

Mechanized Metatheory for the Masses: The POPLMARK Challenge

Brian E. Aydemir¹, Aaron Bohannon¹, Matthew Fairbairn², J. Nathan Foster¹,
Benjamin C. Pierce¹, Peter Sewell², Dimitrios Vytiniotis¹, Geoffrey
Washburn¹, Stephanie Weirich¹, and Steve Zdancewic¹

¹ Department of Computer and Information Science, University of Pennsylvania

² Computer Laboratory, University of Cambridge

Subversion Revision: 171

Document generated on: May 11, 2005 at 15:53

Abstract. How close are we to a world where every paper on programming languages is accompanied by an electronic appendix with machine-checked proofs?

We propose an initial set of benchmarks for measuring progress in this area. Based on the metatheory of System F_{ω} , a typed lambda-calculus with second-order polymorphism, subtyping, and records, these benchmarks embody many aspects of programming languages that are challenging to formalize: variable binding at both the term and type levels, syntactic forms with variable numbers of components (including binders), and proofs demanding complex induction principles. We hope that these benchmarks will help clarify the current state of the art, provide a basis for comparing competing technologies, and motivate further research.

1 Introduction

Many proofs about programming languages are long, straightforward, and tedious, with just a few interesting cases. Their complexity arises from the management of many details rather than from deep conceptual difficulties; yet small mistakes or overlooked cases can invalidate large amounts of work. These effects are amplified as languages scale: it becomes hard to keep definitions and proofs consistent, to reuse work, and to ensure tight relationships between theory and implementations. Automated proof assistants offer the hope of significantly easing these problems. However, despite much encouraging progress in recent years and the availability of several mature tools (ACL2 [26], Coq [3], HOL [18], HOL Light [21], Isabelle [35], Lego [28], NuPRL [7], PVS [40], Twelf [42], etc.), their use is still not commonplace.

We believe that the time is ripe to join the efforts of the two communities, bringing developers of automated proof assistants together with a large pool of eager potential clients—programming language designers and researchers. In particular, we would like to answer two questions:

1. What is the current state of the art in formalizing language metatheory and semantics? What can be recommended as best practices for groups (typically not proof-assistant experts) embarking on formalizing language definitions, either small- or large-scale?
2. What improvements are needed to make the use of tool support commonplace? What can each community contribute?

Over the past several months, we have surveyed the landscape of proof assistants, language representation strategies, and related tools. Collectively, we have applied automated theorem proving technology to a number of problems, including proving transitivity of the algorithmic subtype relation in KernelF_< [6, 5, 9], proving type soundness of Featherweight Java [24], proving type soundness of variants of the simply typed λ -calculus and F_<, and a substantial formalization of the behavior of TCP, UDP, and the Sockets API. We have carried out these case studies using a variety of object-language representation strategies, proof techniques, and proving environments. We have also experimented with lightweight tools designed to make it easier to define and typeset both formal and informal mathematics. Although experts in programming language theory, we are relative outsiders with respect to computer-aided proof.

Our conclusion from these experiments is that the relevant technology has developed *almost* to the point where it can be widely used by language researchers. We seek to push it over the threshold, making the use of proof tools common practice in programming language research—mechanized metatheory for the masses.

Tool support for formal reasoning about programming languages would be useful at many levels:

1. *Machine-checked metatheory.* These are the classic problems: type preservation and soundness theorems, unique decomposition properties for operational semantics, proofs of equivalence between algorithmic and declarative variants of type systems, etc. At present such results are typically proved by hand for small to medium-sized calculi, and are not proved at all for full language definitions. We envision a future in which the papers in conferences such as *Principles of Programming Languages (POPL)* and the *International Conference on Functional Programming (ICFP)* are routinely accompanied by mechanically checkable proofs of the theorems they claim.
2. *Use of definitions as oracles for testing and animation.* When developing a language implementation together with a formal definition one would like to use the definition as an oracle for testing. This requires tools that can decide typing and evaluation relationships, and they might differ from the tools used for (1) or be embedded in the same proof assistant. In some cases one could use a definition directly as a prototype.
3. *Support for engineering large-scale definitions.* As we move to full language definitions—on the scale of Standard ML [30] or larger—pragmatic “software engineering” issues become increasingly important, as do the potential benefits of tool support. For large definitions, the need for elegant and concise

notation becomes pressing, as witnessed by the care taken by present-day researchers using informal mathematics. Even lightweight tool support, without full mechanized proof, could be very useful in this domain, e.g., for sort checking and typesetting of definitions and of informal proofs, automatically instantiating definitions, performing substitutions, etc.

Large scale formalization of languages is already within reach of current technology. For examples, see the work on proofs of correctness of the Damas-Milner type inference algorithm for ML [12, 33], semantics for C [37], semantics for Standard ML [46, 48, 20], and semantics and proofs of correctness for substantial subsets of Java [36, 27, 34]. Some other significant existing applications of mechanized metatheory include Foundational Proof Carrying Code [1] and Typed Assembly Languages [8]. Inspired by these successes, we seek to make mechanized metatheory more accessible to programming languages researchers.

We hope to stimulate progress by providing a framework for comparing alternative technologies. We issue here an initial set of challenge problems, dubbed the POPLMARK Challenge, chosen to exercise some aspects of programming languages that are known to be difficult to formalize: variable binding at both term and type levels, syntactic forms with variable numbers of components (including binders), and proofs demanding complex induction principles. Such challenge problems have been used in the past within the theorem proving community to focus attention on specific areas and to evaluate the relative merits of different tools; these have ranged in scale from benchmark suites and small problems [45, 19, 11, 23, 15, 32] up to the grand challenges of Floyd, Hoare, and Moore [13, 22, 31]. We hope that our challenge will have a similarly stimulating effect.

Our problems are drawn from the basic metatheory of a call-by-value variant of System F_{λ} [5, 9], enriched with records, record subtyping, and record patterns. We provide an informal definition of its type system and operational semantics and outline proofs of some of its metatheory in Appendix A. This language is of moderate scale—significantly more complex than simply typed lambda-calculus or “mini ML,” but much smaller than a full-blown programming language—to keep the work involved in attempting the challenges manageable. (Our challenges therefore explicitly address only points 1 and 2 above; we regard the pragmatic issues of point 3 as equally critical, but it is not yet clear to us how to formulate a useful challenge problem at this larger scale.) As these challenges are met, we anticipate that the set of problems will be extended to emphasize other language features and proof techniques.

We have begun collecting and publicizing solutions to these challenge problems and information related to mechanized metatheory on our web site.³ In the longer run, we hope that this site and accompanying e-mail discussion list⁴ will serve as a forum for promoting and advancing the current best practices in proof assistant technology and making this technology available to the broader programming languages community and beyond. We encourage researchers to

³ <http://www.cis.upenn.edu/proj/plclub/mmm/>

⁴ poplmark@lists.seas.upenn.edu

try out the POPLMARK Challenge using their favorite tools and send us their solutions.

In the next section, we discuss in more detail our reasons for selecting this specific set of challenge problems. Section 3 describes the problems themselves, and Section 4 sketches some avenues for further development of the challenge problem set.

2 Design of the Challenge

This section motivates our choice of challenge problems and discusses the evaluation criteria for proposed solutions to the challenges. Since variable binding is a central aspect of the challenges, we briefly discuss relevant techniques and sketch some of our own experience in this area.

2.1 Problem Selection

The goal of the POPLMARK Challenge is to provide a small, well-defined set of problems that capture many of the most critical issues in formalizing programming language metatheory. By its nature, such a benchmark will not be able to reflect *all* important issues—it is not practical to require challenge participants to formalize a large-scale language, for example. Instead, the POPLMARK problems concentrate on a few important features:

- *Binding*. Most programming languages have some form of binding in their syntax and require a treatment of α -equivalence and capture-avoiding substitution in their semantics. To adequately represent many languages, the representation strategy must support multiple kinds of binders (e.g. term and type), constructs introducing multiple binders over the same scope, and potentially unbounded lists of binders (e.g. for record patterns).
- *Complex inductions*. Programming language definitions often involve complex, mutually recursive definitions. Structural induction over such objects, mutual induction, and induction on heights or pairs of derivations are all commonplace in metatheory.
- *Experimentation*. Proofs about programming languages are just one aspect of formalization; for some applications, experimenting with formalized language designs is equally interesting. It should be easy for the language designer to execute typechecking algorithms, generate sample program behaviors, and—most importantly—test real language implementations against the formalized definitions.
- *Component reuse*. To further facilitate experimentation with and sharing of language designs, the infrastructure should support some way of reusing prior definitions and parts of proofs.

We have carefully constructed the POPLMARK Challenge to stress these features; a theorem-proving infrastructure that addresses the whole challenge

should be applicable across a wide spectrum of programming language theory. While we believe that the features above are essential, our challenge does not address many other interesting and tricky-to-formalize constructs and reasoning principles. We discuss possible extensions to the challenge in Section 4.

2.2 Evaluation Criteria

A solution to the POPLMARK Challenge will consist of appropriate software tools, a language representation strategy, and a demonstration that this infrastructure is sufficient to formalize the problems described in Section 3. Appendix A presents reasonably detailed informal proofs of the challenge properties. Solutions to the challenge should follow the overall structure of these proofs, though we expect that details will vary from prover to prover and across term representations. In all cases, there must be an argument for why the formalization is equivalent to the presentation in Section 3—i.e., an adequacy theorem—which should be as simple as possible.

The primary metric of success (beyond correctness, of course) is that a solution should give us confidence of future success of other formalizations carried out using similar techniques. In particular, this implies that:

- *The technology should impose reasonable overheads.* We accept that there is a cost to formalization, and our goal is *not* to be able to prove things more easily than by hand (although that would certainly be welcome). We are willing to spend more time and effort to use the proof infrastructure, but the overhead of doing so must not be prohibitive. (For example, as we discuss below, our experience is that explicit de Bruijn-indexed representations of variable binding structure fail this test.)
- *The technology should be transparent.* The representation strategy and proof assistant syntax should not depart too radically from the usual conventions familiar to the technical audience, and the content of the theorems themselves should be apparent to someone not deeply familiar with the theorem proving technology used or the representation strategy chosen.
- *The technology should have a reasonable cost of entry.* The infrastructure should be usable (after, say, one semester of training) by someone who is knowledgeable about programming language theory but not an expert in theorem prover technology.

2.3 Representing Binders

The problem of representing and reasoning about inductively-defined structures with binders is central to the POPLMARK challenges. Representing binders has been recognized as crucial by the theorem proving community, and many different solutions to this problem have been proposed. In our (still limited) experience, none emerge as clear winners. In this section we briefly summarize the main approaches and, where applicable, describe our own experiments using them. Our survey is far from complete and we refrain from drawing any hard

conclusions, to give the proponents of each method a chance to try their hand at meeting the challenge.

A first-order, named approach very similar in flavor to standard informal presentations was used by Vestergaard and Brotherston to formalize some metatheory of untyped λ -calculus [49, 50]. Their representation requires that each binder initially be assigned a unique name—one aspect of the so-called Barendregt convention [2].

Another popular concrete representation is de Bruijn’s nameless representation [10]. De Bruijn indices are easy to understand and support the full range of induction principles needed to reason over terms. In our experience, however, de Bruijn representations have two major flaws. First, the *statements* of theorems require complicated clauses involving “shifted” terms and contexts. These extra clauses make it difficult to see the correspondence between informal and formal versions of the same theorem—there is no question of simply typesetting the formal statement and pasting it into a paper. Second, while the notational clutter is manageable for “toy” examples of the size of the simply-typed lambda calculus, we have found it becomes quite a heavy burden even for fairly small languages like F_{\leq} .

In their formalization of properties of pure type systems, McKinna and Pollock use a hybrid approach that combines the above two representation strategies. In this approach, free variables are ordinary names while bound variables are represented using de Bruijn indices [29].

A radically different approach to representing terms with binders is higher-order abstract syntax (HOAS) [41]. In HOAS representations, binders in the meta-language are used to represent binders in the object language. Our experience with HOAS encodings (mainly as realized in Twelf) is that they provide a conveniently high level of abstraction, encapsulating much of the complexity of reasoning about binders. However, the strengths and limitations of the approach are not yet clearly understood, and it can sometimes require significant ingenuity to encode particular language features or proof ideas in this style.

Gordon and Melham propose a way to axiomatize inductive reasoning over untyped lambda-terms [17] and suggest that other inductive structures with binding can be encoded by setting up a correspondence with the untyped lambda terms. Norrish has pursued this direction [38, 39], but observes that these axioms are cumbersome to use without some assistance from the theorem-proving tool. In particular, the axioms use universal quantification in inductive hypotheses where in informal practice “some/any” quantification is used. He has developed a library of lemmas about a system of permutations on top of the axioms that aids reasoning significantly.

Several recent approaches to binding take the concept of “swapping” as a primitive, and use it to build a nominal logic. Gabbay and Pitts proposed a method of reasoning about binders based upon a set theory extended with an intrinsic notion of permutation [14]. Pitts followed this up by proposing a new “nominal” logic based upon the idea of permutations [44]. More recent work by Urban proposes methods based on the same intuitions but carried out within a

conventional logic [47]. Our own preliminary experiments with Urban’s methods have been encouraging.

3 The Challenge

Our challenge problems are taken from the basic metatheory of System F_{\leq} . This system is formed by enriching the types and terms of System F with a subtype relation, refining universal quantifiers to carry subtyping constraints, and adding records, record subtyping, and record patterns. Our presentation is based on Pierce’s *Types and Programming Languages* [43]; other good sources for background information are the seminal papers by Cardelli et al [5] and Curien and Ghelli [9].

The challenge comprises three distinct parts. The first deals just with the type language of F_{\leq} ; the second considers terms, evaluation, and type soundness. Each of these is further subdivided into two parts, starting with definitions and properties for *pure* F_{\leq} and then asking that the same properties be proved for F_{\leq} enriched with records and patterns. This partitioning allows the development to start small, but also—and more importantly—focuses attention on issues of reuse: How much of the first sub-part can be re-used verbatim in the second sub-part? The third problem asks that useful algorithms be extracted from the earlier formal definitions and used to “animate” some simple properties.

Challenge 1A: Transitivity of Subtyping

The first part of this challenge problem deals purely with the type language of F_{\leq} . The syntax for this language is defined by the following grammar and inference rules. Although the grammar is simple—it has only four syntactic forms—some of its key properties require fairly sophisticated reasoning.

Syntax:

$T ::=$	<i>types</i>
X	<i>type variable</i>
Top	<i>maximum type</i>
$T \rightarrow T$	<i>type of functions</i>
$\forall X < : T. T$	<i>universal type</i>
$\Gamma ::=$	<i>type environments</i>
\emptyset	<i>empty type env.</i>
$\Gamma, X < : T$	<i>type variable binding</i>

In $\forall X < : T_1. T_2$, the variable X is a binding occurrence with scope T_2 (X is not bound in T_1). In $\Gamma, X < : T$, the X must not be in the domain of Γ , and the free variables of T must all be in the domain of Γ .

Following standard practice, issues such as the use of α -conversion, capture avoidance during substitution, etc. are left implicit in what follows. There are several ways in which these issues can be formalized: we might take Γ as a

concrete structure (such as an association list of named variables and types) but quotient types and terms up to alpha equivalence, or we could take entire judgments up to alpha equivalence. We might axiomatize the well-formedness of typing environments and types using auxiliary $\Gamma \vdash \text{ok}$ and $\Gamma \vdash T \text{ ok}$ judgments. And so on. We leave these decisions to the individual formalization.

It is acceptable to make small changes to the rules below to reflect these decisions, such as adding well-formedness premises. Changing the presentation of the rules to a notationally different but “obviously equivalent” style such as HOAS is also acceptable, but there must be a clear argument that it is really equivalent. Also, whatever formalization is chosen should make clear that we are only dealing with *well-scoped* terms. For example, it should not be possible to derive $X \prec: \text{Top}$ in the empty typing environment.

The subtyping relation captures the intuition “if S is a subtype of T (written $S \prec: T$) then an instance of S may safely be used wherever an instance of T is expected.” It is defined as the least relation closed under the following rules.

Subtyping

$$\boxed{\Gamma \vdash S \prec: T}$$

$$\Gamma \vdash S \prec: \text{Top} \quad (\text{SA-TOP})$$

$$\Gamma \vdash X \prec: X \quad (\text{SA-REFL-TVAR})$$

$$\frac{X \prec: U \in \Gamma \quad \Gamma \vdash U \prec: T}{\Gamma \vdash X \prec: T} \quad (\text{SA-TRANS-TVAR})$$

$$\frac{\Gamma \vdash T_1 \prec: S_1 \quad \Gamma \vdash S_2 \prec: T_2}{\Gamma \vdash S_1 \rightarrow S_2 \prec: T_1 \rightarrow T_2} \quad (\text{SA-ARROW})$$

$$\frac{\Gamma \vdash T_1 \prec: S_1 \quad \Gamma, X \prec: T_1 \vdash S_2 \prec: T_2}{\Gamma \vdash \forall X \prec: S_1. S_2 \prec: \forall X \prec: T_1. T_2} \quad (\text{SA-ALL})$$

These rules present an *algorithmic* version of the subtyping relation. In contrast to the more familiar *declarative* presentation, these rules are syntax-directed, as might be found in the implementation of a type checker; the algorithmic rules are also somewhat easier to reason with, having, for example, an obvious inversion property (Lemma A.12). The declarative rules differ from these by explicitly stating that subtyping is reflexive and transitive. However, reflexivity and transitivity also turn out to be derivable properties in the algorithmic system. A straightforward induction shows that the algorithmic rules are reflexive. The first challenge is to show that that they are also transitive.

3.1 LEMMA [TRANSITIVITY OF ALGORITHMIC SUBTYPING]: If $\Gamma \vdash S \prec: Q$ and $\Gamma \vdash Q \prec: T$, then $\Gamma \vdash S \prec: T$. \square

The difficulty here lies in the reasoning needed to prove this lemma. Transitivity must be proven simultaneously with another property, called *narrowing*, by an inductive argument with case analyses on the final rules used in the given derivations. Full details of this proof appear in Appendix A.

3.2 LEMMA [NARROWING]: If $\Gamma, X <: Q, \Delta \vdash M <: N$ and $\Gamma \vdash P <: Q$ then $\Gamma, X <: P, \Delta \vdash M <: N$. \square

Challenge 1B: Transitivity of Subtyping with Records

We now extend this challenge by enriching the type language with record types. The new syntax and subtyping rule for record types are shown below. Implicit in the syntax is the condition that the labels $\{l_i^{i \in 1..n}\}$ appearing in a record type $\{l_i : T_i^{i \in 1..n}\}$ are pairwise distinct.

New syntactic forms:

$T ::= \dots$	<i>types</i>
$\{l_i : T_i^{i \in 1..n}\}$	<i>type of records</i>

New subtyping rules

 $\Gamma \vdash S <: T$

$$\frac{\{l_i^{i \in 1..n}\} \subseteq \{k_j^{j \in 1..m}\} \quad \text{if } k_j = l_i, \text{ then } \Gamma \vdash S_j <: T_i}{\Gamma \vdash \{k_j : S_j^{j \in 1..m}\} <: \{l_i : T_i^{i \in 1..n}\}} \quad (\text{SA-RCD})$$

Although it has been shown that records can actually be encoded in pure $F_{<}$ [4, 16], dealing with them directly is a worthwhile task since, unlike other syntactic forms, record types have an arbitrary (finite) number of fields. Also, the informal proof for Challenge 1A extends to record types by only adding the appropriate cases. A formal proof should reflect this.

Challenge 2A: Type Safety for Pure $F_{<}$

The next challenge considers the type soundness of pure $F_{<}$ (without record types, for the moment). Below, we complete the definition of $F_{<}$ by describing the syntax of terms, values, and typing environments with term binders and giving inference rules for the typing relation and a small-step operational semantics.

As usual in informal presentations, we elide the formal definition of substitution and simply assume that the substitutions of a type P for X in T (denoted $[X \mapsto P]T$) and of a term q for x in t (denoted $[x \mapsto q]t$) are capture-avoiding.

Syntax:

$t ::=$	<i>terms</i>
x	<i>variable</i>
$\lambda x : T. t$	<i>abstraction</i>
$t \ t$	<i>application</i>
$\lambda X <: T. t$	<i>type abstraction</i>
$t \ [T]$	<i>type application</i>
$v ::=$	<i>values</i>
$\lambda x : T. t$	<i>abstraction value</i>
$\lambda X <: T. t$	<i>type abstraction value</i>

$\Gamma ::=$

\emptyset
 $\Gamma, x:T$
 $\Gamma, X<:T$

type environments
empty type env.
term variable binding
type variable binding

Typing

$\boxed{\Gamma \vdash t : T}$

$$\frac{x:T \in \Gamma}{\Gamma \vdash x : T} \quad (\text{T-VAR})$$

$$\frac{\Gamma, x:T_1 \vdash t_2 : T_2}{\Gamma \vdash \lambda x:T_1. t_2 : T_1 \rightarrow T_2} \quad (\text{T-ABS})$$

$$\frac{\Gamma \vdash t_1 : T_{11} \rightarrow T_{12} \quad \Gamma \vdash t_2 : T_{11}}{\Gamma \vdash t_1 t_2 : T_{12}} \quad (\text{T-APP})$$

$$\frac{\Gamma, X<:T_1 \vdash t_2 : T_2}{\Gamma \vdash \lambda X<:T_1. t_2 : \forall X<:T_1. T_2} \quad (\text{T-TABS})$$

$$\frac{\Gamma \vdash t_1 : \forall X<:T_{11}. T_{12} \quad \Gamma \vdash T_2 <: T_{11}}{\Gamma \vdash t_1 [T_2] : [X \mapsto T_2]T_{12}} \quad (\text{T-TAPP})$$

$$\frac{\Gamma \vdash t : S \quad \Gamma \vdash S <: T}{\Gamma \vdash t : T} \quad (\text{T-SUB})$$

Evaluation

$\boxed{t \longrightarrow t'}$

$$(\lambda x:T_{11}. t_{12}) v_2 \longrightarrow [x \mapsto v_2]t_{12} \quad (\text{E-APPABS})$$

$$(\lambda X<:T_{11}. t_{12}) [T_2] \longrightarrow [X \mapsto T_2]t_{12} \quad (\text{E-TAPPTABS})$$

Evaluation contexts:

$E ::=$

$[-]$
 $E t$
 $v E$
 $E [T]$

evaluation contexts

hole
app fun
app arg
type fun

Evaluation in context:

$$\frac{t_1 \longrightarrow t'_1}{E[t_1] \longrightarrow E[t'_1]} \quad (\text{E-CTX})$$

The evaluation relation is presented in two parts: the rules E-APPABS and E-TAPPTABS capture the immediate reduction rules of the language, while E-CTX permits reduction under an arbitrary *evaluation context* E . For F_{ζ} , one would have an equally clear definition and slightly simpler proofs using explicit closure rules for the evaluation relation. We use evaluation contexts with an eye to larger languages and languages with non-local control operators such as exceptions, for which (in informal mathematics) they are an important tool for reducing notational clutter in definitions.⁵ Evaluation contexts are also particularly interesting from the point of view of formalization when they include binders, though unfortunately there are no examples of this in call-by-value F_{ζ} .

Type soundness is usually proven in the style popularized by Wright and Felleisen [51], in terms of *preservation* and *progress* theorems. Challenge 2A is to prove these properties for pure F_{ζ} .

3.3 THEOREM [PRESERVATION]: If $\Gamma \vdash t : T$ and $t \longrightarrow t'$, then $\Gamma \vdash t' : T$. \square

3.4 THEOREM [PROGRESS]: If t is a closed, well-typed F_{ζ} term (i.e., if $\vdash t : T$ for some T), then either t is a value or else there is some t' with $t \longrightarrow t'$. \square

Unlike the proof of transitivity of subtyping, the inductive arguments required here are straightforward. However, variable binding becomes a more significant issue, since this language includes binding of both type and term variables. Several lemmas relating to both kinds of binding must also be shown, in particular lemmas about type and term substitutions. These lemmas, in turn, require reasoning about permuting, weakening, and strengthening typing environments.

Challenge 2B: Type Safety with Records and Pattern Matching

The next challenge is to extend the preservation and progress results to cover records and pattern matching. The new syntax and rules for this language appear below. As for record types, the labels $\{l_i^{i \in 1..n}\}$ appearing in a record $\{l_i = t_i^{i \in 1..n}\}$ are assumed to be pairwise distinct. Similarly, the variable patterns appearing in a pattern are assumed to bind pairwise distinct variables.

New syntactic forms:

$t ::=$...	<i>terms</i>
	$\{l_i = t_i^{i \in 1..n}\}$	<i>record</i>
	$t.l$	<i>projection</i>

⁵ This design choice has generated a robust debate on the POPLMARK discussion list as to whether evaluation contexts *must* be used in order for a solution to count as valid, or whether an “obviously equivalent” presentation such as an evaluation relation with additional congruence rules is acceptable. We prefer evaluation contexts for the reasons we have given, but the consensus of the community appears to be that one should accept solutions in other styles. However, a good solution must be formulated in a style that provides similar clarity as the language scales.

$\text{let } p=t \text{ in } t$	<i>pattern binding</i>
$p ::=$	<i>patterns</i>
$x:T$	<i>variable pattern</i>
$\{l_i=p_i \text{ }^{i \in 1..n}\}$	<i>record pattern</i>
$v ::= \dots$	<i>values</i>
$\{l_i=v_i \text{ }^{i \in 1..n}\}$	<i>record value</i>

New typing rules

$\boxed{\Gamma \vdash t : T}$

$$\frac{\Gamma \vdash t_1 : T_1 \quad \vdash p : T_1 \Rightarrow \Delta \quad \Gamma, \Delta \vdash t_2 : T_2}{\Gamma \vdash \text{let } p=t_1 \text{ in } t_2 : T_2} \quad (\text{T-LET})$$

$$\frac{\text{for each } i \quad \Gamma \vdash t_i : T_i}{\Gamma \vdash \{l_i=t_i \text{ }^{i \in 1..n}\} : \{l_i:T_i \text{ }^{i \in 1..n}\}} \quad (\text{T-RCD})$$

$$\frac{\Gamma \vdash t_1 : \{l_i:T_i \text{ }^{i \in 1..n}\}}{\Gamma \vdash t_1.l_j : T_j} \quad (\text{T-PROJ})$$

Pattern typing rules:

$$\vdash (x:T) : T \Rightarrow x : T \quad (\text{P-VAR})$$

$$\frac{\text{for each } i \quad \vdash p_i : T_i \Rightarrow \Delta_i}{\vdash \{l_i=p_i \text{ }^{i \in 1..n}\} : \{l_i:T_i \text{ }^{i \in 1..n}\} \Rightarrow \Delta_n, \dots, \Delta_1} \quad (\text{P-RCD})$$

New evaluation rules

$\boxed{t \longrightarrow t'}$

$$\text{let } p=v_1 \text{ in } t_2 \longrightarrow \text{match}(p, v_1)t_2 \quad (\text{E-LETV})$$

$$\{l_i=v_i \text{ }^{i \in 1..n}\}.l_j \longrightarrow v_j \quad (\text{E-PROJRCD})$$

New evaluation contexts:

$E ::= \dots$	<i>evaluation contexts</i>
$E.l$	<i>projection</i>
$\{l_i=v_i \text{ }^{i \in 1..j-1}, l_j=E, l_k=t_k \text{ }^{k \in j+1..n}\}$	<i>record</i>
$\text{let } p=E \text{ in } t_2$	<i>let binding</i>

Matching rules:

$$\text{match}(x:T, v) = [x \mapsto v] \quad (\text{M-VAR})$$

$$\frac{\{l_i \text{ }^{i \in 1..n}\} \subseteq \{k_j \text{ }^{j \in 1..m}\} \quad \text{if } l_i = k_j, \text{ then } \text{match}(p_i, v_j) = \sigma_i}{\text{match}(\{l_i=p_i \text{ }^{i \in 1..n}\}, \{k_j=v_j \text{ }^{j \in 1..m}\}) = \sigma_n \circ \dots \circ \sigma_1} \quad (\text{M-RCD})$$

Compared to the language of Challenge 2A, the `let` construct is a fundamentally new binding form, since patterns may bind an arbitrary (finite) number of term variables.

Challenge 3: Testing and Animating with Respect to the Semantics

Given a complete formal definition of a language, there are at least two interesting ways in which it can be used (as opposed to being reasoned about). When implementing the language, it should be possible to use the formal definition as an oracle for *testing* the implementation—checking that it does conform to the definition by running test cases in the implementation and confirming formally that the outcome is as prescribed. Secondly, one would like to construct a prototype implementation from the definition and use it for *animating* the language, i.e., exploring the language’s properties on particular examples. In both cases, this should be done without any unverified (and thus error-prone) manual translation of definitions.

Our final challenge is to provide an implementation of this functionality, specifically for the following three tasks (using the language of Challenge 2B):

1. Given F_{\leq} , terms t and t' , decide whether $t \longrightarrow t'$.
2. Given F_{\leq} , terms t and t' , decide whether $t \longrightarrow^* t' \not\rightarrow$, where \longrightarrow^* is the reflexive-transitive closure of \longrightarrow .
3. Given an F_{\leq} term t , find a term t' such that $t \longrightarrow t'$.

The first two subtasks are useful for testing language implementations, while the last is useful for animating the definition. For all three subtasks, the system(s) should accept syntax that is “reasonably close” to that of informal (ASCII) mathematical notation, though it maybe necessary to translate between the syntaxes of a formal environment and an implementation. We will provide an implementation of an interpreter for F_{\leq} with records and patterns at the challenge’s website in order to make this challenge concrete, together with a graded sequence of example terms. To make a rough performance comparison possible, solutions should indicate execution times for these terms.

A solution to this challenge might make use of decision procedures and tactics of a proof assistant or might extract stand-alone code. In general, it may be necessary to combine theorems (e.g. that a rule-based but algorithmic definition of typing coincides with a declarative definition) and proof search (e.g. deciding particular instances of the algorithmic definition).

4 Beyond the Challenge

The POPLMARK Challenge is not meant to be exhaustive: other aspects of programming language theory raise formalization difficulties that are interestingly different from the problems we have proposed—to name a few: more complex binding constructs such as mutually recursive definitions, logical relations proofs, coinductive simulation arguments, undecidability results, and linear handling of

type environments. As time goes on, we will issue a small number of further challenges highlighting the most important of these issues; suggestions from the community would be welcome. However, we believe that a technology that provides a good solution to the POPLMARK challenge as we have formulated it here will be sufficient to attract eager adopters in the programming languages community, beginning with the authors.

So what are you waiting for? It's time to bring mechanized metatheory to the masses!

Acknowledgments

A number of people have joined us in preliminary discussions of these challenge problems, including Andrew Appel, Karl Crary, Frank Pfenning, Bob Harper, Hugo Herbelin, Jason Hickey, Michael Norrish, Andrew Pitts, Randy Pollack, Carsten Schürmann, Phil Wadler, and Dan Wang. We are especially grateful to Randy Pollack for helping guide our earliest efforts at formalization and to Randy Pollack, Michael Norrish, and Carsten Schürmann for their own work on the challenge problems. Discussions on the POPLMARK mailing list have greatly deepened our understanding of the current state of the art.

Sewell and Fairbairn acknowledge support from a Royal Society University Research Fellowship and EPSRC grant GR/T11715.

References

1. Andrew W. Appel. Foundational proof-carrying code. In *IEEE Symposium on Logic in Computer Science (LICS)*, Boston, Massachusetts, pages 247–258, June 2001.
2. Henk P. Barendregt. *The Lambda Calculus*. North Holland, revised edition, 1984.
3. Yves Bertot and Pierre Castéran. *Interactive Theorem Proving and Program Development*, volume XXV of *EATCS Texts in Theoretical Computer Science*. Springer-Verlag, 2004.
4. Luca Cardelli. Extensible records in a pure calculus of subtyping. Research report 81, DEC/Compaq Systems Research Center, January 1992. Also in C. A. Gunter and J. C. Mitchell, editors, *Theoretical Aspects of Object-Oriented Programming: Types, Semantics, and Language Design*, MIT Press, 1994.
5. Luca Cardelli, Simone Martini, John C. Mitchell, and Andre Scedrov. An extension of System F with subtyping. *Information and Computation*, 109(1–2):4–56, 1994. Summary in TACS '91 (Sendai, Japan, pp. 750–770).
6. Luca Cardelli and Peter Wegner. On understanding types, data abstraction, and polymorphism. *Computing Surveys*, 17(4):471–522, December 1985.
7. Robert L. Constable, Stuart F. Allen, Mark Bromley, Rance Cleaveland, James F. Cremer, Robert W. Harper, Douglas J. Howe, Todd B. Knoblock, Paul Mendler, Prakash Panangaden, James T. Sasaki, and Scott F. Smith. *Implementing Mathematics with the NuPRL Proof Development System*. Prentice-Hall, Englewood Cliffs, NJ, 1986.

8. Karl Crary. Toward a foundational typed assembly language. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, New Orleans, Louisiana, pages 198–212, January 2003.
9. Pierre-Louis Curien and Giorgio Ghelli. Coherence of subsumption: Minimum typing and type-checking in F_{\leq} . *Mathematical Structures in Computer Science*, 2:55–91, 1992. Also in C. A. Gunter and J. C. Mitchell, editors, *Theoretical Aspects of Object-Oriented Programming: Types, Semantics, and Language Design*, MIT Press, 1994.
10. N. G. de Bruijn. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. *Indagationes Mathematicae*, 34(5):381–392, 1972.
11. Louise A. Dennis. Inductive challenge problems, 2000. <http://www.cs.nott.ac.uk/~lad/research/challenges>.
12. Catherine Dubois and Valerie Menissier-Morain. Certification of a type inference tool for ML: Damas-Milner within Coq. *Journal of Automated Reasoning*, 23(3-4):319–346, 1999.
13. Robert W. Floyd. Assigning meanings to programs. In J. T. Schwartz, editor, *Mathematical Aspects of Computer Science*, volume 19 of *Proceedings of Symposia in Applied Mathematics*, pages 19–32, Providence, Rhode Island, 1967. American Mathematical Society.
14. Murdoch Gabbay and Andrew Pitts. A new approach to abstract syntax involving binders. In *14th Symposium on Logic in Computer Science*, pages 214–224, 1999.
15. I. P. Gent and T. Walsh. CSPLib: a benchmark library for constraints. Technical Report APES-09-1999, APES, 1999. Available from <http://www.dcs.st-and.ac.uk/~apes/apesreports.html>. A shorter version appears in the Proceedings of the 5th International Conference on Principles and Practices of Constraint Programming (CP-99).
16. Giorgio Ghelli. *Proof Theoretic Studies about a Minimal Type System Integrating Inclusion and Parametric Polymorphism*. PhD thesis, Università di Pisa, 1990. Technical report TD-6/90, Dipartimento di Informatica, Università di Pisa.
17. Andrew D. Gordon and Tom Melham. Five axioms of alpha-conversion. In J. von Wright, J. Grundy, and J. Harrison, editors, *Theorem Proving in Higher Order Logics: 9th International Conference, TPHOLs '96*, volume 1125 of *Lecture Notes in Computer Science*, pages 173–190. Springer-Verlag, 1996.
18. M. J. C. Gordon and T. F. Melham, editors. *Introduction to HOL: a theorem proving environment for higher order logic*. Cambridge University Press, 1993.
19. Ian Green. The dream corpus of inductive conjectures, 1999. <http://dream.dai.ed.ac.uk/dc/lib.html>.
20. Elsa Gunter and Savi Maharaj. Studying the ML module system in HOL. *The Computer Journal: Special Issue on Theorem Proving in Higher Order Logics*, 38(2):142–151, 1995.
21. John Harrison. HOL Light: A tutorial introduction. In Mandayam Srivas and Albert Camilleri, editors, *Proceedings of the First International Conference on Formal Methods in Computer-Aided Design (FMCAD'96)*, volume 1166 of *Lecture Notes in Computer Science*, pages 265–269. Springer-Verlag, 1996. See also <http://www.cl.cam.ac.uk/users/jrh/hol-light/index.html>.
22. Tony Hoare. The verifying compiler: A grand challenge for computing research. *J. ACM*, 50(1):63–69, 2003.
23. Holger Hoos and Thomas Stuetzle. SATLIB. <http://www.intellektik.informatik.tu-darmstadt.de/SATLIB/>.

24. Atsushi Igarashi, Benjamin Pierce, and Philip Wadler. Featherweight Java: A minimal core calculus for Java and GJ. In *ACM SIGPLAN Conference on Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*, October 1999. Full version in *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 23(3), May 2001.
25. J. J. Joyce and C.-J. H. Seger, editors. *Proceedings of the 6th International Workshop on Higher Order Logic Theorem Proving and its Applications (HUG'93)*, volume 780 of *Lecture Notes in Computer Science*, Vancouver, B.C., Canada, August 1993. Springer-Verlag, 1994.
26. Matt Kaufmann, J Strother Moore, and Panagiotis Manolios. *Computer-Aided Reasoning: An Approach*. Kluwer Academic Publishers, 2000.
27. Gerwin Klein and Tobias Nipkow. A machine-checked model for a Java-like language, virtual machine and compiler. Technical Report 0400001T.1, National ICT Australia, Sydney, March 2004.
28. Zhaohui Luo and Robert Pollack. The LEGO proof development system: A user's manual. Technical Report ECS-LFCS-92-211, University of Edinburgh, May 1992.
29. James McKinna and Robert Pollack. Some lambda calculus and type theory formalized. *Journal of Automated Reasoning*, 23(3–4), November 1999.
30. Robin Milner, Mads Tofte, Robert Harper, and David MacQueen. *The Definition of Standard ML*, Revised edition. MIT Press, 1997.
31. J Strother Moore. A grand challenge proposal for formal methods: A verified stack. In Bernhard K. Aichernig and T. S. E. Maibaum, editors, *Formal Methods at the Crossroads. From Panacea to Foundational Support, 10th Anniversary Colloquium of UNU/IIST, Lisbon, Portugal*, volume 2757 of *Lecture Notes in Computer Science*, pages 161–172. Springer, 2002.
32. J Strother Moore and George Porter. The apprentice challenge. *ACM Trans. Program. Lang. Syst.*, 24(3):193–216, 2002.
33. Wolfgang Naraschewski and Tobias Nipkow. Type inference verified: Algorithm W in Isabelle/HOL. *Journal of Automated Reasoning*, 23:299–318, 1999.
34. Tobias Nipkow, David von Oheimb, and Cornelia Pusch. μ Java: Embedding a programming language in a theorem prover. In F.L. Bauer and R. Steinbrüggen, editors, *Foundations of Secure Computation. Proc. Int. Summer School Marktoberdorf 1999*, pages 117–144. IOS Press, 2000.
35. Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL: A Proof Assistant For Higher-Order Logic*, volume 2283 of *Lecture Notes in Computer Science*. Springer-Verlag, 2002.
36. Tobias Nipkow and David von Oheimb. $\text{Java}_{\text{light}}$ is type-safe—definitely. In *POPL '98: Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 161–170, New York, NY, USA, 1998. ACM Press.
37. Michael Norrish. *Formalising C in HOL*. PhD thesis, Computer Laboratory, University of Cambridge, 1998.
38. Michael Norrish. Mechanising Hankin and Barendregt using the Gordon-Melham axioms. In *MERLIN '03: Proceedings Of The 2003 Workshop On Mechanized Reasoning About Languages With Variable Binding*, pages 1–7. ACM Press, 2003.
39. Michael Norrish. Recursive function definition for types with binders. In Konrad Slind, Annette Bunker, and Ganesh Gopalakrishnan, editors, *Theorem Proving in Higher Order Logics: 17th International Conference, TPHOLS 2004*, volume 3223 of *Lecture Notes In Computer Science*, pages 241–256. Springer-Verlag, 2004.
40. Sam Owre, Sreeranga Rajan, John M. Rushby, Natarajan Shankar, and Mandayam K. Srivas. PVS: Combining specification, proof checking, and model checking. In *International Conference on Computer Aided Verification (CAV), New*

- Brunswick, New Jersey*, volume 1102 of *Lecture Notes in Computer Science*, pages 411–414. Springer-Verlag, July 1996.
41. Frank Pfenning and Conal Elliot. Higher-order abstract syntax. In *PLDI '88: Proceedings of the ACM SIGPLAN 1988 Conference on Programming Language Design and Implementation*, pages 199–208. ACM Press, 1988.
 42. Frank Pfenning and Carsten Schürmann. System description: Twelf — A meta-logical framework for deductive systems. In Harald Ganzinger, editor, *Automated Deduction, CADE 16: 16th International Conference on Automated Deduction*, volume 1632 of *Lecture Notes in Artificial Intelligence*, pages 202–206. Springer-Verlag, 1999.
 43. Benjamin C. Pierce. *Types and Programming Languages*. MIT Press, 2002.
 44. Andrew M. Pitts. Nominal logic, a first order theory of names and binding. *Inf. Comput.*, 186(2):165–193, 2003.
 45. Geoff Sutcliffe and Christian Suttner. The TPTP problem library. *Journal of Automated Reasoning*, 21(2):177–203, 1998.
 46. Donald Syme. Reasoning with the formal definition of Standard ML in HOL. In Joyce and Seger [25], pages 43–60.
 47. Christian Urban and Christine Tasson. Nominal techniques in Isabelle/HOL. Accepted at CADE-20 in Tallinn. See <http://www.mathematik.uni-muenchen.de/~urban/nominal/>.
 48. Myra VanInwegen and Elsa Gunter. HOL-ML. In Joyce and Seger [25], pages 61–74.
 49. René Vestergaard and James Brotherston. The mechanisation of Barendregt-style equational proofs (the residual perspective). In *Mechanized Reasoning about Languages with Variable Binding (MERLIN)*, volume 58 of *Electronic Notes in Theoretical Computer Science*. Elsevier, 2001.
 50. René Vestergaard and James Brotherston. A formalised first-order confluence proof for the λ -calculus using one-sorted variable names. *Information and Computation*, 183(2):212 – 244, 2003. Special edition with selected papers from RTA01.
 51. Andrew K. Wright and Matthias Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, November 1994.

A Paper Proofs

For the sake of concreteness and to standardize proof techniques (and thus make solutions easier to compare), here are paper proofs of the properties mentioned in the challenges.

A technical note for type experts: There are two reasonable ways of defining the subtyping relation of $F_{<}$, differing in their formulation of the rule for comparing bounded quantifiers (rule SA-ALL): a more tractable but less flexible version called the *kernel* rule, and a more expressive but technically somewhat problematic *full* subtyping rule. We choose the full variant here as its metatheory is more interesting.

Algorithmic Subtyping (Challenges 1A and 1B)

We give the proof for the full subtyping relation including records. The proof for the pure system is obtained by simply deleting the cases involving records.

A.1 LEMMA [REFLEXIVITY]: $\Gamma \vdash T <: T$ is provable for every Γ and T . \square

Proof: By induction on the structure of T . \square

A.2 LEMMA [PERMUTATION AND WEAKENING]:

1. Suppose that Δ is a well-formed permutation of Γ —that is, Δ has the same bindings as Γ , and their ordering in Δ preserves the scopes of type variables from Γ , in the sense that, if one binding in Γ introduces a type variable that is mentioned in another binding further to the right, then these bindings appear in the same order in Δ .

Now, if $\Gamma \vdash S <: T$, then $\Delta \vdash S <: T$.

2. If $\Gamma \vdash S <: T$ and $\text{dom}(\Delta) \cap \text{dom}(\Gamma) = \emptyset$, then $\Gamma, \Delta \vdash S <: T$. \square

Proof: Routine inductions. Part (1) is used in the SA-ALL case of part (2). \square

A.3 LEMMA [TRANSITIVITY AND NARROWING]:

1. If $\Gamma \vdash S <: Q$ and $\Gamma \vdash Q <: T$, then $\Gamma \vdash S <: T$.
2. If $\Gamma, X <: Q, \Delta \vdash M <: N$ and $\Gamma \vdash P <: Q$ then $\Gamma, X <: P, \Delta \vdash M <: N$. \square

Proof: The two parts are proved simultaneously, by induction on the structure of Q . At each stage of the induction, the argument for part (2) assumes that part (1) has been established already for the Q in question; part (1) uses part (2) only for strictly smaller Q s.

1. We proceed by an inner induction on the structure of $\Gamma \vdash S <: Q$, with a case analysis on the final rules of this derivation and of $\Gamma \vdash Q <: T$.
If the right-hand derivation is an instance of SA-TOP, then we are done, since $\Gamma \vdash S <: \text{Top}$ by SA-TOP. If the left-hand derivation is an instance of SA-TOP, then $Q = \text{Top}$, and, inspecting the algorithmic rules, we see

that the right-hand derivation must also be an instance of SA-TOP. If either derivation is an instance of SA-REFL-TVAR, then we are again done, since the other derivation is exactly the desired result.

If the left-hand derivation ends with an instance of SA-TRANS-TVAR, then we have $S = Y$ with $Y <: U \in \Gamma$ and a subderivation of $\Gamma \vdash U <: Q$. By the inner induction hypothesis, $\Gamma \vdash U <: T$, and, by SA-TRANS-TVAR again, $\Gamma \vdash Y <: T$, as required.

If the left-hand derivation ends with an instance of SA-ARROW, SA-ALL, or SA-RCD, then, since we have already considered the case where the right-hand derivation ends with SA-TOP, it must end with the same rule as the left. If this rule is SA-ARROW, then we have $S = S_1 \rightarrow S_2$, $Q = Q_1 \rightarrow Q_2$, and $T = T_1 \rightarrow T_2$, with subderivations $\Gamma \vdash Q_1 <: S_1$, $\Gamma \vdash S_2 <: Q_2$, $\Gamma \vdash T_1 <: Q_1$, and $\Gamma \vdash Q_2 <: T_2$. We apply part (1) of the outer induction hypothesis twice (noting that Q_1 and Q_2 are both immediate subterms of Q) to obtain $\Gamma \vdash T_1 <: S_1$ and $\Gamma \vdash S_2 <: T_2$, and then use SA-ARROW to obtain $\Gamma \vdash S_1 \rightarrow S_2 <: T_1 \rightarrow T_2$.

In the case where the two derivations end with SA-ALL, we have $S = \forall X <: S_1 . S_2$, $Q = \forall X <: Q_1 . Q_2$, and $T = \forall X <: T_1 . T_2$, with

$$\begin{array}{ll} \Gamma \vdash Q_1 <: S_1 & \Gamma, X <: Q_1 \vdash S_2 <: Q_2 \\ \Gamma \vdash T_1 <: Q_1 & \Gamma, X <: T_1 \vdash Q_2 <: T_2 \end{array}$$

as subderivations. By part (1) of the outer induction hypothesis (Q_1 being an immediate subterm Q), we can combine the two subderivations for the bounds to obtain $\Gamma \vdash T_1 <: S_1$. For the bodies, we need to work a little harder, since the two contexts do not quite agree. We first use part (2) of the outer induction hypothesis (noting again that Q_1 is an immediate subterm of Q) to narrow the bound of X in the derivation of $\Gamma, X <: Q_1 \vdash S_2 <: Q_2$, obtaining $\Gamma, X <: T_1 \vdash S_2 <: Q_2$. Now part (1) of the outer induction hypothesis applies (Q_2 being an immediate subterm of Q), yielding $\Gamma, X <: T_1 \vdash S_2 <: T_2$. Finally, by SA-ALL, $\Gamma \vdash \forall X <: S_1 . S_2 <: \forall X <: T_1 . T_2$.

Finally, if the two derivations both end with SA-RCD, then we have $S = \{l_i : S_i \mid i \in 1..n\}$, $Q = \{h_j : Q_j \mid j \in 1..m\}$, and $T = \{g_k : T_k \mid k \in 1..p\}$, with

$$\begin{array}{l} \{g_k \mid k \in 1..p\} \subseteq \{h_j \mid j \in 1..m\} \subseteq \{l_i \mid i \in 1..n\} \\ \text{if } l_i = h_j, \text{ then } \Gamma \vdash S_i <: Q_j \\ \text{if } h_j = g_k, \text{ then } \Gamma \vdash Q_j <: T_k \end{array}$$

as premises. If $l_i = g_k$, then there is an h_j such that $h_j = g_k = l_i$. Thus, $\Gamma \vdash S_i <: Q_j$ and $\Gamma \vdash Q_j <: T_k$. Observing that Q_j is an immediate subterm of Q , we apply the outer induction hypothesis to obtain $\Gamma \vdash S_i <: T_k$. By SA-RCD, $\Gamma \vdash \{l_i : S_i \mid i \in 1..n\} <: \{g_k : T_k \mid k \in 1..p\}$.

2. We proceed by an inner induction on the structure of the derivation of $\Gamma, X <: Q, \Delta \vdash M <: N$. Most of the cases proceed by straightforward use of the inner induction hypothesis. The interesting case is SA-TRANS-TVAR with $M = X$ and we have $\Gamma, X <: Q, \Delta \vdash Q <: N$ as a subderivation. Applying the inner induction hypothesis to this subderivation yields $\Gamma, X <: P, \Delta \vdash Q <: N$.

Also, applying weakening (Lemma A.2, part 2) to $\Gamma \vdash P <: Q$ yields $\Gamma, X<:P, \Delta \vdash P <: Q$. Now, by part (1) of the outer induction hypothesis (with the same Q), we have $\Gamma, X<:P, \Delta \vdash P <: N$. Rule SA-TRANS-TVAR yields $\Gamma, X<:P, \Delta \vdash X <: N$, as required. \square

Type Safety (Challenges 2A and 2B): Lemmas

Again, we give the proof of type safety for the whole system including records and pattern matching. We begin with some preliminary technical facts about the typing and subtype relations.

A.4 LEMMA [PERMUTATION FOR TYPING]: Suppose that Δ is a well-formed permutation of Γ . If $\Gamma \vdash t : T$, then $\Delta \vdash t : T$. \square

Proof: By straightforward induction on derivations. \square

A.5 LEMMA [WEAKENING FOR SUBTYPING AND TYPING]:

1. If $\Gamma \vdash S <: T$ and $\Gamma, x:U$ is well-formed, then $\Gamma, x:U \vdash S <: T$.
2. If $\Gamma \vdash S <: T$ and $\Gamma, X<:U$ is well-formed, then $\Gamma, X<:U \vdash S <: T$.
3. If $\Gamma \vdash t : T$ and $\Gamma, x:U$ is well-formed, then $\Gamma, x:U \vdash t : T$.
4. If $\Gamma \vdash t : T$ and $\Gamma, X<:U$ is well-formed, then $\Gamma, X<:U \vdash t : T$. \square
5. If $\Gamma \vdash S <: T$ and Γ, Δ is well-formed, then $\Gamma, \Delta \vdash S <: T$.
6. If $\Gamma \vdash t : T$ and Γ, Δ is well-formed, then $\Gamma, \Delta \vdash t : T$.

Proof: The proofs for parts (1) and (2) proceed by straightforward induction on derivations. The SA-ALL cases require permutation on subtyping derivations (Lemma A.2, part 1). The proofs for parts (3) and (4) also proceed by straightforward induction on derivations, using parts (1) and (2) in the T-SUB and T-TAPP cases. The T-ABS, T-TABS, and T-LET cases require permutation on typing derivations (Lemma A.4). Parts (5) and (6) follow by induction on the number of bindings in Δ , using the first four parts. \square

A.6 LEMMA [STRENGTHENING]: If $\Gamma, x:Q, \Delta \vdash S <: T$, then $\Gamma, \Delta \vdash S <: T$. \square

Proof: Typing assumptions play no role in subtype derivations. \square

The proof of type preservation relies on several lemmas relating substitution with the typing and subtype relations. First, we state an analog for the typing relation of the narrowing lemma for subtyping (Lemma A.3, part 2).

A.7 LEMMA [NARROWING FOR THE TYPING RELATION]: If $\Gamma, X<:Q, \Delta \vdash t : T$ and $\Gamma \vdash P <: Q$, then $\Gamma, X<:P, \Delta \vdash t : T$. \square

Proof: Straightforward induction, using Lemma A.3(2) for the T-SUB case. \square

Next, we have the usual lemma relating substitution and the typing relation.

A.8 LEMMA [SUBSTITUTION PRESERVES TYPING]: If $\Gamma, \mathbf{x}:\mathbf{Q}, \Delta \vdash \mathbf{t} : \mathbf{T}$ and $\Gamma \vdash \mathbf{q} : \mathbf{Q}$, then $\Gamma, \Delta \vdash [\mathbf{x} \mapsto \mathbf{q}]\mathbf{t} : \mathbf{T}$. \square

Proof: Induction on a derivation of $\Gamma, \mathbf{x}:\mathbf{Q}, \Delta \vdash \mathbf{t} : \mathbf{T}$, using the properties above. In particular, we use Lemma A.6 in the T-TAPP and T-SUB cases. \square

Since we may substitute types for type variables during reduction, we also need a lemma relating type substitution and typing. The proof of this lemma (specifically, the T-SUB case) depends on a new lemma relating substitution and subtyping.

A.9 DEFINITION: We write $[\mathbf{X} \mapsto \mathbf{S}]\Gamma$ for the context obtained by substituting \mathbf{S} for \mathbf{X} in the right-hand sides of all of the bindings in Γ . \square

A.10 LEMMA [TYPE SUBSTITUTION PRESERVES SUBTYPING]: If $\Gamma, \mathbf{X} <: \mathbf{Q}, \Delta \vdash \mathbf{S} <: \mathbf{T}$ and $\Gamma \vdash \mathbf{P} <: \mathbf{Q}$, then $\Gamma, [\mathbf{X} \mapsto \mathbf{P}]\Delta \vdash [\mathbf{X} \mapsto \mathbf{P}]\mathbf{S} <: [\mathbf{X} \mapsto \mathbf{P}]\mathbf{T}$. \square

Note that we need to substitute for \mathbf{X} only in the part of the environment that *follows* the binding of \mathbf{X} , since our conventions about scoping require that the types to the left of the binding of \mathbf{X} do not contain \mathbf{X} .

Proof: By induction on a derivation of $\Gamma, \mathbf{X} <: \mathbf{Q}, \Delta \vdash \mathbf{S} <: \mathbf{T}$. The only interesting cases are the following:

Case SA-TRANS-TVAR: $\mathbf{S} = \mathbf{Y}$
 $\mathbf{Y} <: \mathbf{U} \in (\Gamma, \mathbf{X} <: \mathbf{Q}, \Delta)$
 $\Gamma, \mathbf{X} <: \mathbf{Q}, \Delta \vdash \mathbf{U} <: \mathbf{T}$

There are two subcases to consider. If $\mathbf{Y} \neq \mathbf{X}$, then the result follows from considering cases on whether $\mathbf{Y} <: \mathbf{U} \in \Gamma$ or $\mathbf{Y} <: \mathbf{U} \in \Delta$, the induction hypothesis, and SA-TRANS-TVAR. (Note that $[\mathbf{X} \mapsto \mathbf{P}]\mathbf{U} = \mathbf{U}$ if $\mathbf{Y} <: \mathbf{U} \in \Gamma$.)

On the other hand, if $\mathbf{Y} = \mathbf{X}$, then we have $\mathbf{U} = \mathbf{Q}$. By the induction hypothesis, $\Gamma, [\mathbf{X} \mapsto \mathbf{P}]\Delta \vdash [\mathbf{X} \mapsto \mathbf{P}]\mathbf{Q} <: [\mathbf{X} \mapsto \mathbf{P}]\mathbf{T}$. Note that $[\mathbf{X} \mapsto \mathbf{P}]\mathbf{Q} = \mathbf{Q}$. We then have $\Gamma \vdash \mathbf{P} <: [\mathbf{X} \mapsto \mathbf{P}]\mathbf{Q}$ since $\Gamma \vdash \mathbf{P} <: \mathbf{Q}$. By weakening, $\Gamma, [\mathbf{X} \mapsto \mathbf{P}]\Delta \vdash \mathbf{P} <: [\mathbf{X} \mapsto \mathbf{P}]\mathbf{Q}$, and by transitivity, $\Gamma, [\mathbf{X} \mapsto \mathbf{P}]\Delta \vdash \mathbf{P} <: [\mathbf{X} \mapsto \mathbf{P}]\mathbf{T}$. Therefore, $\Gamma, [\mathbf{X} \mapsto \mathbf{P}]\Delta \vdash [\mathbf{X} \mapsto \mathbf{P}]\mathbf{X} <: [\mathbf{X} \mapsto \mathbf{P}]\mathbf{T}$.

Case SA-ALL: $\mathbf{S} = \forall \mathbf{Z} <: \mathbf{R}_1 . \mathbf{S}_2$ $\mathbf{T} = \forall \mathbf{Z} <: \mathbf{U}_1 . \mathbf{T}_2$
 $\Gamma, \mathbf{X} <: \mathbf{Q}, \Delta \vdash \mathbf{U}_1 <: \mathbf{R}_1$
 $\Gamma, \mathbf{X} <: \mathbf{Q}, \Delta, \mathbf{Z} <: \mathbf{U}_1 \vdash \mathbf{S}_2 <: \mathbf{T}_2$

By the induction hypothesis,

$\Gamma, [\mathbf{X} \mapsto \mathbf{P}]\Delta, \mathbf{Z} <: [\mathbf{X} \mapsto \mathbf{P}]\mathbf{U}_1 \vdash [\mathbf{X} \mapsto \mathbf{P}]\mathbf{S}_2 <: [\mathbf{X} \mapsto \mathbf{P}]\mathbf{T}_2$, and
 $\Gamma, [\mathbf{X} \mapsto \mathbf{P}]\Delta \vdash [\mathbf{X} \mapsto \mathbf{P}]\mathbf{U}_1 <: [\mathbf{X} \mapsto \mathbf{P}]\mathbf{R}_1$.

By SA-ALL,

$\Gamma, [\mathbf{X} \mapsto \mathbf{P}]\Delta \vdash \forall \mathbf{Z} <: [\mathbf{X} \mapsto \mathbf{P}]\mathbf{R}_1 . [\mathbf{X} \mapsto \mathbf{P}]\mathbf{S}_2 <: \forall \mathbf{Z} <: [\mathbf{X} \mapsto \mathbf{P}]\mathbf{U}_1 . [\mathbf{X} \mapsto \mathbf{P}]\mathbf{T}_2$, i.e.,
 $\Gamma, [\mathbf{X} \mapsto \mathbf{P}]\Delta \vdash [\mathbf{X} \mapsto \mathbf{P}](\forall \mathbf{Z} <: \mathbf{R}_1 . \mathbf{S}_2) <: [\mathbf{X} \mapsto \mathbf{P}](\forall \mathbf{Z} <: \mathbf{U}_1 . \mathbf{T}_2)$. \square

A similar lemma relates type substitution and typing.

A.11 LEMMA [TYPE SUBSTITUTION PRESERVES TYPING]: If $\Gamma, X <: Q, \Delta \vdash t : T$ and $\Gamma \vdash P <: Q$, then $\Gamma, [X \mapsto P]\Delta \vdash [X \mapsto P]t : [X \mapsto P]T$. \square

Proof: By induction on a derivation of $\Gamma, X <: Q, \Delta \vdash t : T$. We give only the interesting cases.

Case T-TAPP: $t = t_1 [T_2]$ $\Gamma, X <: Q, \Delta \vdash t_1 : \forall Z <: T_{11}. T_{12}$
 $T = [Z \mapsto T_2]T_{12}$ $\Gamma, X <: Q, \Delta \vdash T_2 <: T_{11}$

By the induction hypothesis, $\Gamma, [X \mapsto P]\Delta \vdash [X \mapsto P]t_1 : [X \mapsto P](\forall Z <: T_{11}. T_{12})$, i.e., $\Gamma, [X \mapsto P]\Delta \vdash [X \mapsto P]t_1 : \forall Z <: [X \mapsto P]T_{11}. [X \mapsto P]T_{12}$. By the preservation of subtyping under substitution (Lemma A.10), we have $\Gamma, [X \mapsto P]\Delta \vdash [X \mapsto P]T_2 <: [X \mapsto P]T_{11}$. By T-TAPP,

$\Gamma, [X \mapsto P]\Delta \vdash [X \mapsto P]t_1 [X \mapsto P]T_2 : [Z \mapsto [X \mapsto P]T_2]([X \mapsto P]T_{12})$, i.e.
 $\Gamma, [X \mapsto P]\Delta \vdash [X \mapsto P](t_1 [T_2]) : [X \mapsto P]([Z \mapsto T_2]T_{12})$.

Case T-SUB: $\Gamma, X <: Q, \Delta \vdash t : S$ $\Gamma, X <: Q, \Delta \vdash S <: T$

By the induction hypothesis, $\Gamma, [X \mapsto P]\Delta \vdash [X \mapsto P]t : [X \mapsto P]S$. By the preservation of subtyping under substitution (Lemma A.10), we have $\Gamma, [X \mapsto P]\Delta \vdash [X \mapsto P]S <: [X \mapsto P]T$, and the result follows by T-SUB. \square

Next, we establish some simple structural properties of subtyping.

A.12 LEMMA [INVERSION OF SUBTYPING (RIGHT TO LEFT)]:

1. If $\Gamma \vdash S <: X$, then S is a type variable.
2. If $\Gamma \vdash S <: T_1 \rightarrow T_2$, then either S is a type variable or else S has the form $S_1 \rightarrow S_2$, with $\Gamma \vdash T_1 <: S_1$ and $\Gamma \vdash S_2 <: T_2$.
3. If $\Gamma \vdash S <: \forall X <: U_1. T_2$, then either S is a type variable or else S has the form $\forall X <: Q_1. S_2$ with $\Gamma, X <: U_1 \vdash S_2 <: T_2$ and $\Gamma \vdash U_1 <: Q_1$.
4. If $\Gamma \vdash S <: \{l_i : T_i \mid i \in 1..n\}$, then either S is a type variable or else S has the form $\{k_j : Q_j \mid j \in 1..m\}$ with $\{l_i \mid i \in 1..n\} \subseteq \{k_j \mid j \in 1..m\}$ and where $l_i = k_j$ implies $\Gamma \vdash Q_j <: T_i$ for every i and j . \square

Proof: Each part is immediate from the definition of the subtyping relation. \square

Lemma A.12 is used, in turn, to establish one straightforward structural property of the typing relation that is needed in the critical cases of the type preservation proof.

A.13 LEMMA:

1. If $\Gamma \vdash \lambda x : S_1. s_2 : T$ and $\Gamma \vdash T <: U_1 \rightarrow U_2$, then $\Gamma \vdash U_1 <: S_1$ and there is some S_2 such that $\Gamma, x : S_1 \vdash s_2 : S_2$ and $\Gamma \vdash S_2 <: U_2$.
2. If $\Gamma \vdash \lambda X <: S_1. s_2 : T$ and $\Gamma \vdash T <: \forall X <: U_1. U_2$, then $\Gamma \vdash U_1 <: S_1$ and there is some S_2 such that $\Gamma, X <: U_1 \vdash s_2 : S_2$ and $\Gamma, X <: U_1 \vdash S_2 <: U_2$.
3. If $\Gamma \vdash \{l_i = t_i \mid i \in 1..n\} : S$ and $\Gamma \vdash S <: \{k_j : T_j \mid j \in 1..m\}$, then $\{k_j \mid j \in 1..m\} \subseteq \{l_i \mid i \in 1..n\}$ and $l_i = k_j$ implies $\Gamma \vdash t_i : T_j$ for every i and j . \square

Proof: Straightforward induction on typing derivations, using Lemma A.12 for the induction case (rule T-SUB). \square

Type Safety (Challenges 2A and 2B): Progress

The progress theorem for F_{\leq} is relatively straightforward. We begin by recording a canonical forms property telling us the possible shapes of closed values of arrow, record, and quantifier types.

A.14 LEMMA [CANONICAL FORMS]:

1. If v is a closed value of type $T_1 \rightarrow T_2$, then v has the form $\lambda x : S_1 . t_2$.
2. If v is a closed value of type $\{l_i : T_i \mid i \in 1..n\}$, then v has the form $\{k_j = v_j \mid j \in 1..m\}$ with $\{l_i \mid i \in 1..n\} \subseteq \{k_j \mid j \in 1..m\}$.
3. If v is a closed value of type $\forall X < : T_1 . T_2$, then v has the form $\lambda X < : S_1 . t_2$. \square

Proof: All parts proceed by induction on typing derivations; we give the argument only for the third part. (The others are similar.) By inspection of the typing rules, it is clear that the final rule in a derivation of $\vdash v : \forall X < : T_1 . T_2$ must be either T-TABS or T-SUB. If it is T-TABS, then the desired result is immediate. So suppose the last rule is T-SUB. From the premises of this rule, we have $\vdash v : S$ and $\vdash S < : \forall X < : T_1 . T_2$. From the inversion lemma (A.12), we know that S has the form $\forall X < : Q_1 . Q_2$ (S can not be a type variable since the typing environment is empty). The result now follows from the induction hypothesis. \square

We now observe that any non-value term can be decomposed into an evaluation context and a subterm which can take a step.

A.15 LEMMA: If $\vdash t : T$, then either t is a value or there exists an evaluation context E and term t_0 such that $t = E[t_0]$ and $t_0 \longrightarrow t'_0$. \square

Proof: By induction on a derivation of $\vdash t : T$. We give only the interesting cases.

Case T-APP: $t = t_1 \ t_2 \quad \vdash t_1 : T_{11} \rightarrow T_{12} \quad \vdash t_2 : T_{11}$

By the induction hypothesis applied to t_1 , we have two subcases to consider.

Subcase: t_1 is a value

We now apply the induction hypothesis to t_2 . If t_2 is a value, then by the canonical forms lemma (A.14), t_1 must have the form $\lambda x : S_1 . t_3$. We then have $t \longrightarrow [x \mapsto t_2]t_3$ by E-APPABS. Thus, we can take $E = [-]$ and $t_0 = t$. Otherwise, $t_2 = E'[t_3]$ with $t_3 \longrightarrow t'_3$. Thus, we can take $E = t_1 \ E'$ and $t_0 = t_3$.

Subcase: $t_1 = E'[t_3] \quad t_3 \longrightarrow t'_3$

In this case, take $E = E' \ t_2$ and $t_0 = t_3$.

Case T-TAPP: $t = t_1 \ [T_2]$

This case is similar to the T-APP case.

Case T-LET: $\mathbf{t} = \text{let } \mathbf{p}=\mathbf{t}_1 \text{ in } \mathbf{t}_2 \quad \vdash \mathbf{t}_1 : \mathbf{T}_1$

By the induction hypothesis, we have two cases to consider.

Subcase: \mathbf{t}_1 is a value

Then $\mathbf{t} \longrightarrow \text{match}(\mathbf{p}, \mathbf{t}_1)\mathbf{t}_2$ by E-LETV. Take $E = [-]$ and $\mathbf{t}_0 = \mathbf{t}$.

Subcase: $\mathbf{t}_1 = E'[\mathbf{t}_3] \quad \mathbf{t}_3 \longrightarrow \mathbf{t}'_3$

In this case, take $E = \text{let } \mathbf{p}=E' \text{ in } \mathbf{t}_3$ and $\mathbf{t}_0 = \mathbf{t}_3$.

Case T-PROJ: $\mathbf{t} = \mathbf{t}_1.l_j \quad \vdash \mathbf{t}_1 : \{\mathbf{l}_i : \mathbf{T}_i \text{ } i \in 1..n\}$

By the induction hypothesis, we have two cases to consider.

Subcase: \mathbf{t}_1 is a value

By the canonical forms lemma (A.14), \mathbf{t}_1 must have the form $\{\mathbf{h}_k = \mathbf{v}_k \text{ } k \in 1..m\}$ with $\{\mathbf{l}_i \text{ } i \in 1..n\} \subseteq \{\mathbf{h}_k \text{ } k \in 1..m\}$. By E-PROJCD, there is a \mathbf{t}' such that $\mathbf{t} \longrightarrow \mathbf{t}'$. Thus, we can take $E = [-]$ and $\mathbf{t}_0 = \mathbf{t}$ in this case.

Subcase: $\mathbf{t}_1 = E'[\mathbf{t}_2] \quad \mathbf{t}_2 \longrightarrow \mathbf{t}'_2$

In this case, take $E = E'.l_j$ and $\mathbf{t}_0 = \mathbf{t}_2$.

Case T-RCD: $\mathbf{t} = \{\mathbf{l}_i = \mathbf{t}_i \text{ } i \in 1..n\} \quad \mathbf{T} = \{\mathbf{l}_i : \mathbf{T}_i \text{ } i \in 1..n\}$
for each $i \quad \vdash \mathbf{t}_i : \mathbf{T}_i$

If every \mathbf{t}_i is a value, then we're done. Otherwise, by the induction hypothesis, there is a least j such that $\mathbf{t}_j = E'[\mathbf{s}_0]$ and $\mathbf{s}_0 \longrightarrow \mathbf{s}'_0$. In this case, we can take $E = \{\mathbf{l}_i = \mathbf{t}_i \text{ } i \in 1..j-1, \mathbf{l}_j = E', \mathbf{l}_k = \mathbf{t}_k \text{ } k \in j+1..n\}$ and $\mathbf{t}_0 = \mathbf{s}_0$. \square

The proof of progress is now straightforward.

A.16 THEOREM [PROGRESS]: If \mathbf{t} is a closed, well-typed F_{\leq} term (i.e., if $\vdash \mathbf{t} : \mathbf{T}$ for some \mathbf{T}), then either \mathbf{t} is a value or else there is some \mathbf{t}' with $\mathbf{t} \longrightarrow \mathbf{t}'$. \square

Proof: By Lemma A.15, either \mathbf{t} is a value or else there is an evaluation context E and term \mathbf{t}_0 such that $\mathbf{t} = E[\mathbf{t}_0]$ and $\mathbf{t}_0 \longrightarrow \mathbf{t}'_0$. In the latter case, take $\mathbf{t}' = E[\mathbf{t}'_0]$ and observe that $\mathbf{t} \longrightarrow \mathbf{t}'$ by E-CTX. \square

Type Safety (Challenges 2A and 2B): Preservation

We begin the proof of preservation by proving a substitution lemma for pattern matching on records.

A.17 LEMMA [MATCHED PATTERNS PRESERVE TYPING]: Suppose $\vdash \mathbf{p} : \mathbf{T}_1 \Rightarrow \Delta$, that $\Gamma \vdash \mathbf{v}_0 : \mathbf{T}_1$, and that $\Gamma, \Delta \vdash \mathbf{t}_2 : \mathbf{T}_2$. Then $\Gamma \vdash \text{match}(\mathbf{p}, \mathbf{v}_0)\mathbf{t}_2 : \mathbf{T}_2$. \square

Proof: By induction on the derivation of $\vdash \mathbf{p} : \mathbf{T}_1 \Rightarrow \Delta$.

Case P-VAR: $\mathbf{p} = \mathbf{x}:\mathbf{T}_1 \quad \Delta = \mathbf{x} : \mathbf{T}_1$

Then $\text{match}(\mathbf{x}:\mathbf{T}_1, \mathbf{v}_0) = [\mathbf{x} \mapsto \mathbf{v}_0]$, and the result follows from Lemma A.8.

Case P-RCD: $\mathbf{p} = \{\mathbf{l}_i = \mathbf{p}_i \mid i \in 1..n\}$ $\mathbf{T}_1 = \{\mathbf{l}_i : \mathbf{S}_i \mid i \in 1..n\}$
for each i $\vdash \mathbf{p}_i : \mathbf{S}_i \Rightarrow \Delta_i$
 $\Delta = \Delta_n, \dots, \Delta_1$

By Lemma A.13, \mathbf{v}_0 has the form $\{\mathbf{g}_j = \mathbf{v}_j \mid j \in 1..m\}$, with $\{\mathbf{l}_i \mid i \in 1..n\} \subseteq \{\mathbf{g}_j \mid j \in 1..m\}$ and $\mathbf{g}_j = \mathbf{l}_i$ implies $\Gamma \vdash \mathbf{v}_j : \mathbf{S}_i$ for every i and j . By definition, we then have $\text{match}(\mathbf{p}, \mathbf{v}_0) = (\sigma_n \circ \dots \circ \sigma_1)$ where $\mathbf{l}_i = \mathbf{g}_j$ implies $\text{match}(\mathbf{p}_i, \mathbf{v}_j) = \sigma_i$ for every i and j . Starting from $\Gamma, \Delta \vdash \mathbf{t}_2 : \mathbf{T}_2$, we iteratively apply the induction hypothesis with $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n$, to obtain $\Gamma \vdash \text{match}(\mathbf{p}, \mathbf{v}_1)\mathbf{t}_2 : \mathbf{T}_2$. More formally, the result follows by induction on n , applying weakening (Lemma A.5) to the judgments $\Gamma \vdash \mathbf{v}_j : \mathbf{S}_i$ as needed. \square

The following lemma relates evaluation contexts and the typing relation.

A.18 LEMMA:

1. If $\Gamma \vdash E[\mathbf{t}] : \mathbf{T}$ and if for all \mathbf{T}_0 , $\Gamma \vdash \mathbf{t} : \mathbf{T}_0$ implies $\Gamma \vdash \mathbf{t}' : \mathbf{T}_0$, then $\Gamma \vdash E[\mathbf{t}'] : \mathbf{T}$.
2. If $\Gamma \vdash E[\mathbf{t}] : \mathbf{T}$, then $\Gamma \vdash \mathbf{t} : \mathbf{T}_0$ for some \mathbf{T}_0 . \square

Proof: Both parts are proven by induction on the structure of evaluation contexts. In each case, we consider the last rule used in the derivation of $\Gamma \vdash E[\mathbf{t}] : \mathbf{T}$ and apply the induction hypothesis. \square

We now prove that immediate reduction preserves the types of terms.

A.19 LEMMA: If $\Gamma \vdash \mathbf{t} : \mathbf{T}$ and $\mathbf{t} \longrightarrow \mathbf{t}'$, where the derivation of $\mathbf{t} \longrightarrow \mathbf{t}'$ ends in a rule other than E-CTX, then $\Gamma \vdash \mathbf{t}' : \mathbf{T}$. \square

Proof: By induction on a derivation of $\Gamma \vdash \mathbf{t} : \mathbf{T}$. We give only the interesting cases.

Case T-APP: $\mathbf{t} = \mathbf{t}_1 \mathbf{t}_2$ $\Gamma \vdash \mathbf{t}_1 : \mathbf{T}_{11} \rightarrow \mathbf{T}_{12}$ $\mathbf{T} = \mathbf{T}_{12}$ $\Gamma \vdash \mathbf{t}_2 : \mathbf{T}_{11}$

The derivation of $\mathbf{t} \longrightarrow \mathbf{t}'$ must end in E-APPABS. Thus, $\mathbf{t}_1 = \lambda \mathbf{x} : \mathbf{U}_{11}. \mathbf{u}_{12}$ and $\mathbf{t}' = [\mathbf{x} \mapsto \mathbf{t}_2]\mathbf{u}_{12}$. By Lemma A.13, $\Gamma, \mathbf{x} : \mathbf{U}_{11} \vdash \mathbf{u}_{12} : \mathbf{U}_{12}$ for some \mathbf{U}_{12} with $\Gamma \vdash \mathbf{T}_{11} \prec : \mathbf{U}_{11}$ and $\Gamma \vdash \mathbf{U}_{12} \prec : \mathbf{T}_{12}$. By narrowing (Lemma A.7) and the preservation of typing under substitution (Lemma A.8), $\Gamma \vdash [\mathbf{x} \mapsto \mathbf{t}_2]\mathbf{u}_{12} : \mathbf{U}_{12}$, from which we obtain $\Gamma \vdash [\mathbf{x} \mapsto \mathbf{t}_2]\mathbf{u}_{12} : \mathbf{T}_{12}$ by T-SUB.

Case T-PROJ: $\mathbf{t} = \mathbf{t}_1.l_j$ $\mathbf{T} = \mathbf{T}_j$ $\Gamma \vdash \mathbf{t}_1 : \{\mathbf{l}_i : \mathbf{T}_i \mid i \in 1..n\}$

The derivation of $\mathbf{t} \longrightarrow \mathbf{t}'$ must end in E-PROJRCD. Thus, \mathbf{t}_1 is a record value $\{\mathbf{h}_k = \mathbf{v}_k \mid k \in 1..m\}$. By Lemma A.13, we have $\{\mathbf{l}_i \mid i \in 1..n\} \subseteq \{\mathbf{h}_k \mid k \in 1..m\}$ and $\mathbf{l}_i = \mathbf{h}_k$ implies $\Gamma \vdash \mathbf{v}_k : \mathbf{T}_i$ for every i and j . By E-PROJRCD, we have $\mathbf{t}' = \mathbf{v}_k$ where $\mathbf{h}_k = \mathbf{l}_j$. Thus, $\Gamma \vdash \mathbf{t}' : \mathbf{T}_j$ as required.

Case T-TAPP: $\mathbf{t} = \mathbf{t}_1 [\mathbf{T}_2]$ $\Gamma \vdash \mathbf{t} : \forall \mathbf{X} \prec : \mathbf{T}_{11}. \mathbf{T}_{12}$
 $\mathbf{T} = [\mathbf{X} \mapsto \mathbf{T}_2]\mathbf{T}_{12}$ $\Gamma \vdash \mathbf{T}_2 \prec : \mathbf{T}_{11}$

The derivation of $\mathbf{t} \longrightarrow \mathbf{t}'$ must end in E-TAPPTABS. Thus, $\mathbf{t}_1 = \lambda \mathbf{X} \prec : \mathbf{U}_{11}. \mathbf{u}_{12}$ and $\mathbf{t}' = [\mathbf{X} \mapsto \mathbf{T}_2]\mathbf{u}_{12}$. By Lemma A.13, $\Gamma \vdash \mathbf{T}_{11} \prec : \mathbf{U}_{11}$ and $\Gamma, \mathbf{X} \prec : \mathbf{T}_{11} \vdash \mathbf{u}_{12} : \mathbf{U}_{12}$ for some \mathbf{U}_{12} with $\Gamma, \mathbf{X} \prec : \mathbf{T}_{11} \vdash \mathbf{U}_{12} \prec : \mathbf{T}_{12}$. By the preservation of typing under substitution (Lemma A.11), $\Gamma \vdash [\mathbf{X} \mapsto \mathbf{T}_2]\mathbf{u}_{12} : [\mathbf{X} \mapsto \mathbf{T}_2]\mathbf{U}_{12}$, from which $\Gamma \vdash [\mathbf{X} \mapsto \mathbf{T}_2]\mathbf{u}_{12} : [\mathbf{X} \mapsto \mathbf{T}_2]\mathbf{T}_{12}$ follows by Lemma A.10 and T-SUB.

Case T-LET: $\mathfrak{t} = \text{let } \mathfrak{p}=\mathfrak{t}_1 \text{ in } \mathfrak{t}_2$
 $\Gamma \vdash \mathfrak{t}_1 : \mathsf{T}_1 \quad \vdash \mathfrak{p} : \mathsf{T}_1 \Rightarrow \Delta \quad \Gamma, \Delta \vdash \mathfrak{t}_2 : \mathsf{T}$

The derivation of $\mathfrak{t} \longrightarrow \mathfrak{t}'$ must end in E-LETV. Thus, \mathfrak{t}_1 is a value and $\mathfrak{t} \longrightarrow \text{match}(\mathfrak{p}, \mathfrak{t}_1)\mathfrak{t}_2$. The result then follows by Lemma A.17. \square

Finally, we prove the main preservation theorem.

A.20 THEOREM [PRESERVATION]: If $\Gamma \vdash \mathfrak{t} : \mathsf{T}$ and $\mathfrak{t} \longrightarrow \mathfrak{t}'$, then $\Gamma \vdash \mathfrak{t}' : \mathsf{T}$. \square

Proof: By induction on a derivation of $\mathfrak{t} \longrightarrow \mathfrak{t}'$. Suppose the derivation ends in E-CTX. Then, $\mathfrak{t} = E[\mathfrak{t}_0]$, $\mathfrak{t}_0 \longrightarrow \mathfrak{t}'_0$, and $\mathfrak{t}' = E[\mathfrak{t}'_0]$. By Lemma A.18, $\Gamma \vdash \mathfrak{t}_0 : \mathsf{T}_0$ for some T_0 . By the induction hypothesis, for any T_0 such that $\Gamma \vdash \mathfrak{t}_0 : \mathsf{T}_0$, we have $\Gamma \vdash \mathfrak{t}'_0 : \mathsf{T}_0$. The result then follows by Lemma A.18. In all other cases, the result follows by Lemma A.19. \square