical operators):

Rule [halt-spread]. *If $\text{halts}(\text{let } x, y = p \text{ in } F)$ then p computes to a pair.*

Rule [halt-ispair]. *If $\text{halts}(\text{ispair}(t_1, t_2, t_3))$ then $\text{halts}(t_1)$.*

Let us go back to the example presented in Sec. 2.2. We now have enough tools to prove the following lemma:

Lemma 3. *For all terms F in Top, $\text{let } x, y = Ax \text{ in } F \sim \perp$*

Proof. Fig 3 presents our Nuprl proof of that fact. That proof goes as follows: By definition of \sim , we have to prove $\text{let } x, y = Ax \text{ in } F \leq \perp$ and $\perp \leq \text{let } x, y = Ax \text{ in } F$. The second simulation is trivial. Let us prove the first one. Using [convergence], we can assume $\text{halts}(\text{let } x, y = Ax \text{ in } F)$ and using [halt-spread], that Ax is a pair. This reasoning is done by our SqReasoning tactic. Finally, the term $\text{ispair}(Ax, 0, 1)$ computes to 1, and because we deduced that Ax is a pair, it also reduces to 0, and we have an absurdity. \square

4.4 Lifting

Now we describe the following derived proof techniques: lifting, normalization (see Sec. 4.5 below), and strictness (see Sec. 4.6 below) which are used in Sec. 4.7 below. Lifting transforms a term t into t' such that $t \sim t'$ and such that t' has a smaller path to the principal argument of a subterm of t . Let us now provide a few examples. The following bisimulation specifies a lifting operation, where the path to p is shorter in the second term than in the first term:

Lemma 4. *For all terms F and G in Top:*

$$\text{let } c, d = (\text{let } a, b = p \text{ in } F) \text{ in } G \sim \text{let } a, b = p \text{ in } (\text{let } c, d = F \text{ in } G)$$

Proof. To prove that bisimulation, we have to prove that the first term simulates the second one and vice versa. Let us prove that the second one simulates the first one (the other direction is similar), i.e., $\text{let } c, d = (\text{let } a, b = p \text{ in } F) \text{ in } G \leq \text{let } a, b = p \text{ in } (\text{let } c, d = F \text{ in } G)$. Using [convergence], we can assume $\text{halts}(\text{let } c, d = (\text{let } a, b = p \text{ in } F) \text{ in } G)$, from which, using [halt-spread] twice, we obtain that p is a pair. More precisely, we can prove that p is the pair $\langle \pi_1(p), \pi_2(p) \rangle$. By replacing p by $\langle \pi_1(p), \pi_2(p) \rangle$ in the above simulation, and by reducing both sides, we obtain $\text{let } c, d = F[a \setminus \pi_1(p); b \setminus \pi_2(p)] \text{ in } G \leq \text{let } c, d = F[a \setminus \pi_1(p); b \setminus \pi_2(p)] \text{ in } G$, which is true by reflexivity of \leq . \square

Using this lemma, one can, e.g., derive the following chain of rewrites:

$$\begin{aligned} & \text{let } a, b = (\text{let } c, d = p \text{ in } \langle c, d \rangle) \text{ in } F \\ & \sim \text{let } c, d = p \text{ in } (\text{let } a, b = \langle c, d \rangle \text{ in } F) \\ & \sim \text{let } c, d = p \text{ in } F[a \setminus c; b \setminus d] \end{aligned}$$

The following bisimulation specifies another lifting operation where the path to t_1 is shorter in the second term than in the first one:

Lemma 5. *For all terms t_1, t_2, t_3, t_4 , and t_5 in Top:*

$$\begin{aligned} & \text{ispair}(\text{ispair}(t_1, t_2, t_3), t_4, t_5) \\ & \sim \text{ispair}(t_1, \text{ispair}(t_2, t_4, t_5), \text{ispair}(t_3, t_4, t_5)) \end{aligned}$$

The proof of this is similar to the proof of Lemma 4. Because lifting does not always result in a smaller term it must therefore be used in a controlled way.

4.5 Normalization

Normalization allows one to make use of the information given by destructors such as spread or decide, i.e., that some terms are forced to be pairs or injections by the computation system. Normalization achieves some kind of common subexpression elimination, which is a standard optimization technique. For example, the next lemma says that the expression on left-hand-side has a value if and only if p (which can be an arbitrary large term) is a pair, and more precisely in F it has to be the pair $\langle a, b \rangle$:

Lemma 6. *For all terms p and F in Top:*

$$\text{let } x, y = p \text{ in } F[z \setminus p] \sim \text{let } x, y = p \text{ in } F[z \setminus \langle x, y \rangle]$$

The proof of this is similar to the proof of Lemma 4.

4.6 Strictness

Strictness says that if \perp is one of the principal arguments of a term then this term is computationally equal to \perp . For example we proved the following lemma:

Lemma 7. *For all terms F in Top , $(\text{let } x, y = \perp \text{ in } F) \sim \perp$.*

The proof of this is similar to the proof of Lemma 4. Intuitively, such lemmas are true because to evaluate a non-canonical term, one has to evaluate its principal arguments. If one of these principal arguments is \perp , then the computation diverges. Therefore, the entire term is computationally equal to \perp .

4.7 Back To Our List Example

As explained in Sec. 3.2, to prove $\text{map}(f, t) @ \text{nil} \sim \text{map}(f, t)$, we first prove the following lemma:

Lemma 8. *For all terms t in Top , $t @ \text{nil} \leq t$.*

Proof. Because $@$ is defined using fix (see Sec. 2.3), we prove that lemma using the [least-upper-bound] rule (see Sec.4.1). We now have to prove that any approximation of the fixpoint is simulated by t . Let

$$F = \lambda F. \lambda L. \text{ispair}(L, \text{let } x, y = L \text{ in } x \bullet (F y), \text{isaxiom}(L, \text{nil}, \perp))$$

We have $(t @ \text{nil}) = (\text{fix}(F) t)$ by definition of append and beta-reduction. We have to prove that for all natural numbers n , and for all terms t ,

$$F^n \perp t \leq t$$

which we prove by induction on n . The base case boils down to proving that $\perp t \leq t$ which is true using strictness. In the interesting induction case, assuming that for all terms t , $F^{n-1} \perp t \leq t$, we have to prove $F (F^{n-1} \perp) t \leq t$, i.e.,

$$\text{ispair}(t, \text{let } x, y = t \text{ in } x \bullet ((F^{n-1} \perp) y), \text{isaxiom}(t, \text{nil}, \perp)) \leq t \quad (1)$$

Let P be $\text{ispair}(t, \text{let } x, y = t \text{ in } x \bullet ((F^{n-1} \perp) y), \text{isaxiom}(t, \text{nil}, \perp))$. Using [convergence], we can assume $\text{halts}(P)$. Using [halt-ispair], we obtain $\text{halts}(t)$. By Lemma 2, we get $t \sim \langle \pi_1(t), \pi_2(t) \rangle$ or $P \sim \text{isaxiom}(t, \text{nil}, \perp)$.

If $t \sim \langle \pi_1(t), \pi_2(t) \rangle$, we have to prove the following simulation obtained from simulation 1 by replacing t by $\langle \pi_1(t), \pi_2(t) \rangle$ and by reducing:

$$\pi_1(t) \bullet ((F^{n-1} \perp) \pi_2(t)) \leq \langle \pi_1(t), \pi_2(t) \rangle$$

Because the cons operator is defined as the pair constructor, by the context property it remains to prove the following simulation, which is true by induction hypothesis: $((F^{n-1} \perp) \pi_2(t)) \leq \pi_2(t)$.

If $P \sim \text{isaxiom}(t, \text{nil}, \perp)$, we have to prove $\text{isaxiom}(t, \text{nil}, \perp) \leq t$. Using the version of Lemma 2 for isaxiom , we obtain $t \sim \text{Ax}$ or $\text{isaxiom}(t, \text{nil}, \perp) \sim \perp$. Both cases are trivial: in the first case we have to prove $\text{Ax} \leq \text{Ax}$ and in the second we have to prove $\perp \leq t$. \square

Let us now prove the lemma we set out to prove in Sec. 3.2:

Lemma 9. *For all terms t and f in Top , $\text{map}(f, t) @ \text{nil} \sim \text{map}(f, t)$.*

Proof. By definition of \sim , we have to prove $\text{map}(f, t) @ \text{nil} \leq \text{map}(f, t)$ (which is true by Lemma 8), and $\text{map}(f, t) \leq \text{map}(f, t) @ \text{nil}$. Because map is a fixpoint, we can prove the latter using the [least-upper-bound] rule. Let

$$F = \lambda F. \lambda L. \text{ispair}(L, \text{let } x, y = L \text{ in } (f x) \bullet (F y), \text{isaxiom}(L, \text{nil}, \perp))$$

We then have to prove that for all natural numbers n and for all terms f and t ,

$$F^n \perp t \leq \text{map}(f, t) @ \text{nil}$$

which we prove by induction on n . Once again, the base case is trivial. Assume that for all terms t , $F^{n-1} \perp t \leq \text{map}(f, t) @ \text{nil}$, we have to prove that $F(F^{n-1} \perp) t \leq \text{map}(f, t) @ \text{nil}$, i.e., we have to prove the following simulation:

$$\begin{aligned} & \text{ispair}(t, \text{let } x, y = t \text{ in } (f x) \bullet ((F^{n-1} \perp) y), \text{isaxiom}(t, \text{nil}, \perp)) & (2) \\ & \leq \text{map}(f, t) @ \text{nil} \end{aligned}$$

Let $P = \text{ispair}(t, \text{let } x, y = t \text{ in } (f x) \bullet \text{map}(f, y), \text{isaxiom}(t, \text{nil}, \perp))$, which is $\text{map}(f, t)$ unfolded once. We obtain the following sequence of bisimulations by unfolding the definitions of map and $@$ in $(\text{map}(f, t) @ \text{nil})$:

$$\begin{aligned} & \text{map}(f, t) @ \text{nil} \sim P @ \text{nil} \\ & \sim \text{ispair}(P, \text{let } x, y = t \text{ in } x \bullet (y @ \text{nil}), \text{isaxiom}(P, \text{nil}, \perp)) \end{aligned}$$

Using lifting (Lemma 5) and normalization, we obtain the following bisimulation:

$$\begin{aligned} & \text{map}(f, t) @ \text{nil} \\ & \sim \text{ispair}(t, \text{let } x, y = t \text{ in } (f x) \bullet (\text{map}(f, y) @ \text{nil}), \text{isaxiom}(t, \text{nil}, \perp)) \end{aligned}$$

Therefore, given that we have to prove simulation 2, it means that we have to prove the following simulation:

$$\begin{aligned} & \text{ispair}(t, \text{let } x, y = t \text{ in } (f x) \bullet ((F^{n-1} \perp) y), \text{isaxiom}(t, \text{nil}, \perp)) \\ & \leq \text{ispair}(t, \text{let } x, y = t \text{ in } (f x) \bullet (\text{map}(f, y) @ \text{nil}), \text{isaxiom}(t, \text{nil}, \perp)) \end{aligned}$$

which is true by induction hypothesis and the context property. \square

5 Process Optimization

Nuprl implements a Logic of Events (LoE) [10,12,13] to specify and reason about distributed programs, as well as a General Process Model (GPM) [11] to implement them. We have proved a direct relationship between some LoE combinators and some GPM combinators. This allows us to automatically generate processes that are guaranteed to satisfy the logical specifications of LoE.

Using the proof techniques presented in the above section, we were able to optimize many automatically generated GPM processes. For example, we optimized our synthesized version of Paxos, which is used by the ShadowDB replicated database [34]. Because our synthesized Paxos was initially too slow, it was

only used to handle database failures, which are critical to handle correctly but are not frequent. When a failure occurs, Paxos ensures that the replicas agree on the next set of replicas. We can now also use Paxos to consistently order the transactions of the replicated databases. Initially, our synthesized code could only handle one transaction every few seconds. Thanks to our automatic optimizer, the code we synthesize is now about an order of magnitude faster. Our goal is to be able to handle several thousands of transactions per second. Even though we have not yet reached that goal, this work is already an encouraging first step towards generating fast correct-by-construction code.

A GPM process is modeled as a function that takes inputs and computes a new process as well as outputs. For distributed programs based on message passing, these inputs and outputs are messages. Formally, a process that takes inputs of type A , and outputs elements of type B , is an element of (a variant of) the following co-recursive type:

$$\text{corec}(\lambda P.A \rightarrow P \times \text{Bag}(B))$$

where `corec` is defined as follows:

$$\text{corec}(G) = \cap n : \mathbb{N}.\text{fix}(\lambda P.\lambda n.\text{if } n = 0 \text{ then Top else } G (P (n - 1))) n$$

Note the use of bags, also called multisets, formally defined as quotiented lists. The reason for using that type is outside the scope of this paper. However, let us mention that processes can output more than one element and these elements need not be ordered. In the rest of this paper, we use curly braces to denote specific bag instances. Lists and bags have many similar operations such as: `bmap` the map operation on bags, `bnull` the null operation, `bconcat` the concat operation which flattens bags of bags, and `>>=` the bind operation of the bag monad, defined as $b \gg= f = \text{bconcat}(\text{bmap}(f, b))$. For example, $(\{1; 2; 2; 4\} \gg= \lambda x.\{x; x + 1\}) = \{1; 2; 2; 3; 2; 3; 4; 5\} = \{1; 2; 2; 2; 3; 3; 4; 5\}$.

Many of the GPM combinators are defined using `fix`. Because processes are typically defined using several combinators, fixpoints end up being deeply nested which affects the computational complexity of the processes. Using, among other things, the least upper bound property, we can often reduce the number of fixpoints occurring in processes. This is our main process optimization technique.

Let us now present some GPM combinators. Processes often need to maintain an internal state. Therefore, the combinators defined below will all be of the form `fix(λF.λs.λm.G) init`, where `init` is an initial state, and `G` is a transition function that takes the current state of the process (`s`) and an input (`m`), and generates a new process and some output.

5.1 Combinators

Base Combinator It builds a process that applies a function to its inputs:

$$\text{base}(f) = \text{fix}(\lambda F.\lambda s.\lambda m.\langle F s, f m \rangle) \text{Ax}$$

Base processes are stateless, which is modeled using the term `Ax` as the state of the base combinator.

Composition Combinator It builds a process that applies a function f to the outputs of its sub-component X :

$$f \circ X = \text{fix}(\lambda F. \lambda X. \lambda m. \left(\begin{array}{l} \text{let } X', out = X \ m \ \text{in} \\ \text{let } out' ::= \text{bmap}(f, out) \ \text{in} \\ \langle F \ X', out' \rangle \end{array} \right) X)$$

The state maintained by $f \circ X$ is the state maintained by X . Note that for efficiency issues, we use the call-by-value all operator $::=$ in order to generate the outputs out' .

Buffer Combinator From an initial buffer $init$ and a process X producing transition functions, this combinator builds a process that buffers its outputs:

$$\text{buffer}(X, init) = \text{fix}(\lambda F. \lambda s. \lambda m. \left(\begin{array}{l} \text{let } X, buf = s \ \text{in} \\ \text{let } X', b = X \ m \ \text{in} \\ \text{let } b' ::= b \gg= \lambda f. (buf \gg= f) \ \text{in} \\ \langle F \ \langle X', \text{if } \text{bnull}(b') \ \text{then } buf \ \text{else } b' \rangle, b' \rangle \end{array} \right) \langle X, init \rangle)$$

The state maintained by $\text{buffer}(X, init)$ is the pair of the state maintained by X and its previous outputs (initially $init$).

5.2 Example

The following process uses the three combinators presented above:

$$P = \text{buffer}((\lambda n. \lambda buf. \{n + buf\}) \circ \text{base}(\lambda m. \{m\}), \{0\})$$

This process maintains a state constituted of a single integer, initialized to 0. Its inputs are integers. At any point in time, its state is the sum of all the inputs it has received in the past. Because the combinators used in P are defined as fixpoints, P contains three nested occurrences of fix . We will now show that P is computationally equivalent to the following even simpler process:

$$P' = \text{fix}(\lambda F. \lambda s. \lambda m. \text{let } x ::= m + s \ \text{in} \ \langle F \ x, \{x\} \rangle) 0$$

Using Nuprl's powerful tactic mechanism we automatically generate P' from P , and we automatically prove that $P \sim P'$. Our experiments showed that it takes between 100 and 200 computation steps for P to process a single input while it takes less than 10 computation steps for P' to process a single input.

Standard Form To optimize our processes we take advantage of the fact that many of them are of the following form:

$$\text{process}(n, L, S, R, I) = \text{fix}(\lambda F. \lambda s. \lambda m. \left(\begin{array}{l} \text{let } x_1 ::= L \ s \ m \ 1 \ \text{in} \\ \dots \\ \text{let } x_n ::= L \ s \ m \ n \ x_1 \ \dots \ x_{n-1} \ \text{in} \\ \langle F \ (S \ s \ m \ x_1 \ \dots \ x_n), R \ s \ m \ x_1 \ \dots \ x_n \rangle \end{array} \right) I)$$

where L is a sequence of instructions defined as a function, n is the number of instructions that the process executes on each input, S computes the next state of the process, R computes the outputs, and I is its initial state.

and $\text{map}(f, \text{iterate}(f, x))$ are bisimilar, where $\text{iterate}(f, x)$ is defined in `Nuprl` as $\text{fix}(\lambda F. \lambda x. \langle x, F(f\ x) \rangle) x$. `Nuprl`'s corresponding method to prove such results is the least upper bound property. Gordon proves this result using a co-inductive reasoning, while we prove it by induction on the natural number we obtain by approximating the two fixpoints used to define `map` and `iterate`. Apart from that difference, the resulting proofs are similar in spirit.

Note that we have not yet formally proved that the processes returned by our optimizer have a better complexity than the processes it takes as inputs. Using `Isabelle/HOL`, Aspinall, Beringer, and Momigliano [5] developed an optimization validation technique, based on a proof-carrying code approach, to prove that optimized programs use less resources than the non-optimized versions. Currently, we cannot measure the complexity of programs inside `Nuprl` because if t_1 reduces to t_2 then $t_1 \sim t_2$, and hence we cannot distinguish between them in any context.

We hope to solve this issue by either using some kind of reflection, or introducing a subtype of `Base` where equality would be alpha-equality. Also, in order to enhance the usability of our processes in industrial strength systems, we need to identify and verify other optimizations. As mentioned in Sec. 1, we view this work as a step towards making `Nuprl` a usable programming framework. In the meantime, we have built a `Lisp` translator for our processes.

In the last two decades, much work has been done on compiler verification. See Dave [19] for earlier references. To name a few: Using `Coq`, Leroy has developed and certified a compiler for a C-like language [29]. He generated the compiler using `Coq`'s extraction mechanism to `Caml` code. The compiler is certified thanks to “a machine-checked proof of semantic preservation” [29]. Also using `Coq`, Chlipala [15] developed a verified compiler for an impure functional programming language with references and exceptions that produces code in an idealized assembly language. He proved the correctness of the compiler using a big-step operational semantics. Li [30] designed a verified compiler in `HOL`, from an high-level ML-like programming language implemented in `HOL` to ARM assembly code. Each transformation of the compiler generates a correctness argument along with a piece of code.

Following this line of work, we now would like to tackle the task of building a verified compiler for `Nuprl` in `Nuprl`.

Acknowledgements

We would like to thank our colleagues Professor Robert L. Constable, David Guaspari, and Evan Moran for their helpful criticism.

References

1. The Coq Proof Assistant. <http://coq.inria.fr/>.
2. Stuart F. Allen. A non-type-theoretic definition of martin-löf's types. In *LICS*, pages 215–221. IEEE Computer Society, 1987.
3. Stuart F. Allen. *A Non-Type-Theoretic Semantics for Type-Theoretic Language*. PhD thesis, Cornell University, 1987.
4. Stuart F. Allen, Mark Bickford, Robert L. Constable, Richard Eaton, Christoph Kreitz, Lori Lorigo, and Evan Moran. Innovations in computational type theory using `Nuprl`. *J. Applied Logic*, 4(4):428–469, 2006.

5. David Aspinall, Lennart Beringer, and Alberto Momigliano. Optimisation validation. *Electr. Notes Theor. Comput. Sci.*, 176(3):37–59, 2007.
6. Eli Barzilay. *Implementing Reflection in NUPRL*. PhD thesis, Cornell University, 2006.
7. Stefan Berghofer. Program extraction in simply-typed higher order logic. In *Types for Proofs and Programs, 2nd Int'l Workshop, TYPES*, volume 2646 of *LNCS*, pages 21–38. Springer, 2002.
8. Stefan Berghofer and Tobias Nipkow. Executing higher order logic. In *Types for Proofs and Programs, Int'l Workshop, TYPES*, volume 2277 of *LNCS*, pages 24–40. Springer, 2000.
9. Yves Bertot and Pierre Casteran. *Interactive Theorem Proving and Program Development*. SpringerVerlag, 2004.
10. Mark Bickford. Component specification using event classes. In *Component-Based Software Engineering, 12th Int'l Symp.*, volume 5582 of *LNCS*, pages 140–155. Springer, 2009.
11. Mark Bickford, Robert Constable, and David Guaspari. Generating event logics with higher-order processes as realizers. Technical report, Cornell University, 2010.
12. Mark Bickford and Robert L. Constable. Formal foundations of computer security. In *NATO Science for Peace and Security Series, D: Information and Communication Security*, volume 14, pages 29–52. 2008.
13. Mark Bickford, Robert L. Constable, and Vincent Rahli. Logic of events, a framework to reason about distributed systems. In *Languages for Distributed Algorithms Workshop*, 2012.
14. B. Charron-Bost and A. Schiper. The Heard-Of model: computing in distributed systems with benign failures. *Distributed Computing*, 22(1):49–71, 2009.
15. Adam Chlipala. A verified compiler for an impure functional language. In *37th ACM SIGPLAN-SIGACT Symp. on Principles of Programming Languages*, pages 93–106. ACM, 2010.
16. R. L. Constable, S. F. Allen, H. M. Bromley, W. R. Cleaveland, J. F. Cremer, R. W. Harper, D. J. Howe, T. B. Knoblock, N. P. Mendler, P. Panangaden, J. T. Sasaki, and S. F. Smith. *Implementing mathematics with the Nuprl proof development system*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1986.
17. Robert L. Constable and Scott F. Smith. Computational foundations of basic recursive function theory. *Theoretical Computer Science*, 121(1&2):89–112, 1993.
18. Karl Crary. *Type-Theoretic Methodology for Practical Programming Languages*. PhD thesis, Cornell University, Ithaca, NY, August 1998.
19. Maulik A. Dave. Compiler verification: a bibliography. *ACM SIGSOFT Software Engineering Notes*, 28(6):2, 2003.
20. Andrew D. Gordon. Bisimilarity as a theory of functional programming. *Electr. Notes Theor. Comput. Sci.*, 1:232–252, 1995.
21. Jason Hickey, Aleksey Nogin, Robert L. Constable, Brian E. Aydemir, Eli Barzilay, Yegor Bryukhov, Richard Eaton, Adam Granicz, Alexei Kopylov, Christoph Kreitz, Vladimir Krupski, Lori Lorigo, Stephan Schmitt, Carl Witty, and Xin Yu. MetaPRL - a modular logical environment. In *TPHOLS*, volume 2758 of *LNCS*, pages 287–303. Springer, 2003.
22. Jason J. Hickey. *The MetaPRL Logical Programming Environment*. PhD thesis, Cornell University, Ithaca, NY, January 2001.
23. Douglas J. Howe. Equality in lazy computation systems. In *Proceedings of Fourth IEEE Symposium on Logic in Computer Science*, pages 198–203. IEEE Computer Society, 1989.
24. Douglas J. Howe. Proving congruence of bisimulation in functional programming languages. *Inf. Comput.*, 124(2):103–112, 1996.
25. Alexei Kopylov. Dependent intersection: A new way of defining records in type theory. In *LICS*, pages 86–95. IEEE Computer Society, 2003.
26. Alexei Kopylov. *Type Theoretical Foundations for Data Structures, Classes, and Objects*. PhD thesis, Cornell University, Ithaca, NY, 2004.
27. Christoph Kreitz. *The Nuprl Proof Development System, Version 5, Reference Manual and User's Guide*. Cornell University, Ithaca, NY, 2002. www.nuprl.org/html/02cucs-NuprlManual.pdf.
28. Leslie Lamport. The part-time parliament. *ACM Trans. Comput. Syst.*, 16(2):133–169, 1998.
29. Xavier Leroy. Formal certification of a compiler back-end or: programming a compiler with a proof assistant. In *33rd ACM SIGPLAN-SIGACT Symp. on Principles of Programming Languages*, pages 42–54. ACM, 2006.
30. Guodong Li. *Formal Verification of Programs and Their Transformations*. PhD thesis, University of Utah, 2010.
31. John McCarthy. Recursive functions of symbolic expressions and their computation by machine, Part I. *Commun. ACM*, 3(4):184–195, 1960.
32. P.F. Mendler. *Inductive Definition in Type Theory*. PhD thesis, Cornell University, Ithaca, NY, 1988.
33. Vincent Rahli, Nicolas Schiper, Robbert Van Renesse, Mark Bickford, and Robert L. Constable. A diversified and correct-by-construction broadcast service. In *The 2nd Int'l Workshop on Rigorous Protocol Engineering (WRiPE)*, Austin, TX, October 2012.
34. Nicolas Schiper, Vincent Rahli, Robbert Van Renesse, Mark Bickford, and Robert L. Constable. ShadowDB: A replicated database on a synthesized consensus core. In *Eighth Workshop on Hot Topics in System Dependability, HotDep'12*, 2012.
35. Scott F. Smith. *Partial Objects in Type Theory*. PhD thesis, Cornell University, Ithaca, NY, 1989.