Just-in-Time Static Type Checking for Dynamic Languages

Brianna M. Ren Jeffrey S. Foster

University of Maryland, College Park, USA

{bren, jfoster}@cs.umd.edu

Abstract

Dynamic languages such as Ruby, Python, and JavaScript have many compelling benefits, but the lack of static types means subtle errors can remain latent in code for a long time. While many researchers have developed various systems to bring some of the benefits of static types to dynamic languages, prior approaches have trouble dealing with metaprogramming, which generates code as the program executes. In this paper, we propose Hummingbird, a new system that uses a novel technique, *just-in-time static type checking*, to type check Ruby code even in the presence of metaprogramming. In Hummingbird, method type signatures are gathered dynamically at run-time, as those methods are created. When a method is called, Hummingbird statically type checks the method body against current type signatures. Thus, Hummingbird provides thorough static checks on a per-method basis, while also allowing arbitrarily complex metaprogramming. For performance, Hummingbird memoizes the static type checking pass, invalidating cached checks only if necessary. We formalize Hummingbird using a core, Ruby-like language and prove it sound. To evaluate Hummingbird, we applied it to six apps, including three that use Ruby on Rails, a powerful framework that relies heavily on metaprogramming. We found that all apps typecheck successfully using Hummingbird, and that Hummingbird's performance overhead is reasonable. We applied Hummingbird to earlier versions of one Rails app and found several type errors that had been introduced and then fixed. Lastly, we demonstrate using Hummingbird in Rails development mode to typecheck an app as live updates are applied to it.

Categories and Subject Descriptors F.3.2 [Semantics of Programming Languages]: Program analysis

General Terms Languages; Theory; Verification

Keywords Type checking; dynamic languages; Ruby

1. Introduction

Many researchers have explored ways to bring the benefits of static typing to dynamic languages [1–4, 12, 14, 18, 20–22, 25, 28, 37, 38]. However, many of these prior systems do not work well in the presence of *metaprogramming*, in which code the program relies on is generated as the program executes. The challenge is that purely static systems cannot analyze metaprogramming code, which is often complicated and convoluted; and prior mixed static/dynamic systems are either cumbersome or make certain limiting assumptions. (Section 6 discusses prior work in detail.)

In this paper, we introduce Hummingbird,¹ a type checking system for Ruby that solves this problem using a new approach we call *just-in-time static type checking*. In Hummingbird, user-provided type annotations actually execute at *run-time*, adding types to an environment that is maintained during execution. As metaprogramming code creates new methods, it, too, executes type annotations to assign types to dynamically created methods. Then whenever a method *m* is called, Hummingbird *statically* type checks *m*'s body in the current dynamic type environment. More precisely, Hummingbird checks that *m* calls methods at their types as annotated, and that *m* itself matches its annotated type. Moreover, Hummingbird caches the type check so that it need not recheck *m* at the next call unless the dynamic type environment has changed in a way that affects *m*.

Just-in-time static type checking provides a highly effective tradeoff between purely dynamic and purely static type checking. On the one hand, metaprogramming code is very challenging to analyze statically, but in our experience it is easy to create type annotations at run time for generated code. On the other hand, by statically analyzing whole method bodies, we catch type errors earlier than a purely

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

PLDI'16, June 13–17, 2016, Santa Barbara, CA, USA
© 2016 ACM. 978-1-4503-4261-2/16/06...\$15.00
<http://dx.doi.org/10.1145/2908080.2908127>

¹ A hummingbird can dynamically flap its wings while statically hovering.

dynamic system, and we can soundly reason about all possible execution paths within type checked methods. (Section 2 shows how several examples of metaprogramming are handled by Hummingbird.)

To ensure our approach to type checking is correct, we formalize Hummingbird using a core, Ruby-like language in which method creation and method type annotation can occur at arbitrary points during execution. We provide a flow-sensitive type checking system and a dynamic semantics that invokes the type system at method entry, caching the resulting typing proof. Portions of the cache may be invalidated as new methods are defined or type annotations are changed. We prove soundness for our type system. (Section 3 presents the formalism.)

Our implementation of Hummingbird piggybacks on two prior systems we developed. We use the Ruby Intermediate Language [8, 13] to parse input Ruby files and translate them to simplified control-flow graphs. We use RDL [26, 34], a Ruby contract system, to intercept method calls and to represent and store method type signatures at run time. Hummingbird supports an extensive set of typing features, including union types, intersection types, code blocks (anonymous functions), generics, modules, and type casts, among others. (Section 4 describes our implementation.)

We evaluated Hummingbird by applying it to six Ruby apps. Three use Ruby on Rails (just “Rails” below), a popular, sophisticated web app framework that uses metaprogramming heavily both to make Rails code more compact and expressive and to support “convention over configuration.” We should emphasize that Rails’s use of metaprogramming makes static analysis of it very challenging [17]. Two apps use other styles of metaprogramming, and the last app does not use metaprogramming, as a baseline.

We found that all of our subject apps type check successfully using Hummingbird, and that dynamically generated types are essential for the apps that use metaprogramming. We also found that Hummingbird’s performance overhead ranges from 19% to 469%, which is much better than prior approaches [18, 28], and that caching is essential to achieving this performance. For one Rails app, we ran type checking on many prior versions, and we found a total of six type errors that had been introduced and then later fixed. We also ran the app in Rails *development mode*, which reloads files as they are edited, to demonstrate how Hummingbird type check caching behaves in the presence of modified methods. (Section 5 reports on our results.)

In summary, we believe Hummingbird is an important step forward in our ability to bring the benefits of static typing to dynamic languages while still supporting flexible and powerful metaprogramming features.

2. Overview

We begin our presentation by showing some uses of metaprogramming in Ruby and the corresponding Hummingbird

```

1 class Talk < ActiveRecord::Base
2   belongs_to :owner, :class_name => "User"
3   ....
4   type :owner?, "(User) → %bool"
5   def owner?(user)
6     return owner == user
7   end end
8
9 module ActiveRecord::Associations::ClassMethods
10  pre(: belongs_to) do |*args|
11    hmi = args[0]
12    options = args[1]
13    hm = hmi.to_s
14    cn = options[:class_name] if options
15    hmu = cn ? cn : hm.singularize.camelize
16    type hm.singularize, "()" → #{hmu}"
17    type "#{hm.singularize}=",
18          "({#{hmu}) → #{hmu}"
19    true
20  end end

```

Figure 1. Ruby on Rails Metaprogramming.

type checking process. The examples below are from the experiments in Section 5.

Rails Associations. The top of Figure 1 shows an excerpt from the *Talks* Rails app. This code defines a class *Talk* that is a *model* in Rails, meaning an instance of *Talk* represents a row in the *talks* database table. The change in case and pluralization here is not an accident—Rails favors “convention over configuration,” meaning many relationships that would otherwise be specified via configuration are instead implicitly expressed by using similar or the same name for things.

In this app, every *talk* is owned by a *user*, which in implementation terms means a *Talk* instance has a foreign key *owner_id* indicating the owner, which is an instance of class *User* (not shown). The existence of that relationship is defined on line 2. Here it may look like *belongs_to* is a keyword, but in fact it is simply a method call. The call passes the *symbol* (an interned string) *:owner* as the first argument, and the second argument is a hash that maps symbol *:class_name* to string “*User*”.

Now consider the *owner?* method, defined on line 5. Just before the method, we introduce a type annotation indicating the method takes a *User* and returns a boolean. Given such an annotation, Hummingbird’s goal is to check whether the method body has the indicated type.² This should be quite simple in this case, as the body of *owner?* just calls no-argument method *owner* and checks whether the result is equal to *user*.

However, if we examine the remaining code of *Talk* (not shown), we discover that *owner* is not defined anywhere in

² In practice type takes another argument to tell Hummingbird to type check the body, in contrast to library and framework methods whose types are trusted. We elide this detail for simplicity.

the class! Instead, this method is created at run-time by `belongs_to`. More specifically, when `belongs_to` is called, it defines several convenience methods that perform appropriate SQL queries for the relationship [23], in this case to get the User instance associated with the Talk’s owner. Thus, as we can see, it is critical for Hummingbird to handle such dynamically created methods even to type check simple Rails code.

Our solution is to instrument `belongs_to` so that, just as it creates a method dynamically, it also creates method type signatures dynamically. The code on lines 9–20 of Figure 1 accomplishes this. Hummingbird is built on RDL, a Ruby contract system for specifying pre- and postconditions [26, 34]. The precondition is specified via a *code block*—an anonymous function (i.e., a lambda) delimited by `do...end`—passed to `pre`. Here the code block trivially returns `true` so the precondition is always satisfied (last line) and, as a side effect, creates method type annotations for `belongs_to`.

In more detail, `hmi` is set to the first argument to `belongs_to`, and `options` is either `nil` or the hash argument, if present. (Here `hm` is shorthand for “has many,” i.e., since the Talk belongs to a User, the User has many Talks.) Then `hmu` is set to either the `class_name` argument, if present, or `hmi` after singularizing and camel-casing it. Then `type` is called twice, once to give a type to a getter method created by `belongs_to`, and once for a setter method (whose name ends with `=`). Notation `#{e}` inside a string evaluates the expression `e` and inserts the result in the string. In this particular case, these two calls to type evaluate to

```
type "owner", "()" → User"
type "owner=", "(User) → User"
```

Now consider executing this code. When `Talk` is loaded, `belongs_to` will be invoked, adding those type signatures to the class. Then when `owner?` is called, Hummingbird will perform type checking using currently available type information, and so it will be able to successfully type check the body. Moreover, notice this approach is very flexible. Rails does not require that `belongs_to` be used at the beginning of a class or even during this particular class definition. (In Ruby, it is possible to “re-open” a class later on and add more methods to it.) But no matter where the call occurs, it must be before `owner?` is called so that `owner` is defined. Thus in this case, Hummingbird’s typing strategy matches well with Ruby’s semantics.

Type Checking Dynamically Created Methods. In the previous example, we trusted Rails to dynamically generate code matching the given type signature. Figure 2 shows an example, extracted from *Rolify*, in which user code dynamically generates a method. The first part of the figure defines a module (aka mixin) with a two-argument method `define_dynamic_method`. The method body calls `define_method` to create a method named using the first argu-

```
1 module Rolify :: Dynamic
2   def define_dynamic_method(role_name, resource)
3     class_eval do
4       define_method("is_#{role_name}?" .to_sym) do
5         has_role?(" #{role_name}")
6       end if
7       !method_defined?("is_#{role_name}?" .to_sym)
8       ...
9     end end
10
11   pre :define_dynamic_method do |role_name, resource|
12     type "is_#{role_name}?", "()" → %bool"
13     true
14   end end
15
16   class User; include Rolify :: Dynamic end
17   user = User.first
18   user.define_dynamic_method("professor", ...)
19   user.define_dynamic_method("student", ...)
20   user.is_professor ?
21   user.is_student ?
```

Figure 2. Methods Dynamically Created by User Code.

ment, as long as that method does not exist (note the postfix `if` on line 7). Similarly to earlier, line 11 adds a precondition to `define_dynamic_method` that provides an appropriate method type. (We do not check for a previous type definition since adding the same type again is harmless.)

The code starting at line 16 uses the module. This particular code is not from our experiment but is merely for expository purposes. Here we (re)open class `User` and mix in the module. Then we create a user; call `define_dynamic_method` twice; and then call the generated methods `is_professor?` and `is_student?`.

In this case, since the generated methods have type annotations and are in user code, Hummingbird type checks their bodies when they are called, just like any other user-defined method with a type. For example, consider the call to `is_professor?`, which is given type `() → %bool`. At the call, Hummingbird type checks the code block at line 4 and determines that it has no arguments and that its body returns a boolean, i.e., it type checks.

User-provided Type Signatures. In the examples so far, the types for dynamically created methods could be determined automatically. However, consider Figure 3, which shows an excerpt from *CCT* that uses *Struct* from the Ruby core library. Line 1 creates a new class, instances of which are defined to have getters `type`, `account_name`, and `amount`, and setters `type=`, `account_name=`, and `amount=`. The `process_transactions` method iterates through instance field `@transactions` (whose type is provided on line 8), and calls the `account_name` method of each one.

From line 1 we know the `account_name` method exists, but we do not know its type. Indeed, a “struct field” can hold

```

1 Transaction=Struct.new(:type,:account_name,:amount)
2 class ApplicationRunner
3   def process_transactions
4     @transactions.each do |t|
5       name = t.account_name
6       ...
7     end ... end
8     field_type :@transactions, "Array<Transaction>"
9   end
10
11 class Struct
12   def self.add_types(*types)
13     members.zip(types).each {|name, t|
14       self.class_eval do
15         type name, "() → #{t}"
16         type "#{name}=", "(t) → #{t}"
17       end
18     }
19   end
20 end
21 Transaction.add_types("String", "String", "String")

```

Figure 3. Type Signatures for Struct.

any type by default. Thus, to fully type check the body of `process_transactions`, we need more information from the programmer to specify the type of `account_name`.

The bottom part of Figure 3 defines a new method, `add_types`, that the programmer can call to indicate desired struct field types. The types are given in the same order as the constructor arguments, and the body of `add_types` uses `zip` to pair up the constructor arguments (retrieved via `members`) and the types, and then iterates through the pairs, creating the appropriate type signatures for the getters and setters. The last line of the figure uses `add_types` to create type signatures for this example, allowing us to type check `process_transactions` when it is called.

In this particular case, we could have individually specified type signatures for the methods of `Transaction`. However, because Hummingbird lets programmers write arbitrary Ruby programs to generate types, we were able to develop this much more elegant solution.

3. Formalism

We formalize Hummingbird using the core, Ruby-like language shown at the top of Figure 4. *Values* v include `nil`, which can be treated as if it has any type, and $[A]$, which is an instance of class A . Note that we omit both fields and inheritance from our formalism for simplicity, but they are handled by our implementation.

Expressions e include values, variables x , the special variable `self`, assignments $x = e$, and sequencing $e; e$. Objects are created with $A.new$. Conditional `if  $e_1$  then  $e_2$  else  $e_3$` evaluates to e_2 unless e_1 evaluates to `nil`, in which case it

values	$v ::= \text{nil} \mid [A]$
expressions	$e ::= v \mid x \mid \text{self} \mid x = e \mid e; e \mid A.new$ $\mid \text{if } e \text{ then } e \text{ else } e \mid e.m(e)$ $\mid \text{def } A.m = b \mid \text{type } A.m : \tau_m$
premtys	$b ::= \lambda x.e$
val tys	$\tau ::= A \mid \text{nil}$
meth tys	$\tau_m ::= \tau \rightarrow \tau$

$x \in \text{var ids} \quad m \in \text{meth ids} \quad A \in \text{cls ids}$

dyn env	E	:	$\text{var ids} \rightarrow \text{vals}$
dyn cls tab	DT	:	$\text{cls ids} \rightarrow \text{meth ids} \rightarrow \text{premtys}$
contexts	C	$::=$	$\square \mid x = C \mid C.m(e) \mid v.m(C)$ $\mid C; e \mid \text{if } C \text{ then } e \text{ else } e$
stack	S	$::=$	$\cdot \mid (E, C) :: S$
type env	Γ, Δ	:	$\text{var ids} \rightarrow \text{val tys}$
type tab	TT	:	$\text{cls ids} \rightarrow \text{meth ids} \rightarrow \text{meth tys}$
cache	X	$::=$	$\text{cls ids} \rightarrow \text{meth ids} \rightarrow \mathcal{D}_M \times \mathcal{D}_\leq$
typ chk deriv	$\mathcal{D}_M$	$::=$	$TT \vdash \langle \Gamma, e \rangle \Rightarrow \langle \Gamma', \tau \rangle$
subtyp deriv	$\mathcal{D}_\leq$	$::=$	$\tau_1 \leq \tau_2$

Figure 4. Source Language and Auxiliary Definitions.

evaluates to e_3 . Method invocation $e_1.m(e_2)$ is standard, invoking the m method based on the run-time type of e_1 .

Expression `def  $A.m = \lambda x.e$` , defines method m of class A as taking argument x and returning e . (We refer to $\lambda x.e$ as a *premethod*.) This form allows methods to be defined anywhere during execution, thus it combines the features of Ruby’s `def` and `define_method`. As in Ruby, if $A.m$ is already defined, `def` overwrites the previous definition. The `def` expression itself evaluates to `nil`.

Finally, expression `type  $A.m : \tau \rightarrow \tau'$` asserts that method m of class A has domain type τ and range type τ' . Types may be either classes A or `nil`, the type of expression `nil`. The type expression overwrites the previous type of $A.m$, if any. Like Hummingbird, there is no ordering dependency between `def` and `type`—the only requirement is that a method’s type must be declared by the time the method is called. The type expression itself evaluates to `nil`.

Type Checking. Figure 5 gives the static type checking rules. As in Hummingbird, static type checking is performed at run time at method entry—thus these rules will be invoked as a subroutine by the dynamic semantics (below). The bottom part of Figure 4 defines the sets and maps used in this figure and in the dynamic semantics.

In these rules, TT is a *type table* mapping class and method ids $A.m$ to their corresponding types, as declared by `type`, and Γ is a *type environment* mapping local variables to their types. These rules prove judgments of the form $TT \vdash \langle \Gamma, e \rangle \Rightarrow \langle \Gamma', \tau \rangle$, meaning with type table TT , in type environment Γ , expression e has type τ , and after evaluating e , the new type environment is Γ' . Using an “output”

$TT \vdash \langle \Gamma, e \rangle \Rightarrow \langle \Gamma', \tau \rangle$			
(TNil)	(TObject)	(TSelf)	(TVar)
$TT \vdash \langle \Gamma, \text{nil} \rangle \Rightarrow \langle \Gamma, \text{nil} \rangle$	$TT \vdash \langle \Gamma, [A] \rangle \Rightarrow \langle \Gamma, A \rangle$	$TT \vdash \langle \Gamma, \text{self} \rangle \Rightarrow \langle \Gamma, \Gamma(\text{self}) \rangle$	$TT \vdash \langle \Gamma, x \rangle \Rightarrow \langle \Gamma, \Gamma(x) \rangle$
(TSeq)	(TAssn)	(TDef)	
$TT \vdash \langle \Gamma, e_1 \rangle \Rightarrow \langle \Gamma_1, \tau_1 \rangle$ $TT \vdash \langle \Gamma_1, e_2 \rangle \Rightarrow \langle \Gamma_2, \tau_2 \rangle$ $TT \vdash \langle \Gamma, e_1; e_2 \rangle \Rightarrow \langle \Gamma_2, \tau_2 \rangle$	$TT \vdash \langle \Gamma, e \rangle \Rightarrow \langle \Gamma', \tau \rangle$ $TT \vdash \langle \Gamma, x = e \rangle \Rightarrow \langle \Gamma'[x \mapsto \tau], \tau \rangle$	$TT \vdash \langle \Gamma, \text{def } A.m = \lambda x.e \rangle \Rightarrow \langle \Gamma, \text{nil} \rangle$	
	(TNew)	(TType)	
	$TT \vdash \langle \Gamma, A.\text{new} \rangle \Rightarrow \langle \Gamma, A \rangle$	$TT \vdash \langle \Gamma, \text{type } A.m : \tau_m \rangle \Rightarrow \langle \Gamma, \text{nil} \rangle$	
(TApp)	(TIf)		
$TT \vdash \langle \Gamma, e_0 \rangle \Rightarrow \langle \Gamma_0, A \rangle$ $TT \vdash \langle \Gamma_0, e_1 \rangle \Rightarrow \langle \Gamma_1, \tau \rangle$ $TT(A.m) = \tau_1 \rightarrow \tau_2 \quad \tau \leq \tau_1$ $TT \vdash \langle \Gamma, e_0.m(e_1) \rangle \Rightarrow \langle \Gamma_1, \tau_2 \rangle$	$TT \vdash \langle \Gamma, e_0 \rangle \Rightarrow \langle \Gamma', \tau \rangle$ $TT \vdash \langle \Gamma', e_1 \rangle \Rightarrow \langle \Gamma_1, \tau_1 \rangle$ $TT \vdash \langle \Gamma', e_2 \rangle \Rightarrow \langle \Gamma_2, \tau_2 \rangle$ $TT \vdash \langle \Gamma, \text{if } e_0 \text{ then } e_1 \text{ else } e_2 \rangle \Rightarrow \langle \Gamma_1 \sqcup \Gamma_2, \tau_1 \sqcup \tau_2 \rangle$		

Figure 5. Type Checking System.

type environment Γ' allows us to build a flow-sensitive type system, in which variables' types can change at assignments. Note there is no output TT because the type table does not change during static type checking—it only changes as the program is executed by the dynamic semantics.

The type rules are largely standard. (TNil) and (TObject) give nil and instances the obvious types. (TSelf) and (TVar) give self and local variables their types according to the type environment. Since none of these four expressions updates the state, the output type environment is the same as the input environment.

(TSeq) types sequencing, threading the type environment from the output of e_1 to the input of e_2 . (TAssn) types an assignment, updating the output type environment to bind the assigned variable x to the type of the right-hand side. (TNew) types object creation in the obvious way. (TDef) trivially type checks method definitions. Notice we do not type check the method body; that will happen at run time when the method is actually called. (TType) type checks a type expression, which has no effect during type checking. Such expressions are only evaluated at run-time, when they update the type table (see below).

One consequence of (TType) is that our type system forbids typing a method and then immediately calling it in the same method body. For example, the following method body would fail to type check:

```

1 def A.m =  $\lambda x$ .
2   def B.m = ...;   # define B.m
3   type B.m : ...;  # give B.m a type
4   B.new.m          # type error! B.m not in type table

```

Here we type check A.m's body at the first call to it, so the type expression has not been run—and hence has not bound

a type to B.m—yet. Thus it is a type error to invoke B.m in the method body.

While we could potentially solve this problem with a more complex type system, in our experience (Section 5) we have not needed such a feature.

Next, (TApp) types method invocation $e_0.m(e_1)$, where we look up the method's type in TT based on the compile-time type of e_0 . (Note that since there is no inheritance, we need not search the inheritance hierarchy to find the type of $A.m$.) Here subtyping is defined as $\text{nil} \leq A$ and $A \leq A$ for all A . Thus, as is standard in languages with nil, the type system may accept a program that invokes a non-existent method of nil even though this is a run-time error. However, notice that if e_0 evaluates to a non-nil value, then (TApp) guarantees e_0 has method m .

Finally, (TIf) types conditionals. Like Ruby, the guard e_0 may have any type. The type of the conditional is the least upper bound of the types of the two branches, defined as $A \sqcup A = A$ and $\text{nil} \sqcup \tau = \tau \sqcup \text{nil} = \tau$. The output environment of the conditional is the least upper bound of the output environments of the branches, defined as $(\Gamma_1 \sqcup \Gamma_2)(x) = \Gamma_1(x) \sqcup \Gamma_2(x)$ if $x \in \text{dom}(\Gamma_1) \wedge x \in \text{dom}(\Gamma_2)$ and $(\Gamma_1 \sqcup \Gamma_2)(x)$ is undefined otherwise.

Dynamic Semantics. Figure 6 gives a small-step dynamic semantics for our language. The semantics operates on *dynamic configurations* of the form $\langle X, TT, DT, E, e, S \rangle$. The first two components are the key novelties to support run-time static type checking. X is a *cache* mapping $A.m$ to the type checking proofs for its method body (more details below). TT is the *type table*, which is updated at run time by calls to type. The last four components are standard. DT is a *dynamic class table* mapping $A.m$ to its premethod. E is the *dynamic environment* mapping local variables to values. e is the expression being reduced. Lastly, S is a *stack*

$$\langle X, TT, DT, E, e, S \rangle \rightarrow \langle X', TT', DT', E', e', S' \rangle$$

(ESelf)	$\langle X, TT, DT, E, \text{self}, S \rangle \rightarrow \langle X, TT, DT, E, E(\text{self}), S \rangle$	
(EVar)	$\langle X, TT, DT, E, x, S \rangle \rightarrow \langle X, TT, DT, E, E(x), S \rangle$	
(EAssn)	$\langle X, TT, DT, E, x = v, S \rangle \rightarrow \langle X, TT, DT, E[x \mapsto v], v, S \rangle$	
(ENew)	$\langle X, TT, DT, E, A.\text{new}, S \rangle \rightarrow \langle X, TT, DT, E, [A], S \rangle$	
(ESeq)	$\langle X, TT, DT, E, (v; e_2), S \rangle \rightarrow \langle X, TT, DT, E, e_2, S \rangle$	
(ElfTrue)	$\langle X, TT, DT, E, \text{if } v \text{ then } e_1 \text{ else } e_2, S \rangle \rightarrow \langle X, TT, DT, E, e_1, S \rangle$	if $v \neq \text{nil}$
(ElfFalse)	$\langle X, TT, DT, E, \text{if nil then } e_1 \text{ else } e_2, S \rangle \rightarrow \langle X, TT, DT, E, e_2, S \rangle$	
(EDef)	$\langle X, TT, DT, E, \text{def } A.m = \lambda x.e, S \rangle \rightarrow \langle X \setminus A.m, TT, DT[A.m \mapsto \lambda x.e], E, \text{nil}, S \rangle$	
(EType)	$\langle X, TT, DT, E, \text{type } A.m : \tau_m, S \rangle \rightarrow \langle (X \setminus A.m)[TT'], TT', DT, E, \text{nil}, S \rangle$	$TT' = TT[A.m \mapsto \tau_m] \text{ and } A.m \notin \text{TApp}(S)$
(EAppMiss)	$\langle X, TT, DT, E, C[v_1.m(v_2)], S \rangle \rightarrow \langle X', TT, DT, [\text{self} \mapsto v_1, x \mapsto v_2], e, (E, C) :: S \rangle$	if $A.m \notin \text{dom}(X)$ and $v_1 = [A]$ and $DT(A.m) = \lambda x.e$ and $TT(A.m) = \tau_1 \rightarrow \tau_2$ and $\text{type_of}(v_2) \leq \tau_1$ and $\mathcal{D}_M = (TT \vdash \langle [x \mapsto \tau_1, \text{self} \mapsto A], e \rangle \Rightarrow \langle \Gamma', \tau \rangle)$ holds and $\mathcal{D}_\leq = (\tau \leq \tau_2)$ holds and $X' = X[A.m \mapsto (\mathcal{D}_M, \mathcal{D}_\leq)]$
(EAppHit)	$\langle X, TT, DT, E, C[v_1.m(v_2)], S \rangle \rightarrow \langle X, TT, DT, [\text{self} \mapsto v_1, x \mapsto v_2], e, (E, C) :: S \rangle$	if $A.m \in \text{dom}(X)$ and $v_1 = [A]$ and $DT(A.m) = \lambda x.e$ and $TT(A.m) = \tau_1 \rightarrow \tau_2$ and $\text{type_of}(v_2) \leq \tau_1$
(ERet)	$\langle X, TT, DT, E', v, (E, C) :: S \rangle \rightarrow \langle X, TT, DT, E, C[v], S \rangle$	
(EContext)	$\frac{\langle X, TT, DT, E, e, S \rangle \rightarrow \langle X', TT', DT', E', e', S' \rangle \quad \nexists v_1, v_2, e'. e = (v_1.m(v_2)) \vee e = v_1 \vee e = C[e']}{\langle X, TT, DT, E, C[e], S \rangle \rightarrow \langle X', TT', DT', E', C[e'], S' \rangle}$	

Figure 6. Dynamic Semantics.

of pairs (E, C) , where E is the dynamic environment and C is the *evaluation context* (defined in the usual way) at a call site. The semantics pushes onto the stack at calls and pops off the stack at returns.

The first seven rules in the semantics are standard. (ESelf) and (EVar) evaluate self and variables by looking them up in the environment. (EAssn) binds a variable to a value in the environment. Notice that, like Ruby, variables can be written without first declaring them, but it is an error to try to read a variable that has not been written. (ENew) creates a new instance. Note that since objects do not have fields, we do not need a separate heap. (ESeq) discards the left-hand side of a sequence if it has been fully evaluated. (ElfTrue) reduces to the true branch if the guard is non-nil, and (ElfFalse) reduces to the false branch otherwise.

The next four rules are the heart of just-in-time static type checking. Our goal is to statically type check methods once at the first call, and then avoid rechecking them unless something has changed. To formalize this notion, we define the cache X as a map from $A.m$ to a pair of typing derivations $(\mathcal{D}_M, \mathcal{D}_\leq)$. Here $\mathcal{D}_M$ is a type checking derivation from Figure 5 for the body of $A.m$, and $\mathcal{D}_\leq$ is a subtyping judgment showing that the type of e is a subtype of the declared return type. We need $\mathcal{D}_\leq$ because our type system is syntax-directed and hence does not include a standalone subsumption rule.

(EDef) reduces to nil, updating the dynamic class table to bind $A.m$ to the given premethod along the way. Recall that we allow a method to be redefined with def. Hence we need to *invalidate* anything in the cache relating to $A.m$ so

that $A.m$ will be checked the next time it is called. More precisely:

Definition 1 (Cache invalidation). *We write $X \setminus A.m$ to indicate a new cache that is the same as X , except $A.m$ has been invalidated, meaning:*

1. Any entries with $A.m$ as the key are removed.
2. Any entries with a $\mathcal{D}_M$ that apply (TApp) with $A.m$ are removed.

Thus, in (EDef), the output cache is the same as the input cache but with $A.m$ invalidated.

(EType) also reduces to nil, updating the type table to be TT' , which is the same as TT but with new type information for $A.m$. As with (EDef), we invalidate $A.m$ in the cache. However, there is another subtlety. Recall that cached typing derivations $\mathcal{D}_M$ include the type table TT . This is potentially problematic, because we are changing the type table to TT' . However, cache invalidation removes any derivations that refer to $A.m$. Hence, cached type derivations that use TT can safely use TT' . Formally, we define:

Definition 2 (Cache upgrading). *We write $X[TT']$ to indicate a new cache that is the same as X , except the type table in every derivation is replaced by TT' .*

Thus, in (EType), the output cache is upgraded to the new type table after invalidation.

The next two rules use the type cache. Both rules evaluate a method call in a context, written $C[v_1.m(v_2)]$; we will discuss the other rule for contexts shortly. In both rules, the receiver v_1 is a run-time object $[A]$. (EAppMiss) applies

when $A.m$ is not in the cache. In this case, we look up the type of $A.m$ in TT , yielding some type $\tau_1 \rightarrow \tau_2$. We type check the method body e in an environment in which formal variable x is bound to τ_1 and `self` is bound to A , yielding a derivation $\mathcal{D}_M$. We check that the resulting type τ of e is a subtype of the declared type e_2 , with subtyping derivation $\mathcal{D}_\leq$. Finally, we check that the run-time type of v_2 —defined as $\text{type_of}(\text{nil}) = \text{nil}$ and $\text{type_of}([A]) = A$ —is a subtype of τ_1 . If all this holds, then it is type-safe to call the method. Hence we update the cache with the method’s typing derivations and start evaluating the method body, pushing the context C and the environment E on the stack.

(EAppHit) is similar but far simpler. This rule applies when $A.m$ is in the cache. In this case we know its method body has been successfully type checked, so we need only check that the run-time type of v_2 is a subtype of the declared domain type of v_1 . If so, we allow the method call to proceed.

However a method is called, the return, handled by (ERet), is the same. This rule applies when an expression has been fully evaluated and is at the top level. In this case, we pop the stack, replacing E' with E from the stack and plugging the value v into the context C from the stack.

Finally, (EContext) takes a step in an subexpression inside a context C . This rule only applies if the subexpression is not a method call (since that case is handled by (EApp*), which must push the context on the stack) and not a fully evaluated value (which is handled by (ERet), which must pop the context from the stack). We also do not allow the subexpression to itself be a context, since that could cause (EApp*) and (ERet) to misbehave.

Soundness. Our type system forbids invoking non-existent methods of objects. However, there are three kinds of errors the type system does not prevent: invoking a method on `nil`; calling a method whose body does not type check at run time; and calling a method that has a type signature but is itself undefined. (We could prevent the latter error by adding a side condition to (TApp) that requires the method to be defined, but we opt not to to keep the formalism slightly simpler.) To state a soundness theorem, we need to account for these cases, which we do by extending the dynamic semantics with rules that reduce to *blame* in these three cases. After doing so, we can state soundness:

Theorem 1 (Soundness). *If $\emptyset \vdash \langle \emptyset, e \rangle \Rightarrow \langle \Gamma', \tau \rangle$ then either e reduces to a value, e reduces to blame, or e diverges.*

We show soundness using a standard progress and preservation approach. The key technical challenge is preservation, in which we need to show that not only are expression types preserved, but also the validity of the cache and types of contexts pushed on the stack. We omit the proof due to lack of space, but it can be found in a companion technical report [27].

4. Implementation

Hummingbird is implemented using a combination of Ruby and OCaml. On the OCaml side, we use the Ruby Intermediate Language (RIL) [13] to parse input Ruby files and translate them to control-flow graphs (CFG) on which we perform type checking. On the Ruby side, we extend RDL [26], a contract system for Ruby, to perform static type checking. We next discuss the major challenges of implementing Hummingbird.

RIL. RIL is essentially the front-end of Diamondback Ruby (DRuby) [8, 14]. Given an input Ruby program, RIL produces a CFG that simplifies away many of the tedious features of Ruby, e.g., multiple forms of conditionals. We modified DRuby so it emits the RIL CFG as a JSON file and then exits. When loading each application file at run-time, we read the corresponding JSON file and store a mapping from class and method names and positions (file and line number) to the JSON CFG. At run-time we look up CFGs in this map to perform static type checking.

RDL and Type Checking. Like standard RDL, Hummingbird’s type annotation stores type information in a map and wraps the associated method to intercept calls to it. We should emphasize that RDL does not perform any static checking on its own—rather, it solely enforces contracts dynamically. In Hummingbird, when a wrapped method is called, Hummingbird first checks to see if it has already been type checked. If not, Hummingbird retrieves the method’s CFG and type and then statically checks that the CFG matches the given type.

Hummingbird uses RDL’s type language, which includes nominal types, intersection types, union types, optional and variable length arguments, block (higher-order method) types, singleton types, structural types, a self type, generics, and types for heterogeneous arrays and hashes. Hummingbird supports all of these kinds of types except structural types, self types, heterogeneous collections, and some variable length arguments. In addition, Hummingbird adds support for both instance field types (as seen in Figure 3) and class field types.

There is one slight subtlety in handling union types: If in a method call the receiver has a union type, Hummingbird performs type checking once for each arm of the union and the unions the possible return types. For example if in call $e.m(\dots)$ the receiver has type $A \cup B$, then Hummingbird checks the call assuming e has type of $A.m$, yielding a return type τ_A ; checks the call assuming $B.m$, yielding return type τ_B ; and then sets the call’s return type to $\tau_A \cup \tau_B$.

Eliminating Dynamic Checks. Recall the (EApp*) rules dynamically check that a method’s actual arguments have the expected types before calling a statically typed method. This check ensures that an untrusted caller cannot violate the assumptions of a method. However, observe that if the immediate caller is itself statically checked, then we know

the arguments are type-safe. Thus, as a performance optimization, Hummingbird only dynamically checks arguments of statically typed methods if the caller is itself not statically checked. As a further optimization, Hummingbird also does not dynamically check calls from Ruby standard library methods or the Rails framework, which are assumed to be type-safe. The one exception is that Hummingbird does dynamically check types for the Rails params hash, since those values come from the user's browser and hence are untrusted.

Numeric Hierarchy. Ruby has a Numeric tower that provides several related types for numbers. For example, `Fixnum < Integer < Numeric` and `Bignum < Integer < Numeric`. Adding two `Fixnums` normally results in another `Fixnum`, but adding two large `Fixnums` could result in a `Bignum` in the case of numeric overflow. To keep the type checking system simple, Hummingbird omits the special overflow case and does not take `Bignum` into consideration. (This could be addressed by enriching the type system [33].) Numeric overflow does not occur in our experiments.

Code Blocks. As mentioned earlier, Ruby code blocks are anonymous functions delimited by `do...end`. Hummingbird allows methods that take code block arguments to be annotated with the block's type. For example:

```
type :m, "()" { (T) → U } → nil"
```

indicates that `m` takes no regular arguments; one code block argument where the block takes type `T` and returns type `U`; and `m` itself returns type `nil`.

There are two cases involving code blocks that we need to consider. First, suppose Hummingbird is statically type checking a call `m() do |x| body end`, and `m` has the type given just above. Then at the call, Hummingbird statically checks that the code block argument matches the expected type, i.e., assuming `x` has type `T`, then `body` must produce a value of type `U`. Second, when statically type checking `m` itself, Hummingbird should check that calls to the block are type correct. Currently this second case is unimplemented as it does not arise in our experiments.

Recall from above that Hummingbird sometimes needs to dynamically check the arguments to a statically typed method. While this test is easy to do for objects, it is hard to do for code blocks, which would require higher-order contracts [11]. Currently Hummingbird does not implement this higher order check, and simply assumes code block arguments are type safe. Also, Hummingbird currently assumes the `self` inside a code block is the same as in the enclosing method body. This assumption holds in our experiments, but it can be violated using `instance_eval` and `instance_exec` [34]. In the future, we plan to address this limitation by allowing the programmer to annotate the `self` type of code blocks.

Type Casts. While Hummingbird's type system is quite powerful, it cannot type check every Ruby program, and thus in some cases we need to insert type casts. Hummingbird includes a method `o.rdl_cast(t)` that casts `o`'s type to `t`. After such a call, Hummingbird assumes that `o` has type `t`. At run-time, the call dynamically checks that `o` has the given type.

In our experience, type casts have two main uses. First, sometimes program logic dictates that we can safely down-cast an object. For example, consider the following method from one of our experiments:

```
def self.load_cache
  f = datafile_path(['cache', 'countries'])
  t = Marshal.load(File.binread(f))
  @@cache ||= t.rdl_cast('Hash<String, %any>')
end
```

`Marshal.load` returns the result of converting its serialized data argument into a Ruby object of arbitrary type. However, in our example, the argument passed to `Marshal.load` is always an application data file that will be converted to the annotated `Hash`.

Second, by default Hummingbird gives instances of generic classes their “raw” type with no type parameters. To add parameters, we use type casts, as in the following code:

```
a = [] # a has type Array
a.rdl_cast("Array<Fixnum>") # cast to Array<Fixnum>
a.push(0) # ok
a.push("str") # type error due to cast
```

Here without the type annotation the last line would succeed; with the annotation it triggers a type error. Note that when casting an array or hash to a generic type, `rdl_cast` iterates through the elements to ensure they have the given type.

Modules. Ruby supports mixins via *modules*, which are collections of methods that can be added to a class. Recall that Hummingbird caches which methods have been statically type checked. Because a module can be mixed in to multiple different classes—and can actually have different types in those different classes—we need to be careful that module method type checks are cached via where they are mixed in rather than via the module name.

For example, consider the following code, where the method `foo` defined in module `M` calls `bar`, which may vary depending on where `M` is mixed in:

```
module M
  def foo(x) bar(x) end
end
class C; include M; def bar(x) x + 1 end
class D; include M; def bar(x) x.to_s end
```

Here method `foo` returns `Fixnum` when mixed into `C` and `String` when mixed into `D`. Thus, rather than track the type checking status of `M#foo`, Hummingbird separately tracks the statuses of `C#foo` and `D#foo`.

Cache Invalidation. Recall from Section 3 that Hummingbird needs to invalidate portions of the cache when certain typing assumptions change. While Hummingbird currently does not support cache invalidation in general, it does support one important case. In Rails *development mode*, Rails automatically reloads modified files without restarting, thus redefining the methods in those files but leaving other methods intact [24]. In Rails development mode, Hummingbird intercepts the Rails reloading process and performs appropriate cache invalidation. More specifically, when a method is called, if there is a difference between its new and old method body (which we check using the RIL CFGs), we invalidate the method and any methods that depend on it. We also maintain a list of methods defined in each class, and when a class is reloaded we invalidate dependencies of any method that has been removed. In the next section, we report on an experiment running a Rails app under Hummingbird as it is updated.

We plan to add more general support for cache invalidation in future work. There are two main cases to consider. The first is when a method is redefined or is removed (which never happens in our experiments except in Rails development mode). Ruby provides two methods, `method_added` and `method_removed`, that can be used to register callbacks when the corresponding actions occur, which could be used for cache invalidation.

The second case of cache invalidation is method’s type changes. However, in RDL and Hummingbird, multiple calls to `type` for the same method are used to create intersection types. For example, the core library `Array#[]` method is given its type with the following code:

```
type Array, :[], '(Fixnum or Float) → t'
type Array, :[], '(Fixnum, Fixnum) → Array<t>'
type Array, :[], '(Range<Fixnum>) → Array<t>'
```

meaning if given a `Fixnum` or `Float`, method `Array#[]` returns the array contents type; and, if given a pair of `Fixnums` or a `Range<Fixnum>`, it returns an array.

In this setting, we cannot easily distinguish adding a new arm of an intersection type from replacing a method type. Moreover, adding a new arm to an intersection type should not invalidate the cache, since the other arms are still in effect. Thus, full support of cache invalidation will likely require an explicit mechanism for replacing earlier type definitions.

5. Experiments

We evaluated Hummingbird by applying it to six Ruby apps:

- *Talks*³ is a Rails app, written by the second author, for publicizing talk announcements. *Talks* has been in use in the UMD CS department since February 2012.

³<https://github.com/jeffrey-s-foster/talks>

- *Boxroom*⁴ is a Rails implementation of a simple file sharing interface.
- *Pubs* is a Rails app, developed several years ago by the second author, for managing lists of publications.
- *Rolify*⁵ is a role management library for Rails. For this evaluation, we integrated *Rolify* with *Talks* on the User resource.
- Credit Card Transactions (*CCT*)⁶ is a library that performs simple credit card processing tasks.
- *Countries*⁷ is an app that provides useful data about each country.

We selected these apps for variety rather than for being representative. We chose these apps because their source code is publicly available (except *Pubs*); they work with the latest versions of Ruby and RDL; and they do not rely heavily on other packages. Moreover, the first three apps use Rails, which is an industrial strength web app framework that is widely deployed; the next two use various metaprogramming styles in different ways than Rails; and the last one does not use metaprogramming, as a baseline.

Table 1 summarizes the results of applying Hummingbird to these apps. On the left we list the app name, version number or date (if no version number is available), and lines of code as measured with `sloccount` [41]. For the Rails apps, we ran `sloccount` on all ruby files in the model, controller, helper, and mailer directories. We do not include lines of code for views, as we do not type check views. For *Countries* and *CCT*, we ran `sloccount` on all files in the `lib` directory. For *Rolify*, we only statically checked several methods, comprising 84 lines of code, that use `define_method` in an interesting way.

Type Annotations. For all apps, we used common type annotations from RDL for the Ruby core and standard libraries. For several apps, we also added type annotations for third-party libraries and for the Rails framework. We trusted the annotations for all these libraries, i.e., we did not statically type check the library methods’ bodies.

We also added code to dynamically generate types for metaprogramming code. For Rails, we added code to dynamically generate types for model getters and setters based on the database schema; for finder methods such as `find_by_name` and `find_all_by_password` (the method name indicates which field is being searched); and for Rails associations such as `belongs_to`.

In Figure 2, we showed code we added to *Rolify* to generate types for a method created by calling `define_dynamic_method`. Calling `define_dynamic_method` also dynam-

⁴<http://boxroomapp.com>

⁵<https://github.com/RolifyCommunity/rolify>

⁶https://github.com/daino3/credit_card_transactions

⁷<https://github.com/hexorx/countries>

App	LoC	Static types			Dynamic types		Casts	Phs	Running time (s)			
		Chk'd	App	All	Gen'd	Used			Orig	No\$	Hum	Or. Ratio
<i>Talks-1/4/2013</i>	1,055	111	201	363	990	45	31	1	162	1,590	256	1.6×
<i>Boxroom-1.7.1</i>	854	127	221	306	534	93	17	1	263	705	327	1.2×
<i>Pubs-1/12/2015</i>	620	47	86	171	445	33	13	1	72.0	4,470	217	3.0×
<i>Rolify-4.0.0</i>	84	14	24	71	26	2	15	12	5.63	7.79	6.71	1.2×
<i>CCT-3/23/2014</i>	172	23	27	75	6	3	6	1	3.06	78.2	17.4	5.7×
<i>Countries-1.1.0</i>	227	33	40	111	0	0	22	1	1.02	18.1	4.62	4.5×

Table 1. Type checking results.

ically creates another method, `is_{role_name}_of(arg)?`, which we also provide types for in the pre block.

In *CCT*, we used the code in Figure 3 to generate types for Struct getters and setters.

Finally, we wrote type annotations for the app’s own methods that were included in the lines of code count in Table 1. We marked those methods to indicate Hummingbird should statically type check their bodies. Developing these annotations was fairly straightforward, especially since we could quickly detect incorrect annotations by running Hummingbird.

Type Checking Results. For each program, we performed type checking while running unit tests that exercised all the type-annotated app methods. For *Talks* and *Pubs*, we wrote unit tests with the goal of covering all application methods. For *Boxroom*, we used its unit tests on models but wrote our own unit tests on controllers, since it did not have controller tests. For *Rolify*, we wrote a small set of unit tests for the dynamic method definition feature. For *CCT* and *Countries*, we used the unit tests that came with those apps.

In all cases, the app methods type check correctly in Hummingbird; there were no type errors. The middle group of columns summarizes more detailed type checking data.

The “Static types” columns report data on static type annotations. The count under “Chk’d” is the number of type annotations for the app’s methods whose bodies we statically type checked. The count under “App” is that number plus the number of types for app-specific methods with (trusted) static type annotations, e.g., some Rails helper functions have types that we do not currently dynamically generate. The count under “All” reports the total number of static type annotations we used in type checking each app. This includes the “App” count plus standard, core, and third-party library type annotations for methods referred to in the app.

The “Dynamic types” columns report the number of types that were dynamically generated (“Gen’d”) and the number of those that were actually used during type checking (“Used”). These numbers differ because we tried to make the dynamic type information general rather than app-specific, e.g., we generate both the getter and setter for `belongs_to` even if only one is used by the app.

These results show that having types for methods generated by metaprogramming is critical for successfully typing these programs—every app except *Countries* requires at least a few, and sometimes many, such types.

The “Casts” column reports the number of type casts we needed to make these programs type check; essentially this measures how often Hummingbird’s type system is overly conservative. The results show we needed a non-trivial but relatively small number of casts. All casts were for the reasons discussed in Section 4: downcasting and generics.

The “Phs” column in Table 1 shows the number of type checking phases under Hummingbird. Here a *phase* is defined as a sequence of type annotation calls with no intervening static type checks, followed by a sequence of static type checks with no intervening annotations. We can see that almost all apps have only a single phase, where the type annotations are executed before any static type checks. Investigating further, we found this is due to the way we added annotations. For example, we set up our Rails apps so the first loaded application file in turn loads all type annotation files. In practice the type annotations would likely be spread throughout the app’s files, thus increasing the number of phases.

Rolify is the only application with multiple phases. Most of the phases come from calling `define_dynamic_method`, which dynamically defines other methods and adds their type annotations. The other phases come from the order in which the type annotation files are required—unlike the Rails apps, the *Rolify* type annotation files are loaded piecemeal as the application loads.

Performance. The last four columns of Table 1 report the overhead of using Hummingbird. The “Orig” column shows the running time without Hummingbird. The next two columns report the running time with Hummingbird, with caching disabled (“No\$”) and enabled (“Hum”). The last column lists the ratio of the “Hum” and “Orig” column.

For *Talks*, *Boxroom*, and *Pubs*, we measured the running time of a client script that uses curl to connect to the web server and exercise a wide range of functionality. For *CCT*, we measured the time for running its unit tests 100 times. For *Countries* and *Rolify*, we measured the time for running the unit tests once (since these take much more time than

CCT's tests). For all apps, we performed each measurement three times and took the arithmetic mean.

These results show that for the Rails apps, where IO is significant, Hummingbird slows down performance from 24% to 201% (with caching enabled). We think these are reasonable results for an early prototype that we have not spent much effort optimizing. Moreover, across all apps, the ratios are significantly better than prior systems that mix static and dynamic typing for Ruby [18, 28], which report orders of magnitude slowdowns.

Investigating further, we found that the main Hummingbird overhead arises from intercepting method calls to statically type checked methods. (Note the interception happens regardless of the cache state.) The higher slowdowns for *CCT* and *Countries* occur because those applications spend much of their time in code with intercepted calls, while the other applications spend most of their time in framework code, whose methods are not intercepted. We expect performance can be improved with further engineering effort.

We can also see from the results that caching is an important performance optimization: without caching, performance slows down $1.4\times$ to $62\times$. We investigated *pubs*, the app with the highest no-caching slowdown, and found that while running the application with large array inputs, certain application methods are called more than 13,000 times while iterating through the large arrays. This means that each of these application methods are statically type checked more than 13,000 times when caching is disabled.

Type Errors in Talks. We downloaded many earlier versions of *Talks* from its github repository and ran Hummingbird on them using mostly the same type annotations as for the latest version, changed as necessary due to program changes. Cumulatively, we found six type errors that were introduced and later removed as *Talks* evolved. Below the number after the date indicates which checkin it was, with 1 for the first checkin of the day, 2 for the second, etc.

- 1/8/12-4: This version misspells `compute_edit_fields` as `copute_edit_fields`. Hummingbird reported this error because the latter was an unbound local variable and was also not a valid method.
- 1/7/12-5: Instead of calling `@list.talks.upcoming.sort{ |a, b| ... }`, this version calls `@list.talks.upcoming{ |a, b| ... }` (leaving off the `sort`). Hummingbird detects this error because `upcoming`'s type indicates it does not take a block. Interestingly, this error would not be detected at run-time by Ruby, which simply ignores unused block arguments.
- 1/26/12-3: This version calls `user.subscribed_talks(true)`, but `subscribed_talks`'s argument is a `Symbol`.
- 1/28/12: This version calls `@job.handler.object`, but `@job.handler` returns a `String`, which does not have an object method.

- 2/6/12-2: This version uses undefined variable `old_talk`. Thus, Hummingbird assumes `old_talk` is a no-argument method and attempts to look up its type, which does not exist.
- 2/6/12-3: This version uses undefined variable `new_talk`

We should emphasize that although we expected there would be type errors in *Talks*, we did not know exactly what they were or what versions they were in. While the second author did write *Talks*, the errors were made a long time ago, and the second author rediscovered them independently by running Hummingbird.

Updates to Talks Finally, we performed an experiment in which we launched one version of *Talks* in Rails development mode and then updated the code to the next six consecutive versions of the app. (We skipped versions in which none of the Ruby application files changed) Notice that cache invalidation is particularly useful here, since in typical usage only a small number of methods are changed by each update.

In more detail, after launching the initial version of the app, we repeated the following sequence six times: Reset the database (so that we run all versions with the same initial data); run a sequence of `curl` commands that access the same *Talks* functionalities as the ones used to measure the running time of *Talks* in Table 1; update the code to the next version; and repeat.

Table 2 shows the results of our experiment. The “ Δ Meth” column lists the number of methods whose bodies or types were changed compared to the previous version. Note there are no removed methods in any of these versions. The “Added” column lists the number of methods added; such methods will be checked when they are called for the first time but do not cause any cache invalidations. The “Deps” column counts the number of *dependent* methods that call one or more of the changed methods. These methods plus the changed methods are those whose previous static type check are invalidated by the update. The last column, “Chk’d,” reports how many methods are newly or re-type checked after the update. Currently, Hummingbird always rechecks Rails helper methods, due to a quirk in the Rails implementation—the helper methods’ classes get a new name each time the helper file is reloaded, causing Hummingbird to treat their methods as new. Thus (except for the first line, since this issue does not arise on the first run), we list two numbers in the column: the first with all rechecks, including the helper methods, and the second excluding the helper methods.

These results show that in almost all cases, the second number in “Chk’d” is equal to the sum of the three previous columns. There is one exception: in 8/24/12-1, there 14 rechecked methods but 18 changed/added/dependent methods. We investigated and found that the 14 rechecks are composed of six changed methods that are rechecked once; two changed methods that are rechecked twice because they have

Version	Δ Meth	Added	Deps	Chk'd
5/14/12	N/A	N/A	N/A	77
7/24/12	1	-	4	15 / 5
8/24/12-1	8	2	8	24 / 14
8/24/12-2	-	1	-	11 / 1
8/24/12-3	1	1	-	12 / 2
9/14/12	1	-	-	15 / 1
1/4/13	4	-	-	13 / 4

Table 2. Talks Update Results

dependencies whose updates are interleaved with calls to those methods; one added method that is checked; and three dependent methods that are rechecked. The remaining added method is not called by the curl scripts, and the remaining dependent methods are also changed methods (this is the only version where there is overlap between the changed and dependent columns).

Finally, as there are no type errors in this sequence of updates, we confirmed that this streak of updates type checks under Hummingbird.

6. Related Work

There are several threads of related work.

Type Systems for Ruby. We have developed several prior type systems for Ruby. Diamondback Ruby (DRuby) [14] is the first comprehensive type system for Ruby that we are aware of. Because Hummingbird checks types at run-time, we opted to implement our own type checker rather than reuse DRuby for type checking, which would have required some awkward shuffling of the type table between Ruby and OCaml. Another reason to reimplement type checking was to keep the type system a little easier to understand—DRuby performs type inference, which is quite complex for this type language, in contrast to Hummingbird, which implements much simpler type checking.

DRuby was effective but did not handle highly dynamic language constructs well. *PRuby* [12] solves this problem using profile-based type inference. To use *PRuby*, the developer runs the program once to record dynamic behavior, e.g., what methods are invoked via `send`, what strings are passed to `eval`, etc. *PRuby* then applies DRuby to the original program text plus the profiled strings, e.g., any string that was passed to `eval` is parsed and analyzed like any other code. While *PRuby* can be effective, we think that Hummingbird’s approach is ultimately more practical because Hummingbird does not require a separate, potentially cumbersome, profiling phase. We note that Hummingbird does not currently handle `eval`, because it was not used in our subject apps’ code, but it could be supported in a straightforward way.

We also developed DRails [17], which type checks Rails apps by applying DRuby to translated Rails code. For example, if DRails sees a call to `belongs_to`, it outputs Ruby

code that explicitly contains the methods generated from the call, which DRuby can then analyze. While DRails was applied to a range of programs, its analysis is quite brittle. Supporting each additional Rails feature in DRails requires implementing, in OCaml, a source-to-source transformation that mimics that feature. This is a huge effort and is hard to sustain as Rails evolves. In contrast, Hummingbird types are generated in Ruby, which is far easier. DRails is also complex to use: The program is combined into one file, then run to gather profile information, then transformed and type checked. Using Hummingbird is far simpler. Finally, DRails is Rails-specific, whereas Hummingbird applies readily to other Ruby frameworks. Due to all these issues, we feel Hummingbird is much more lightweight, agile, scalable, and maintainable than DRails.

Finally, RubyDust [18] implements type inference for Ruby at run time. RubyDust works by wrapping objects to annotate them with type variables. More precisely, consider a method `def m(x) ... end`, and let α be the type variable for x . RubyDust’s wrapping is approximately equal to adding `x = Wrap.new(x,  $\alpha$ )` to the beginning of `m`. Uses of the wrapped `x` generate type constraints on α and then delegate to the underlying object. The Ruby Type Checker [28] (*rtc*) is similar but implements type checking rather than type inference.

Hummingbird has several important advantages over RubyDust and *rtc*. First, RubyDust and *rtc* can only report errors on program paths they observe. In contrast, Hummingbird type checks all paths through methods it analyzes. Second, wrapping every object with a type annotation is extremely expensive. By doing static analysis, Hummingbird avoids this overhead. Finally, RubyDust and *rtc* have no special support for metaprogramming. In RubyDust, dynamically created methods could have their types inferred in a standard way, though RubyDust would likely not infer useful types for Rails-created methods. In *rtc*, dynamically created methods would lack types, so their use would not be checked. (Note that it would certainly be possible to add Hummingbird-style support for metaprogramming-generated type annotations to either RubyDust or *rtc*.) In sum, we think that Hummingbird strikes the right compromise between the purely static DRuby approach and the purely dynamic RubyDust/*rtc* approach.

Type Systems for Other Dynamic Languages. Many researchers have proposed type systems for dynamic languages, including Python [3], JavaScript [2, 21, 37], Racket [33, 38, 39], and Lua [22], or developed new dynamic languages or dialects with special type systems, such as Thorn [5], TypeScript [4, 25], and Dart [7]. To our knowledge, these type systems are focused on checking the core language and can have difficulty in the face of metaprogramming.

One exception is RPython [1], which introduces a notion of load time, during which highly dynamic features may

be used, and run time, when they may not be. In contrast, Hummingbird does not need such a separation.

Lerner et al [20] propose a system for type checking programs that use JQuery, a very sophisticated Javascript framework. The proposed type system has special support for JQuery's abstractions, making it quite effective in that domain. On the other hand, it does not easily apply to other frameworks.

Feldthaus et al's TSCHECK [10] is a tool to check the correctness of TypeScript interfaces for JavaScript libraries. TSCHECK discovers a library's API by taking a snapshot after executing the library's top-level code. It then performs checking using a separate static analysis. This is similar to Hummingbird's tracking of type information at run-time and then performing static checking based on it. However, Hummingbird allows type information to be generated at any time and not just in top-level code.

Related Uses of Caching. Several researchers have proposed systems that use caching in a way related to Hummingbird. Koukoutos et al [19] reduce the overhead of checking data structure contracts (e.g., "this is a binary search tree") at run time by modifying nodes to hold key verification properties. This essentially caches those properties. However, because the properties are complex, the process of caching them is not automated.

Stulova et al [35] propose memoizing run-time assertion checking to improve performance. This is similar to Hummingbird's type check caching, but much more sophisticated because the cached assertions arise from a rich logic.

Hermenegildo et al [16] proposed a method to incrementally update analysis results at run-time as code is added, deleted, or changed. Their analysis algorithms are designed for constraint logic programming languages, and are much more complicated than Hummingbird's type checking.

Staged Analysis. MetaOCaml [36] is a multi-stage extension of OCaml in which code is compiled in one stage and executed in a later stage. The MetaOCaml compiler performs static type checking on any such delayed code, which is similar to Hummingbird's just-in-time type checking. A key difference between MetaOCaml and Hummingbird is that Ruby programs do not have clearly delineated stages.

Chugh et al's staged program analysis [6] performs static analysis on as much code as is possible at compile time, and then computes a set of remaining checks to be performed at run time. Hummingbird uses a related idea in which no static analysis is performed at compile time, but type checking is always done when methods are called. Hummingbird is simpler because it need not compute which checks are necessary, as it always does the same kind of checking.

Other. Several researchers have explored other ways to bring the benefits of static typing to dynamic languages. Contracts [11] check assertions at function or method entry and exit. In contrast, Hummingbird performs static analysis

of method bodies, which can find bugs on paths before they are run. At the same time, contracts can encode richer properties than types.

Gradual typing [32] lets developers add types gradually as programs evolve; Vitousek et al recently implemented gradual typing for Python [40]. Like types [42] bring some of the flexibility of dynamic typing to statically typed languages. The goal of these systems is to allow mixing of typed and untyped code. This is orthogonal to Hummingbird, which focuses on checking code with type annotations.

Richards et al [29, 30] have explored how highly dynamic language features are used in JavaScript. They find such features, including eval, are used extensively in a wide variety of ways, including supporting metaprogramming.

The GHC Haskell compiler lets developers defer type errors until run-time to suppress type errors on code that is never actually executed [15]. Hummingbird provides related behavior in that a method that is never called will never be type checked by Hummingbird. Template Haskell [31] can be used for compile-time metaprogramming. Since Haskell programs contain types, template Haskell is often used to generate type annotations, analogously to the type annotations generated using Hummingbird. Similarly, F# type providers [9] allow users to create compile time types, properties and methods. A key difference between these Haskell/F# features and Hummingbird is that Ruby does not have a separate compile time.

7. Conclusion

We presented Hummingbird, a novel tool that type checks Ruby apps using an approach we call just-in-time static type checking. Hummingbird works by tracking type information *dynamically*, but then checking method bodies *statically* at run time as each method is called. As long as any metaprogramming code is extended to generate types as it creates methods, Hummingbird will, in a very natural way, be able to check code that uses the generated methods. Furthermore, Hummingbird can cache type checking so it need not be unnecessarily repeated at later calls to the same method.

We formalized Hummingbird using a core, Ruby-like language that allows methods and their types to be defined at arbitrary (and arbitrarily separate) points during execution, and we proved type soundness. We implemented Hummingbird on top of RIL, for parsing Ruby source code, and RDL, for intercepting method calls and storing type information. We applied Hummingbird to six Ruby apps, some of which use Rails. We found that Hummingbird's approach is effective, allowing it to successfully type check all the apps even in the presence of metaprogramming. We ran Hummingbird on earlier versions of one app and found several type errors. Furthermore, we ran Hummingbird while applying a sequence of updates to a Rails app in development mode to demonstrate cache invalidation under Hummingbird. Fi-

nally, we measured Hummingbird’s run-time overhead and found it is reasonable.

In sum, we think that Hummingbird takes a strong step forward in bringing static typing to dynamic languages.

Acknowledgments

Thanks to ThanhVu Nguyen and the anonymous reviewers for their helpful comments. This research was supported in part by NSF CCF-1319666 and Subcontract to Northeastern University, NSF CCF-1518844.

References

- [1] D. Ancona, M. Ancona, A. Cuni, and N. D. Matsakis. RPython: A Step Towards Reconciling Dynamically and Statically Typed OO Languages. In *Proceedings of the 2007 Symposium on Dynamic Languages, DLS ’07*, pages 53–64, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-868-8. doi: 10.1145/1297081.1297091.
- [2] C. Anderson, P. Giannini, and S. Drossopoulou. Towards Type Inference for Javascript. In *Proceedings of the 19th European Conference on Object-Oriented Programming, ECOOP’05*, pages 428–452, Berlin, Heidelberg, 2005. Springer-Verlag. ISBN 3-540-27992-X, 978-3-540-27992-1. doi: 10.1007/11531142_19.
- [3] J. Aycock. Aggressive Type Inference. In *International Python Conference*, 2000.
- [4] G. Bierman, M. Abadi, and M. Torgersen. Understanding typescript. In *ECOOP 2014—Object-Oriented Programming*, pages 257–281. Springer, 2014.
- [5] B. Bloom, J. Field, N. Nystrom, J. Östlund, G. Richards, R. Strniša, J. Vitek, and T. Wrigstad. Thorn: Robust, Concurrent, Extensible Scripting on the JVM. In *Proceedings of the 24th ACM SIGPLAN Conference on Object Oriented Programming Systems Languages and Applications, OOPSLA ’09*, pages 117–136, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-766-0. doi: 10.1145/1640089.1640098.
- [6] R. Chugh, J. A. Meister, R. Jhala, and S. Lerner. Staged Information Flow for Javascript. In *Proceedings of the 30th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI ’09*, pages 50–62, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-392-1. doi: 10.1145/1542476.1542483.
- [7] Dart. Dart, 2015. <https://www.dartlang.org>.
- [8] DRuby. Diamondback Ruby, 2009. <http://www.cs.umd.edu/projects/PL/druby/>.
- [9] F#. Type Provider, 2016. <https://msdn.microsoft.com/en-us/library/hh156509.aspx>.
- [10] A. Feldthaus and A. Möller. Checking Correctness of TypeScript Interfaces for JavaScript Libraries. In *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA ’14*, pages 1–16, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2585-1. doi: 10.1145/2660193.2660215.
- [11] R. B. Findler and M. Felleisen. Contracts for Higher-order Functions. In *Proceedings of the Seventh ACM SIGPLAN International Conference on Functional Programming, ICFP ’02*, pages 48–59, New York, NY, USA, 2002. ACM. ISBN 1-58113-487-8. doi: 10.1145/581478.581484.
- [12] M. Furr, J. hoon (David) An, and J. S. Foster. Profile-Guided Static Typing for Dynamic Scripting Languages. In *ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA)*, pages 283–300, Orlando, Florida, October 2009. doi: 10.1145/1639949.1640110. Best student paper award.
- [13] M. Furr, J. hoon (David) An, J. S. Foster, and M. Hicks. The Ruby Intermediate Language. In *Dynamic Languages Symposium (DLS)*, pages 89–98, Orlando, Florida, October 2009. doi: 10.1145/1837513.1640148.
- [14] M. Furr, J. hoon (David) An, J. S. Foster, and M. Hicks. Static Type Inference for Ruby. In *Object-Oriented Program Languages and Systems (OOPS) Track at ACM Symposium on Applied Computing (SAC)*, pages 1859–1866, Honolulu, Hawaii, March 2009. doi: 10.1145/1529282.1529700.
- [15] GHCLanguageFeatures. Deferring Type Errors to Runtime, 2016. https://downloads.haskell.org/~ghc/latest/docs/html/users_guide/defer-type-errors.html.
- [16] M. Hermenegildo, G. Puebla, K. Marriott, and P. J. Stuckey. Incremental Analysis of Constraint Logic Programs. *ACM Trans. Program. Lang. Syst.*, 22(2):187–223, Mar. 2000. ISSN 0164-0925. doi: 10.1145/349214.349216.
- [17] J. hoon (David) An, A. Chaudhuri, and J. S. Foster. Static Typing for Ruby on Rails. In *IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 590–594, Auckland, New Zealand, November 2009. Short paper.
- [18] J. hoon (David) An, A. Chaudhuri, J. S. Foster, and M. Hicks. Dynamic Inference of Static Types for Ruby. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 459–472, Austin, TX, USA, January 2011. doi: 10.1145/1926385.1926437.
- [19] E. Koukoutos and V. Kuncak. Checking Data Structure Properties Orders of Magnitude Faster. In *Runtime Verification - 5th International Conference, RV 2014, Toronto, ON, Canada, September 22-25, 2014. Proceedings*, pages 263–268, 2014. doi: 10.1007/978-3-319-11164-3_22.
- [20] B. S. Lerner, L. Elbert, J. Li, and S. Krishnamurthi. Combining Form and Function: Static Types for JQuery Programs. In *Proceedings of the 27th European Conference on Object-Oriented Programming, ECOOP’13*, pages 79–103, Berlin, Heidelberg, 2013. Springer-Verlag. ISBN 978-3-642-39037-1. doi: 10.1007/978-3-642-39038-8_4.
- [21] B. S. Lerner, J. G. Politz, A. Guha, and S. Krishnamurthi. TeJaS: Retrofitting Type Systems for JavaScript. In *Proceedings of the 9th Symposium on Dynamic Languages, DLS ’13*, pages 1–16, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-2433-5. doi: 10.1145/2508168.2508170.
- [22] A. M. Maidl, F. Mascarenhas, and R. Ierusalimsky. Typed Lua: An Optional Type System for Lua. In *Proceedings of the Workshop on Dynamic Languages and Applications, Dyla’14*, pages 3:1–3:10, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2916-3. doi: 10.1145/2617548.2617553.

- [23] RailsGuides. Active Record Associations, 2015. http://guides.rubyonrails.org/association_basics.html.
- [24] RailsGuides. Autoloading and Reloading Constants, 2016. http://guides.rubyonrails.org/autoloading_and_reloading_constants.html.
- [25] A. Rastogi, N. Swamy, C. Fournet, G. Bierman, and P. Vekris. Safe & Efficient Gradual Typing for TypeScript. In *Proceedings of the 42Nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '15, pages 167–180, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3300-9. doi: 10.1145/2676726.2676971.
- [26] RDL. RDL, 2015. <https://github.com/plum-umd/rdl>.
- [27] B. Ren and J. S. Foster. Just-in-Time Static Type Checking for Dynamic Languages, 2016. preprint, <http://arxiv.org/abs/1604.03641>.
- [28] B. M. Ren, J. Toman, T. S. Strickland, and J. S. Foster. The Ruby Type Checker. In *Object-Oriented Program Languages and Systems (OOPS) Track at ACM Symposium on Applied Computing*, pages 1565–1572, Coimbra, Portugal, March 2013. doi: 10.1145/2480362.2480655.
- [29] G. Richards, S. Lebesne, B. Burg, and J. Vitek. An Analysis of the Dynamic Behavior of JavaScript Programs. In *Proceedings of the 31st ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '10, pages 1–12, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0019-3. doi: 10.1145/1806596.1806598.
- [30] G. Richards, C. Hammer, B. Burg, and J. Vitek. The Eval That Men Do: A Large-scale Study of the Use of Eval in Javascript Applications. In *Proceedings of the 25th European Conference on Object-oriented Programming*, ECOOP'11, pages 52–78, Berlin, Heidelberg, 2011. Springer-Verlag. ISBN 978-3-642-22654-0.
- [31] T. Sheard and S. P. Jones. Template Meta-programming for Haskell. *SIGPLAN Not.*, 37(12):60–75, Dec. 2002. ISSN 0362-1340. doi: 10.1145/636517.636528.
- [32] J. G. Siek and W. Taha. Gradual Typing for Functional Languages. In *Scheme and Functional Programming Workshop*, 2006.
- [33] V. St-Amour, S. Tobin-Hochstadt, M. Flatt, and M. Felleisen. Typing the Numeric Tower. In *Proceedings of the 14th International Conference on Practical Aspects of Declarative Languages*, PADL'12, pages 289–303, Berlin, Heidelberg, 2012. Springer-Verlag. ISBN 978-3-642-27693-4. doi: 10.1007/978-3-642-27694-1_21.
- [34] T. S. Strickland, B. Ren, and J. S. Foster. Contracts for Domain-Specific Languages in Ruby. In *Dynamic Languages Symposium (DLS)*, Portland, OR, October 2014. doi: 10.1145/2661088.2661092.
- [35] N. Stulova, J. F. Morales, and M. V. Hermenegildo. Practical Run-time Checking via Unobtrusive Property Caching. *CoRR*, abs/1507.05986, 2015.
- [36] W. Taha. MetaOcaml, 2016. <http://www.cs.rice.edu/~taha/MetaOCaml>.
- [37] P. Thiemann. Towards a Type System for Analyzing Javascript Programs. In *Proceedings of the 14th European Conference on Programming Languages and Systems*, ESOP'05, pages 408–422, Berlin, Heidelberg, 2005. Springer-Verlag. ISBN 3-540-25435-8, 978-3-540-25435-5. doi: 10.1007/978-3-540-31987-0_28.
- [38] S. Tobin-Hochstadt and M. Felleisen. The Design and Implementation of Typed Scheme. In *Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '08, pages 395–406, New York, NY, USA, 2008. ACM. ISBN 978-1-59593-689-9. doi: 10.1145/1328438.1328486.
- [39] S. Tobin-Hochstadt and M. Felleisen. Logical Types for Untyped Languages. In *Proceedings of the 15th ACM SIGPLAN International Conference on Functional Programming*, ICFP '10, pages 117–128, New York, NY, USA, 2010. ACM. ISBN 978-1-60558-794-3. doi: 10.1145/1863543.1863561.
- [40] M. M. Vitousek, A. M. Kent, J. G. Siek, and J. Baker. Design and Evaluation of Gradual Typing for Python. In *Proceedings of the 10th ACM Symposium on Dynamic Languages*, DLS '14, pages 45–56, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-3211-8. doi: 10.1145/2661088.2661101.
- [41] D. A. Wheeler. SLOccount, 2015. <http://www.dwheeler.com/sloccount>.
- [42] T. Wrigstad, F. Z. Nardelli, S. Lebesne, J. Östlund, and J. Vitek. Integrating Typed and Untyped Code in a Scripting Language. In *Proceedings of the 37th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '10, pages 377–388, New York, NY, USA, 2010. ACM. ISBN 978-1-60558-479-9. doi: 10.1145/1706299.1706343.