A Vision for Online Verification-Validation

Matthew A. Hammer

University of Colorado Boulder, USA
matthew.hammer@colorado.edu

Bor-Yuh Evan Chang

University of Colorado Boulder, USA
evan.chang@colorado.edu

David Van Horn

University of Maryland,
College Park, USA
dvanhorn@cs.umd.edu

Abstract

Today’s programmers face a false choice between creating software that is extensible and software that is correct. Specifically, dynamic languages permit software that is richly extensible (via dynamic code loading, dynamic object extension, and various forms of reflection), and today’s programmers exploit this flexibility to “bring their own language features” to enrich extensible languages (e.g., by using common JavaScript libraries). Meanwhile, such library-based language extensions generally lack enforcement of their abstractions, leading to programming errors that are complex to avoid and predict.

To offer verification for this extensible world, we propose online verification-validation (OVV), which consists of language and VM design that enables a “phaseless” approach to program analysis, in contrast to the standard static-dynamic phase distinction. Phaseless analysis freely interposes abstract interpretation with concrete execution, allowing analyses to use dynamic (concrete) information to prove universal (abstract) properties about future execution.

In this paper, we present a conceptual overview of OVV through a motivating example program that uses a hypothetical database library. We present a generic semantics for OVV, and an extension to this semantics that offers a simple gradual type system for the database library primitives. The result of instantiating this gradual type system in an OVV setting is a checker that can progressively type successive continuations of the program until a continuation is fully verified. To evaluate the proposed vision of OVV for this example, we implement the VM semantics (in Rust), and show that this design permits progressive typing in this manner.

Categories and Subject Descriptors D.3.1 [Programming Languages]: Formal Definitions and Theory; D.3.3 [Pro-

gramming Languages]: Language Constructs and Features; F.3.2 [Logics and Meanings of Programs]: Semantics of Programming Languages

Keywords Online Verification-Validation, Virtual Machines, For-all Analysis, Exists Analysis, Phaseless Analysis, Type Systems, Call-by-Push-Value (CBPV).

1. Introduction

We consider the problem of typing unknown, dynamically-determined data obtained from the environment. To illustrate what we mean by unknown, dynamically-determined data, consider the code in Figure 1a that loads a comma-separated value (CSV) file `authors.csv` using the `openDb("authors.csv")` call on line 1. The contents of this file are organized into lines, where each line is a row of field values. The first line of the file is special and contains a list of field names instead of field value—see Figure 1b for example content. In line 1 of Figure 1a, the programmer filters the author list down to those that have US citizenship using the field projection `author.citizenship`.

Suppose the programmer merely wants to know that their code *will not access undefined fields*—that, in this respect, the program is well-typed. Given that these object fields are defined by the *dynamic-generation of data structures* via `openDb`, the validity of this field projection `author.citizenship` is generally unknowable until after line 1 when the structure of the `authors` table is defined based on the contents of `authors.csv`. At the same time, this field projection is clearly valid because of the special first line in specifically this `authors.csv` shown in Figure 1b.

There is a tension here between static and dynamic checking. On one hand, once the `authors` table has been determined, the programmer would like the field projections from the rows to be statically verified. But on other hand, how would static checking be feasible if the schema of `authors` database is dynamically-determined—what would be the type of the `openDb` function?

To resolve this tension, our key insight is to imagine pausing the concrete execution of Figure 1a right after line 1 and before line 2. Then, imagine reflecting the continuation of the program’s execution to apply abstract interpretation or

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.

GPCE’16, October 31 – November 1, 2016, Amsterdam, Netherlands
© 2016 ACM. 978-1-4503-4446-3/16/10...\$15.00
<http://dx.doi.org/10.1145/2993236.2993255>

```

1 let authors      = openDb("authors.csv")
2 let authorsUS    = filterDb(authors, (author) =>
                                author.citizenship === "US")

```

(a) Open a database with an implicit schema and filter the content.

```

name, citizenship, institution
Evan Chang, US, University of Colorado
Roly Perera, UK, University of Glasgow
Matthew Hammer, US, University of Colorado
David Van Horn, US, University of Maryland

```

(b) File authors.csv, example content.

Figure 1. Programming with dynamically-determined tables and objects in JavaScript

other static techniques to prove the validity of *all* subsequent field projections on the rows of the authors table.

Our vision centers around a new paradigm for verification and validation that we call *online verification-validation* (OVV). In contrast to today’s *phasic* analysis techniques, where static and dynamic checking are separated into distinct, non-overlapping phases, OVV is *phaseless*: “static” analysis is freely interposed with “dynamic” execution. By virtue of this mixed approach, OVV transcends the conventional phase distinctions of “static” and “dynamic” analysis. To avoid confusion in this phaseless setting, we refer to static analysis techniques as those for $\forall$ -analysis, since they demonstrate *universal* properties of future program’s states. Likewise, we refer to the techniques used in dynamic analysis as those for $\exists$ -analysis, since they demonstrate *existential* properties about the past program’s states.

In this paper, we make the following contributions:

- We define λ -VMF, an abstract machine semantics that enables a form of phaseless, online verification-validation (Section 3). The key idea is annotating potentially faulting operations (e.g., field projection) with a *certain* (!) or *uncertain* (?) flag, indicating whether or not that operation can be verified. Executing uncertain operations gets stuck, so it is up to the OVV program to progressively rewrite uncertain operations into certain ones by proving the safety of the potentially faulting operation.
- We present a case study of instantiating online verification-validation with a simple, bidirectional gradual type system for dynamic field projection and databases with dynamic schemas (Section 4). The result of instantiating a gradual type system in an online verification-validation setting is a checker that can progressively type successive continuations of the program until a continuation is fully verified (i.e., “statically” typed).¹
- We implement the proposed design for λ -VMF in Rust to demonstrate that the proposed instantiation of OVV

¹ A *continuation* of a running program represents its state, including the remainder of the code to be executed in the future.

```

1 let authors      = openDb("authors.csv")
2 let authorsUS    = filterDb(authors, (author) =>
                                author.citizenship === "US")
3 let books        = openDb("books.csv")
4 let authbooksUS  = joinDb(authorsUS, "name",
                            books, "author")

```

(a) The joinDb(db1, key1, db2, key2) library function accesses the key fields using run-time reflection (e.g., db1[i][key1]).

```

author, title, year, publisher
...

```

(b) File books.csv, example content.

Figure 2. Continuing the example from Figure 1a with subsequent dynamically-determined data.

Figure 3. Typical type checkers are phasic with an offline static verifier and an online dynamic validator.

indeed realizes progressive typing. Our implementation is public: <https://github.com/cuplv/vmfuture>.

In the next section, we dig deeper into what we term online verification-validation by following the *progressive typing* of the example program from Figure 1a.

2. Overview

In Figure 2a, consider an extension of the example from Figure 1a with two additional lines. The file books.csv is opened on line 3 and contains a CSV file with book information (shown in Figure 2b). The final line (line 4) creates a table of books written by US authors, along with the information about those authors.

Observe that this code alternates between *dynamically determining* the types of the rows and tables (by reading the two files authors.csv and books.csv on lines 1 and 3) and computing over those tables (on lines 2 and 4). This example is a simple version of a pervasive pattern in dynamic languages, where execution interleaves dynamic steps that create new data types (the tables loaded in lines 1 and 3) and steps that compute over previously-defined data types (the filtering and joining steps in lines 2 and 4).

Phaseless Analysis. The essence of online verification-validation is pausing concrete execution to interleave it with

abstract interpretation. Pictorially, this interleaving of concrete and abstract execution can be visualized as a two-dimensional grid, as in Figure 3a. The vertical axis represents the extent of concrete execution and dynamic, $\exists$ -analysis (measured by program line), and for each such point, the horizontal axis represents the extent of abstract interpretation and static, $\forall$ -analysis performed at this dynamic execution point. As shown, abstract interpretation explores the states after lines 2–4 after pausing at line 1, but revisits the state after line 4 again after pausing at line 3. That is, we imagine suspending the concrete execution of the program in Figure 2a after the `openDb` call on line 1 and then interpreting the continuation to statically compute, under an abstraction, the set of reachable future states from this suspended current state (the horizontal axis).

This online look-ahead would permit us to check the projections of `author.citizenship` and `db1[i]["name"]` (on lines 2 and 4, respectively) before the running program executes them concretely. However, the other projection in the `joinDb` call on line 4—the projection of the `author` field (i.e., `db2[j]["author"]`)—cannot be proven valid in this continuation because the `books.csv` database has not yet been loaded. But imagine similarly suspending the concrete execution again after line 3 (shown as the bottom horizontal execution in Figure 3a). Now in this continuation, an abstract interpretation can easily verify this last remaining projection on all future concrete executions from this point.

For exposition, this example is short, and the distance between the concrete points for static checks and the future concrete points of potential failure (a bad projection) are tiny: they only consist of one or two lines. In the general case, however, the distance between these points can be arbitrarily large. For instance, a long-running scientific workload may last *days* or *weeks*, and in these cases, it is critical to know about possible future execution failures as soon as possible, to minimize interruption due to programming errors.

Related Work: Traditional Analyses are Phasic. Typical, existing program analyses impose *phase distinctions* between static and dynamic steps, forcing each phase to use one approach or the other. At one extreme, today’s techniques for static verification explore all possible execution paths, but have no knowledge of the dynamic execution environment. At the other extreme, dynamic validation explores one path of execution: the path determined by concrete execution.

Type system design is, in general, a tradeoff in checking in an offline, static phase (e.g., are function application expressions well-typed?) and checking in the online, dynamic phase (e.g., is an array-index in bounds?). We illustrate this phasic architecture in Figure 3b where the static and dynamic phases are sequenced and independent.

Some techniques attempt to “blend” these static and dynamic phases, so that information gleaned from one phase

```

1 function openDb(file) {
2   var table = openDbInternal(file)
3   {{ reflect/cc (chk_state, return table) }}
4 }

```

Figure 4. The `openDb` library function synthesizes a type for the loaded table before returning via `reflect/cc`.

feeds into the other. For instance, gradual typing [36, 37] enables shifting type checking between the static and dynamic phases. Or, the notion of using a set of dynamic runs to glean information *before* static verification appears several times in the literature for call resolution [15], reflection instantiation [9] and for `eval` [20, 41]. We consider such techniques phasic if there is some sequencing (rather than interleaving) of static, $\forall$ -analysis and dynamic, $\exists$ -analysis phases.

While phasic analyses dominate the literature, there exist some non-phasic analysis techniques, such as the “proofs-from-tests” approach [6], that mix $\exists$ -analysis information during $\forall$ -analysis (e.g., via directed-random automated testing [22]). In these works, the goal is to perform *offline* static verification whose abstraction selection leverages $\exists$ -analysis information from testing. This is distinct from our vision, where concrete execution is interposed with analyses (not vice versa). However, this work shares our concern with incremental exploration of a state space, and it possible that our proposed incremental substrate would also be beneficial in the context of “proofs-from-tests.”

Online Verification-Validation. OVV programs consist of two stratified layers that interact during execution. First, the *object layer* expresses ordinary execution. Execution can *escape* into the *meta layer*, which is capable of inspecting the run-time representation of the object layer. Code at the meta layer expresses $\forall$ -analysis and $\exists$ -analysis over object programs. During ordinary execution, the meta layer plays a passive role until special primitives transfer control. In particular, `reflect/cc` transfers control, along with a reflected view of the current (object layer) continuation. Further, the meta layer has access to read and write hidden *annotations* on object layer values (e.g., to store program facts, such as types). Collectively, these hidden annotations can be viewed as providing a “shadow heap” for tracking dynamic, meta-level information, in the service of performing $\exists$ -analyses. For maximum extensibility, the object layer lacks a static type system, relying on checking at the other layer. For convenience in expressing analysis over object programs, the meta layer may employ a language that employs a static type system (a la ML), but this is not a requirement. For concreteness in presentation, we use JavaScript syntax for object layer code and Rust-like syntax for meta layer code, and we may consider language choice as an orthogonal concern.

`reflect/cc.` The implementor of `openDb` uses `reflect/cc` to mediate between concrete execution in the object layer, and interposed code in the meta layer that performs on-

Figure 5. Illustrating online verification-validation.

```
enum Stk {
  Halt,
  FrLet (Env, Var, Exp, Stk),
  FrApp (Val, Stk)
}

enum CTyp {
  Ret (VTyp),
  Arr (VTyp, CTyp)
}

meta chk_stk (st:StTyp, stk:Stk, ct:CTyp) -> Option<Stk>
{ match (ct, stk) {
  -, Halt => Some(stk),
  Ret(vt), FrLet(env, x, e, stk2) => ...
  Arr(vt, ct), FrApp(v, stk2) => ...
}}
```

Figure 6. Type checking continuations.

line $\forall$ -analysis (shown in Figure 4). First, `openDb` uses `openDbInternal` to load the given file and construct a table from its CSV content (line 2). Next, line 3 uses `reflect/cc` to invoke the meta-level function `chk_state`, which determines types for the table’s content (failing with an error if this content is malformed). As a side effect, it writes to a meta layer field on the `table` variable to record the type of the rows of the table so that it can be consumed in a subsequent meta layer execution. In particular, the `reflect/cc` primitive pauses the execution of the object program, giving control to a meta language function, here `chk_state`. Before transferring control, the primitive reflects the current continuation as a *first-class data structure* and uses it as an argument to the given meta language code, here `chk_state`, as we illustrate in Figure 5.

By using this reflected structure, the meta language code can perform arbitrarily complex $\forall$ -analysis. When it finishes, it returns control to the object program by returning a transformed program state. In this case, `openDb` uses `chk_stk` to type-check the program’s continuation at each call, given the type of the loaded table. Based on this type information, `openDb` either signals errors (if there is a type error), or transforms future validation checks (if type analysis succeeds in proving that these validation checks are redundant).

The Meta Program Expresses Online Meta Theory. Figure 6 lists the meta-layer function `chk_stk`, which checks a

VM stack against a store typing and computation type. We adopt Rust-like syntax (a recent dialect of ML). In particular, this meta-level code interacts with the store typing, stack and computation types as values, as if they are persistent (purely functional, applicative) inductively-defined structures, a la ordinary ML. (The top of Figure 6 defines these structures in Rust-like syntax). Because the meta layer has *meta-level access* to the semantic structures of the object layer, it is sufficiently powerful to compute meta-theoretical properties (either $\forall$ -analysis or $\exists$ -analysis). In particular, meta layer function `chk_stk` performs a $\forall$ -analysis (type-inference and checking), and uses these universal program facts to improve the efficiency of $\exists$ -analyses that will occur in the future.

At a high level, `chk_stk` traverses the frames of a VM stack (from top to bottom), and checks whether the form of the frame is consistent with the given computation type `ct`. There are three cases to consider, depending on whether the stack is empty (`Halt`), or has a top-most frame for a `let` binding or function application.

Defining and using OVV. Below, Section 3 defines the syntax and dynamic semantics of λ -VMF in detail; building on these definitions, Section 4 revisits the algorithm shown in Figure 6, showing it in the context of a larger system for gradually-typing programs that compute with databases, like the one shown in Figure 5.

3. OVV Machine Semantics

We present λ -VMF, an abstract machine semantics for libraries and programs that employ OVV. In Section 4, we instantiate this OVV framework with a gradual type system for checking programs that compute with simple databases.

Program Syntax. Figure 7 gives the syntax for λ -VMF programs. To streamline the definition of analyses and dynamic interpretation, λ -VMF syntactically separates program structure into expressions e and values v . Further, to permit a meta-layer, extension-defined analysis to annotate λ -VMF programs (as illustrated in Section 4), the recursive syntax of expressions and values consists of annotated pre-expressions $\dot{e}$ and annotated pre-values $\dot{v}$, respectively. The operational semantics of λ -VMF programs, defined below, does not directly depend on these annotations. However, an extension may use its annotations to prove properties that verify sound online program transformations. In these cases, the annotations can indirectly impact program behavior, i.e., by aiding a meta-level program transformation.

Pre-values consist of open and closed thunks, which represent suspended expressions, including all higher-order data (`othunk  $e$` and `thunk  $\rho e$`); closed thunks employ a closing environment ρ that maps the free variables of e to (closed) values. Base types consist of numbers (`num  $n$`), strings (`str  $s$`), boolean bits (`bool  $b$`) and reference cells (`loc  $\ell$`). Dictionaries map values to values (`dict  $\delta$`), modeling a row

Value	$v ::= \dot{v} @: a^v$	Annotated pre-value
Pre-Val.	$\dot{v} ::= \mathbf{othunk} \ e$	Open thunk
	$\quad \mathbf{thunk} \ \rho \ e$	Closed thunk
	$\quad \mathbf{dict} \ \delta$	Dictionary
	$\quad \mathbf{num} \ n$	Number
	$\quad \mathbf{str} \ s$	String
	$\quad \mathbf{bool} \ b$	Boolean
	$\quad \mathbf{loc} \ \ell$	Store location
	$\quad ()$	Unit value
	$\quad x$	Value variable
Dict.	$\delta ::= \varepsilon \mid \delta, v_1 \mapsto v_2$	Dictionaries of values
Expr.	$e ::= \dot{e} @: a^e$	Annotated pre-expression
Pre-Expr.	$\dot{e} ::= \mathbf{force} \ v$	Unsuspend (force) thunk
	$\quad \lambda x. e$	Function abstraction
	$\quad e \ v$	Function application
	$\quad \mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2$	Bind computed value
	$\quad \mathbf{ret} \ v$	Produce a value
	$\quad \mathbf{ref} \ v$	Allocate store reference
	$\quad \mathbf{set} \ v_1 \ v_2$	Store mutation
	$\quad \mathbf{get} \ v$	Store projection
	$\quad \mathbf{ext} \ v_1 \ v_2 \ v_3$	Dictionary extension
	$\quad v_1 [v_2]_{a^{\text{op}}}$	Dictionary projection
Annot.	$\quad e ? : a^e$	Annotation ascription
	$\quad \mathbf{rcc} \ e_m \ e$	Reflect current continuation
Meta Expr.	$a^{\text{op}} ::= ? \mid !$	Uncertain vs. certain
	$a^v ::= \dots$	Value annotation
	$a^e ::= \dots$	Expression annotation
	$e_m ::= \dots$	Meta-level programs

Figure 7. Syntax of λ -VMF Programs

of a database, or a record, where the typical notion of a field name is generalized to any value in λ -VMF.

Pre-expressions consist of forcing a suspended expression (**force** v), function abstraction ($\lambda x. e$), function application ($e \ v$), let-binding a returned value (**let** $x = e_1$ **in** e_2), returning a value (**ret** v), allocating, mutating and accessing mutable storage (**ref** v , **set** $v_1 \ v_2$, **get** v , respectively), updating the field of a record (**ext** $v_1 \ v_2 \ v_3$), projecting the field of a record ($v_1 [v_2]_{a^{\text{op}}}$).

Finally, λ -VMF includes special forms for ascribing a sub-expression with a manual annotation ($e ? : a^e$), and reflectively inspecting (and transforming) the current continuation via **reflect/cc** as core primitive, **rcc** $e_m \ e$. Notably, execution pauses *before* executing the local continuation e , and a common idiom consists of using a manual ascription there, to be discharged via the use of **rcc**.

As explained below, using **rcc** to prove and discharge ascriptions is actually necessary in λ -VMF, since they have no other form of dynamic semantics. For this purpose, the meta-level program e_m transforms the program state before its continuation resumes. We do not model the model-level programming language here; our current implementation uses Rust.

State	$\sigma ::= \langle \mu; \kappa; \rho; \dot{e} \rangle$	
Store	$\mu ::= \varepsilon \mid \mu, \ell \mapsto v$	Maps locations to values
Stack	$\kappa ::= \mathbf{halt}$	Empty stack
	$\quad \kappa :: (\rho, x. e)$	Waiting for return
	$\quad \kappa :: v$	Fun. application argument
Environment	$\rho ::= \varepsilon \mid \rho, x \mapsto v$	Maps variables to values

Figure 8. VM State: The store, stack and environment.

VM State syntax. Figure 8 defines the global state of the λ -VMF program: It consists of a store, mapping locations to mutable values (μ); a stack of evaluation context frames (κ), an environment mapping variables to values (ρ), and a pre-expression $\dot{e}$ that gives the current local continuation. Non-empty stacks give evaluation contexts for **let** bodies ($\kappa :: (\rho, x. e)$) and function application ($\kappa :: v$).

Dynamics of λ -VMF Programs. Figure 9 defines a small-step operational semantics over λ -VMF states. The rules for **let** and function application each push the stack with a frame that is eliminated by the rules for **ret** and function abstraction, respectively. In both cases, eliminating the frame consists of binding a value to a variable, and continuing the program. Forcing a thunk consists of unpacking its environment and expression, and continuing execution with them. Rules for allocating, mutating and accessing a reference cell in the store are each standard. Dictionary extension adds a field to a (possibly empty) dictionary; and dictionary projection selects a given field's associated value, returning it.

What makes λ -VMF particularly interesting is that there are no stepping rules for ?-mode field projection or for ascription. To avoid getting stuck at these operations, these operations should be either verified progressively or validated, perhaps immediately before executing, when all relevant information is available. To do so, the program uses **rcc** $e_m \ e$, which runs the meta-level program e_m on a reflected version of the current VM state: the rule constructs the current continuation, reflects this program state into a data structure, runs the meta-level term e_m , and then injects the resulting program state into a transformed continuation σ' . In OVV, this step sometimes verifies and validates operations before the program attempts to execute them, transforming the program to remove or modify them. In contrast to traditional phasic static verification or dynamic validation (cf. Figure 3b), this step does not need to be either eagerly before the entire execution or lazily just before the potentially faulting operation. By choosing the placement of **rcc**, the program can choose how eager or lazy the checking should be anywhere between these two extremes. We give a detailed example below.

4. Gradual Typing for Simple Databases

In this section we present a gradual type system for λ -VMF and libDb, an extension that permits us to express the

$\sigma \longrightarrow \sigma'$ λ -VMF state σ steps to state σ' .

$\langle \mu; \kappa; \rho; \text{let } x = (\dot{e}_1 @: -) \text{ in } e_2 \rangle$	$\longrightarrow$	$\langle \mu; \kappa :: (\rho, x.e_2); \rho; \dot{e}_1 \rangle$	
$\langle \mu; \kappa; \rho; (\dot{e} @: -) v \rangle$	$\longrightarrow$	$\langle \mu; \kappa :: \rho(v); \rho; \dot{e} \rangle$	
$\langle \mu; \kappa :: (\rho, x.(\dot{e} @: -)); \rho'; \text{ret } v \rangle$	$\longrightarrow$	$\langle \mu; \kappa; \rho, x \mapsto \rho(v); \dot{e} \rangle$	
$\langle \mu; \kappa :: v; \rho; \lambda x.(\dot{e} @: -) \rangle$	$\longrightarrow$	$\langle \mu; \kappa; \rho, x \mapsto v; \dot{e} \rangle$	
$\langle \mu; \kappa; \rho; \text{force } v \rangle$	$\longrightarrow$	$\langle \mu; \kappa; \rho'; \dot{e} \rangle$	when $\rho(v) = (\text{thunk } \rho' (\dot{e} @: -) @: -)$
$\langle \mu; \kappa; \rho; \text{ref } v \rangle$	$\longrightarrow$	$\langle \mu, \ell \mapsto v; \kappa; \rho; \text{ret } (\text{loc } \ell @: ?) \rangle$	when $\ell \notin \mu$
$\langle \mu; \kappa; \rho; \text{set } v_1 v_2 \rangle$	$\longrightarrow$	$\langle \mu, \ell \mapsto \rho(v_2); \kappa; \rho; \text{ret } (() @: ?) \rangle$	when $\rho(v_1) = (\text{loc } \ell @: -)$
$\langle \mu, \ell \mapsto v_2; \kappa; \rho; \text{get } v_1 \rangle$	$\longrightarrow$	$\langle \mu, \ell \mapsto v_2; \kappa; \rho; \text{ret } v_2 \rangle$	when $\rho(v_1) = (\text{loc } \ell @: -)$
$\langle \mu; \kappa; \rho; \text{ext } v_1 v_2 v_3 \rangle$	$\longrightarrow$	$\langle \mu; \kappa; \rho; \text{ret } (\text{dict } (\delta, \rho(v_2) \mapsto \rho(v_3)) @: ?) \rangle$	when $\rho(v_1) = (\text{dict } \delta @: -)$
$\langle \mu; \kappa; \rho; v_1 [v_2]_1 \rangle$	$\longrightarrow$	$\langle \mu; \kappa; \rho; \text{ret } v_3 \rangle$	when $\rho(v_1) = (\text{dict } (\delta, \rho(v_2) \mapsto v_3) @: -)$
$\langle \mu; \kappa; \rho; v_1 [v_2]_2 \rangle$			
$\langle \mu; \kappa; \rho; e ? : a^e \rangle$			
$\langle \mu; \kappa; \rho; \text{rcc } e_m (\dot{e} @: -) \rangle$	$\longrightarrow$	σ'	when $e_m(\llbracket \langle \mu; \kappa; \rho; \dot{e} \rangle \rrbracket) \Downarrow_{\text{meta}} \llbracket \sigma' \rrbracket$

Figure 9. Small-step, abstract machine semantics of λ -VMF

$e ::= \dots$	Existing forms (Figure 7)
$\text{openDb}_{a^{\text{op}}} v$	Open database by file path
$\text{filterDb}_{a^{\text{op}}} v_1 v_2$	Filter DB by predicate
$\text{joinDb}_{a^{\text{op}}} v_1 v_2 v_3 v_4$	Join DBs using keys' value

Figure 10. Database Library Forms

motivating example from Figure 2a. We tour the type system and illustrate how OVV progressively types and validates the operations in this example.

Syntax for libDb. Figure 10 extends the syntax from Figure 7 with the three operations implemented by libDb, for opening databases, filtering them with a predicate, and joining them using named fields. Each operation is parameterized by one or more argument values, and an operation annotation a^{op} that determines how to type operation. As with record field projection, this annotation determines whether the operation's pre-conditions for success have been fully verified via OVV. The gradual type system for libDb presented here uses different rules for certain (!) versus uncertain (?) reasoning modes.

Types for λ -VMF and libDb. Figure 11 instantiates the λ -VMF framework for a gradual type system. This system has (bidirectional, algorithmic) rules to reason about λ -VMF code as well as the libDb extension. Value types consist of types for thunked computations ($\mathbf{U} C$)², dictionaries (**Dict** Δ , where Δ maps field values to field types), numbers (**Num**), strings (**Str**), booleans (**Bool**), reference cells (**Ref** A), unit (1), unknown (?) and databases (**Db** A). Computation types consist of the arrow type for

Annotations	$a^v ::= A$	Value annotation
	$a^e ::= C$	Expression annotation
Value Types	$A, B ::= \mathbf{U} C$	Thunked computation
	Dict Δ	Dictionary
	Num	Number
	Str	String
	Bool	Boolean
	Ref A	Reference cell
	1	Unit
	?	Unknown value type
	Db A	Database; multiset of A s
Dictionary	$\Delta ::= \varepsilon \mid \Delta, v \mapsto A$	Maps values to types
Computation	$C, D ::= A \rightarrow C$	Function abstraction
Types	F A	Value production

Figure 11. Type Syntax: Annotations for Values and Expressions

functions ($A \rightarrow C$) and value types for value-producing computations (**F** A).

Typing λ -VMF program states. Figure 12 lists typing judgement forms for λ -VMF program states and for stacks. To type a program state, we assume that the stored values are annotated, and we use these annotations as a store typing, written $|\mu|$, which maps reference locations to values types. To type a program state, we assume this store typing $|\mu|$ and attempt to verify the other VM machinery, consisting of the current environment ρ , program term $\dot{e}$, and stack κ . Three judgements compute type properties for these components: Assuming a store typing Γ , the judgement $\Gamma \vdash \rho \Rightarrow \Gamma'$ computes from an environment (mapping variables to values), a typing context Γ' (mapping variables to types). Assuming a typing context Γ , the judgement $\Gamma \vdash e \Rightarrow C$ computes a type from a term e . Assuming

² We follow conventions from the literature on call-by-push-value (CBPV) in our type syntax for thunked and value-returning computations, which uses special letters **U** and **F**, respectively [30, 31].

$\sigma \text{ ok}$ State σ is well-typed.

$$\frac{|\mu| \vdash \rho \Rightarrow \Gamma \quad |\mu|, \Gamma \vdash \dot{e} \Rightarrow C \quad |\mu| \vdash \kappa \Leftarrow C}{\langle \mu; \kappa; \rho; \dot{e} \rangle \text{ ok}} \text{ state}$$

$\Gamma \vdash \kappa \Leftarrow C$ Under Γ , stack κ eliminates a computation of type C .

$$\begin{array}{c} \frac{}{\Gamma \vdash \text{halt} \Leftarrow C} \text{ k-emp} \quad \frac{\Gamma_1 \vdash \rho \Rightarrow \Gamma_2 \quad \Gamma_2, x : A \vdash e \Rightarrow C \quad \Gamma_1 \vdash \kappa \Leftarrow C}{\Gamma_1 \vdash \kappa :: (\rho, x.e) \Leftarrow \mathbf{F} A} \text{ k-let} \\[10pt] \frac{\Gamma \vdash v \Leftarrow A \quad \Gamma \vdash \kappa \Leftarrow C}{\Gamma \vdash \kappa :: v \Leftarrow A \rightarrow C} \text{ k-app} \end{array}$$

Figure 12. Stack typing and State typing

a typing context Γ and computation type C for a terminal computation, the judgement $\Gamma \vdash \kappa \Leftarrow C$ checks that the stack κ either correctly continues execution or halts.

The remainder of the figure gives three rules for type-checking the stack. First, k-emp says that halting stacks are always permitted. Next, k-let and k-app handle the recursive cases of the stack, where the topmost frame can be viewed as eliminating the terminal computation type, call it D . In the case of k-let, we have that D is $\mathbf{F} A$, which types the terminal computation that returns a value of type A ; we check that the top of the stack holds the body of the let, which can use the let-bound variable (of type A), for which we can synthesize another computation type C that checks against the rest of the stack. In the case of k-app, we have that D is $A \rightarrow C$, the type of a function abstraction; we check that the top of the stack is an argument value of type A , and the rest of the stack checks against the type of the abstraction's body, C .

Transforming λ -VMF program states. Though the gradual type system defined here is stated propositionally, it constitutes an algorithm, and we demonstrate this fact by implementing these relational definitions as a (mutually) recursive total functions. However, instead of merely returning `true` or `false` to indicate the success or failure of the relation to hold, in the case of `true`, we also construct an annotated term, possibly with transformations (e.g., changing operation annotations from uncertain $?$ to certain $!$).

For instance, we implement the type relation for program states as a total function from program states to (optional) program states with annotations; and when the algorithm fails, it returns `None`. Furthermore, this algorithm plays the role of e_m in `libDb`'s use of `rec e_m e`.

The ability to phrase the typing relations as functional algorithms stems the fact that the rules treat certain positions of their (bidirectional) relations consistently as inputs and outputs, and that outputs are determined functionally from

inputs. As an example, Figure 6 gives the algorithmic version of the stack-checking relation $\Gamma \vdash \kappa \Leftarrow C$, which resembles an ordinary function in ML. The remainder of the rules transform in a similar manner, so that checking relations produce an optional, annotated term structure, while synthesizing relations produce an optional pair of annotated term structure and synthesized type. When these functions produce `None`, the corresponding typing relation is not derivable. The dynamic semantics of λ -VMF do not permit execution to continue when this occurs.

Typing core λ -VMF terms bidirectionally. For simplicity, we use a bidirectional type system to encode the gradual type systems of the λ -VMF core calculus and its `libDb` extension. Figure 13 defines type checking (above) and synthesis (below) for program terms. For space reasons, we elide some synthesis cases, as well as the checking and synthesis judgements for value forms; the rules shown give a representative flavor for the complete definition.

The analytical (checking) judgement form $\Gamma \vdash e \Leftarrow C$ can be read as, “Under typing context Γ , term e checks against computation type C .” Specifically, the type C is given as an input to the checking judgement, when viewed as an algorithm. The synthesizing judgement form $\Gamma \vdash e \Rightarrow C$ can be read as, “the typing context Γ and term e synthesize the computation type C .” Specifically, the algorithm *computes* the type C , when given Γ and e . For the core forms of λ -VMF, the bidirectional rules for values and computations follow the usual patterns found in bidirectional type systems [10, 16]. We show several standard-looking rules, sub, lam, app and annot. In particular, the annotation form of λ -VMF, $e ? : a^e$, which asserts the annotation a^e correctly describes the program e , plays the role of type ascription in the bidirectional rules; the annot rule says that terms are checked against their type annotations, and these annotated terms synthesize the annotation type. Because it has no dynamic semantics, λ -VMF uses OVV to prove and discharge this form earlier by rewriting it to e sometime before evaluation; if this rewrite fails, then the program terminates (by failing) early, as a result of OVV failing, not execution.

As is customary in bidirectional systems, type annotations mediate between synthesizing and checking. This provides one the ability to place checking-only terms (such as lambda abstractions) in positions that require the sub-term synthesize a type (such as the abstraction position of an application). See typing rules lam and app, respectively, to see the details; both are standard. Finally, as is customary, type *subsumption* allows less specific types to check against terms that synthesize more specific types. Rule sub uses a definition of type *consistency* (written $C \approx? D$), which behaves like type equality, modulo the uncertain type $?$, which is consistent with all other types. This notion of consistency is standard in some gradual typing literature [36].

Gradual typing for dictionary projection. The typing rules for uncertain and certain projection differ in what is known

$\boxed{\Gamma \vdash e \Leftarrow C}$ Under Γ , expression e checks against type C .

$$\frac{\Gamma \vdash e \Rightarrow C \quad C \approx_? D}{\Gamma \vdash e \Leftarrow D} \text{sub} \quad \frac{\Gamma, x : A \vdash e \Leftarrow C}{\Gamma \vdash \lambda x. e \Leftarrow A \rightarrow C} \text{lam}$$

$\boxed{\Gamma \vdash e \Rightarrow C}$ Under Γ , expression e synthesizes type C .

$$\frac{\Gamma \vdash e \Leftarrow C}{\Gamma \vdash e \text{ ?} : C \Rightarrow C} \text{annot} \quad \frac{\Gamma \vdash e \Rightarrow A \rightarrow C \quad \Gamma \vdash v \Leftarrow A}{\Gamma \vdash e v \Rightarrow C} \text{app}$$

$$\frac{\Gamma \vdash v_1 \Rightarrow ? \quad \Gamma \vdash v_2 \Rightarrow B}{\Gamma \vdash v_1 [v_2]_? \Rightarrow \mathbf{F} ?} \text{p?} \quad \frac{\Gamma \vdash v_1 \Rightarrow \mathbf{Dict}(\Delta, v_2 \mapsto A) \quad \Gamma \vdash v_2 \Rightarrow B}{\Gamma \vdash v_1 [v_2]_! \Rightarrow \mathbf{F} A} \text{p!}$$

$$\frac{\Gamma \vdash v \Rightarrow \mathbf{Str}}{\Gamma \vdash \text{openDb}_? v \Rightarrow \mathbf{F}(\mathbf{Db} ?)} \text{openDb?}$$

$$\frac{\Gamma \vdash v_1 \Rightarrow \mathbf{Db} ? \quad \Gamma \vdash v_2 \Leftarrow \mathbf{U} (? \rightarrow \mathbf{F} \mathbf{Bool})}{\Gamma \vdash \text{filterDb}_? v_1 v_2 \Rightarrow \mathbf{F}(\mathbf{Db} ?)} \text{filterDb?}$$

$$\frac{\Gamma \vdash v_1 \Rightarrow \mathbf{Db} A \quad ? \notin A \quad \Gamma \vdash v_2 \Leftarrow \mathbf{U}(A \rightarrow \mathbf{F} \mathbf{Bool})}{\Gamma \vdash \text{filterDb}_! v_1 v_2 \Rightarrow \mathbf{F}(\mathbf{Db} A)} \text{filterDb!}$$

$$\frac{\Gamma \vdash v_1 \Rightarrow \mathbf{Db} ? \quad \Gamma \vdash v_2 \Rightarrow B_2 \quad \Gamma \vdash v_3 \Rightarrow \mathbf{Db} ? \quad \Gamma \vdash v_4 \Rightarrow B_4}{\Gamma \vdash \text{joinDb}_? v_1 v_2 v_3 v_4 \Rightarrow \mathbf{F}(\mathbf{Db} ?)} \text{joinDb?}$$

$$\frac{\Delta = \Delta_1, v_2 \mapsto A, \Delta_3, v_4 \mapsto A \quad \Gamma \vdash v_1 \Rightarrow \mathbf{Db}(\mathbf{Dict}(\Delta_1, v_2 \mapsto A)) \quad \Gamma \vdash v_2 \Rightarrow B_2 \quad \Gamma \vdash v_3 \Rightarrow \mathbf{Db}(\mathbf{Dict}(\Delta_3, v_4 \mapsto A)) \quad \Gamma \vdash v_4 \Rightarrow B_4}{\Gamma \vdash \text{joinDb}_! v_1 v_2 v_3 v_4 : \mathbf{F}(\mathbf{Db}(\mathbf{Dict} \Delta))} \text{joinDb!}$$

Figure 13. Selected typing rules for computation typing (checking and synthesis).

Typing (horizontal) across **Execution** (vertical).

	Line 1	Line 2	Line 3	Line 4	...
	openDb	filterDb	openDb	joinDb	—
1: openDb	?	?	?	?	?
2: filterDb	—	!	?	?	! / ?
3: openDb	—	—	?	?	! / ?
4: joinDb	—	—	—	!	!

Figure 14. Progressive typing for Figure 2a: As execution progresses (vertically) over Lines 1 and 3, the continuation's typing becomes more certain: After Line 1 executes, Line 2 types in the certain modality (!) instead of the uncertain modality (?); similarly, after Line 3, Lines 4 onward type using the certain modality. (The dashed horizontal lines indicate these progressions).

about the record and field values, and illustrate a form of gradual typing. In the uncertain case, rule p? synthesizes return type ?, since nothing is known about the dictionary of values being projected. By contrast, in the certain case, the dictionary type is known to rule p!, and this dictionary maps the given field value to a corresponding field type. In this case, the soundness of the type system means that the projection must succeed in all possible future program states, and moreover, that the projected value has the given type.

Typing the libDb operations. The rule openDb? is uncertain and has no certain counterpart: The type of the database is not known until after the operation completes, just before execution resumes with its continuation; before then, the database could hold any type, so the rule types the returned database as Db ?.

Following similar reasoning, since filtering and joining databases occur *after* a database is loaded, it is possible to type these operations in both uncertain and certain modes. The rule filterDb? says that filtering a database of uncertain values leads to another database of values with an uncertain type; since it merely assumes the type of the database is ?, it does not prove that the predicate will not “go wrong”, e.g., by projecting the wrong field from its argument. By contrast, the rule filterDb! says that filtering a database of known type using a predicate that checks against this type leads to a database with the same known type; in this case, the soundness of the type system means that the predicate must always succeed.

Similarly, the rules for joinDb? and joinDb! follow the pattern set above: the uncertain rule assumes nothing about the argument values, beyond the arguments actually consisting of databases. The certain rule assumes that the database arguments' types are fully known, that the chosen field values are mapped in these types, and that the chosen fields share a common type (we want to compare values of this field for equality to perform the join).

Gradual Typing, Progressively via OVV. Figure 14 illustrates using the typing rules of Figure 13 to perform progressive typing our four-line motivating example (Figure 2a). The vertical and horizontal dimensions of the table list each of the four lines; the vertical axis represents concrete execution, and the horizontal axis represents typing the four right-hand-sides of the program's four let-bindings, and in particular, for filterDb (on Line 2) and joinDb (on Line 4), the table indicates whether the operation was typed in the uncertain or certain modality. As execution progresses (vertically) over Lines 1 and 3, the typing of the program continuation's becomes more certain: After Line 1 executes, Line 2 types in the certain modality (!) instead of the uncertain modality (?); similarly, after Line 3, Lines 4 onward type using the certain modality. As the final column shows, the certainty of code using these tables in the remainder of the program (Line 5 onwards) increases after each of the two calls to openDb. Between the two calls, some information is

Figure 15. Progressive verification and regressive validation are complementary ways of seeing the interplay between $\forall$ -analysis and $\exists$ -analysis enabled by online verification-validation.

known (relating to the first two tables defined on Lines 1 and 2), but some information is still missing (relating to the two tables defined in Lines 3 and 4).

Our current Rust-based prototype of λ -VMF is powerful enough to express this example, including the progressive typing discussed above. In Section 5, we discuss the potential to use incremental computation in the context of such progressive typing; the goal is to improve performance by exploiting the redundancy of re-typing the program’s continuation.

5. Discussion

In this section, we discuss future challenges and directions for the vision of OVV presented in this paper. Specifically, we discuss the design of the meta-level programming language, and its use in expressing *progressive verification* and *regressive validation*.

Incremental Computation for Progressive Verification.

In the motivating example from Section 2, the `chk_stk` calls performed by `openDb` use progressive typing to check their continuations. In fact, these two continuations are related: The earlier version lacks type information about the table loaded in line 3, whereas the later version has access to this type information. Progressive typing could exploit this incremental relationship to avoid re-computing all of the typing facts about the program state that have *not changed*.

Progressive typing is a specific instance of progressive verification. Pictorially, progressive verification relates distinct abstract executions (shown horizontally in Figure 15a), by exploiting their similarity (a small change, depicted as δ , extending vertically). The task of a progressive verifier is a $\forall$ -analysis, just like a classical static verifier, which attempts to prove that *all* executions to an assertion satisfy a particular safety property.

More Aggressive Regressive Validation. After type checking the continuation for line 1, *regressive validation* con-

sists of eliminating dynamic checks within the library calls of lines 2 and 4 (`filterDb` and `joinDb`, respectively). In particular, if the continuation type-checks under the partially-known type information, the known type information can be used to elide run-time $\exists$ -analysis checks that concern the authors table, including the `asserts` for the projections of `author.name`. Pictorially, we think of regressive validation as introducing a third dimension that consists of all possible outcomes of a program transformation on the object program (shown in Figure 15b). After performing an online verification, the meta layer transforms the continuation, either eliding certain downstream validation checks (labeled **validate**), or introducing residual checks that reduce the original checks’ complexity (labeled **residual**). As illustrated in Section 4, progressive typing can, before executing the operation, eliminate residual checks by rewriting uncertain operations to certain operations (which require no run-time type checks). More aggressive verification techniques can hope to regress even more aggressive validation checks to simpler forms. For instance, global heap-based properties present an interesting challenge.

Implicitly-Incremental Meta-Level Computation. In sum, the example `chk_stk` above encodes a theory about *online typing*, along with a mechanism for using this type information to optimize the dynamic run-time checks that would otherwise be used. The meta layer should have an in-built ability to implicitly express progressive verification as ordinary verification, so that the system, not the programmer, takes into account execution environment changes across these progressive stages. Further, when the meta layer uses this progressive verification to enable regressive validation in the future execution down stream, the system, not the programmer, accounts for these changes when doing future stages of progressive verification. In other words, the VM that runs the meta language and object language should have an in-built ability to express interaction among the levels in terms of implicitly-incremental computation. In Section 6, we discuss the challenges that OVV poses to work on general-purpose incremental computation.

6. Related Work

This section supplements the related work in Section 2.

General-Purpose Incremental Computation. Section 5 proposes an implicitly incremental meta-level language for λ -VMF, which challenges current research on (general-purpose, programming language-based) incremental computation (IC). Consider the desired incremental behavior of `chk_stk` in the motivating example, where it occurs after lines 1 and 3, when the program states are similar, but not identical. In particular, both continuations include the AST of the call in line 4, and onward, which `chk_stk` will process in *both* verification stages. The central challenge is reusing the redundant work performed by `chk_stk`, despite the fact

that the AST and store typing *are not equal to that in the prior stage*, which creates challenges for incremental computing via *memoization*, a key implementation mechanism used across many specific IC approaches.

To understand why these “small” changes are challenging for typical memoization, consider the structural recursion of the `chk_stk` function from Figure 6. One approach to memoization identifies each saved invocation by the entire store typing `st` and entire stack `stk`, including all of their recursive sub-structure, e.g., via *hash-consing* [18]. This approach is commonly taken by past work on incremental computation [7, 8, 17, 24, 25, 33, 34, 34], however, it is brittle and overly sensitive to small changes, since they alter the identity of the whole recursive structure. Recent work addresses this shortcoming by introducing *unique names* that are special to incremental computing [26]. This naming mechanism can overcome the challenges outlined above for `chk_stk`, since the presence of names isolates changed components of the store typing, stack and local environments.

However, several key challenges remain before these techniques can fully realize OVV: We want to use these names *correctly* (to avoid unsound incremental results), use them *efficiently* (to isolate changes and avoid sub-redundant computations) and use them *implicitly* (so that the meta-level programs look like ML). Further, we may want to control how fine-grained the IC techniques track program dependencies, to reduce constant-factor overhead.

Reflective Towers of Interpreters. As proposed in Section 5, future work on λ -VMF should permit library extension authors to write meta programs and object programs in an integrated way. Fortunately, many researchers have proposed designs that allow interesting interplay between the interpreter’s viewpoint (where the meta-level program runs) and the program being interpreted (where the object program runs). Conceptually, this work begins with 3-LISP [14, 38], which gives the programmer access to an *infinite tower* of (so-called *meta-circular*) interpreters, allowing them to re-define the language from within the language. Following (theoretical) work on 3-LISP, researchers give various approaches that attack practical concerns in how to express and implement reflective towers in simpler terms; these efforts are named after various hair colors: Brown [19, 40], Blond [13] and most recently, Black [3, 4].

Compared to the impressive and mind-bending work on metacircular interpreters, the vision for λ -VMF is more modest: Two levels suffice to perform OVV. Having said that, if meta-level programmers want to verify their meta-level programs as object programs (to “bootstrap” a typed meta level), the work mentioned above will likely provide further insights.

Program Analysis for Dynamic Languages. The ultimate aim of online verification-validation is offer “strong checking” in an extensible, dynamic language environment. And thus we seek to build on the substantial amount of research

activity on program analysis for dynamic languages. Since by definition, dynamic languages lack a built-in static typing discipline, much of the static verification work focuses on either retrofitting rich typing or specification disciplines [11, 12, 21] or applying whole-program flow analysis for inferring and checking type properties (e.g., for JavaScript [5, 27, 28] or for Ruby [1]). Researchers have also used static type inference to instrument dynamic programs with assertions that fail *preemptively* [23]; while similar in spirit to OVV, such approaches are still phasic, since type analysis occurs only once, before program execution begins.

The dynamic language features that make widely-used libraries like jQuery possible also make retrofitting static techniques incredibly challenging [2, 29, 35, 39]. Much of this work focuses on finding the right kinds of context-sensitivity to try to more precisely resolve the flow of values to dynamic features like dynamic property read in a static analysis [2, 32, 39] or to determine when dynamically-observed information is sufficient to apply in a static verification [35]. We expect such techniques to be not only applicable and useful but strengthened in an OVV context. In the end, static techniques in a phasic setting are limited by what is indeed available statically, and dynamic techniques are limited by what can be observed in testing runs. As exhibited in Section 2, the vision of OVV enables these techniques to be strengthened with a flexible interleaving of “static” and “dynamic” analysis (i.e., $\forall$ -analysis and $\exists$ -analysis).

7. Conclusion

This paper presents a vision for online verification-validation (OVV), an approach to ease the tension between extensibility (of dynamic languages) and safety (of static languages). The key insight of OVV is that analysis in a VM can be *phaseless*, allowing analyses to run progressively on the object program by pausing execution, reflecting on the current continuation, and transforming the continuation to replace *uncertain* (?) operations with *certain* (!) ones.

In this paper, we formalize an approach for OVV as a language semantics and Rust-based implementation called λ -VMF. We explore a proof-of-concept instantiation of OVV by defining a gradual type system for dynamic field projection and databases with dynamic schemas, and we observed that the result is a progressive type checker.

Acknowledgments

This material is based in part upon work supported by a gift from Mozilla, and support from the National Science Foundation under grant numbers CCF-1619282, CCF-1618756, CCF-1055066 and CCF-1218208. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of Mozilla, or the National Science Foundation.

References

- [1] Jong-hoon (David) An, Avik Chaudhuri, and Jeffrey S. Foster. Static typing for Ruby on Rails. In *Automated Software Engineering (ASE)*, 2009.
- [2] Esben Andreasen and Anders Møller. Determinacy in static analysis for jQuery. In *Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*, 2014.
- [3] Kenichi Asai. Reflection in direct style. In *Generative Programming and Component Engineering (GPCE)*, 2011.
- [4] Kenichi Asai, Satoshi Matsuoaka, and Akinori Yonezawa. Duplication and partial evaluation for a better understanding of reflective languages. *Lisp and Symbolic Computation*, 9(2-3), 1996.
- [5] SungGyeong Bae, Hyunghun Cho, Inho Lim, and Sukyoung Ryu. SAFEWAPI: web API misuse detector for web applications. In *Foundations of Software Engineering (FSE)*, 2014.
- [6] Nels E. Beckman, Aditya V. Nori, Sriram K. Rajamani, Robert J. Simmons, Sai Deep Tetali, and Aditya V. Thakur. Proofs from tests. *IEEE Transactions on Software Engineering*, 2010.
- [7] Pramod Bhatotia, Alexander Wieder, Rodrigo Rodrigues, Umut A. Acar, and Rafael Pasquin. Incoop: MapReduce for incremental computations. In *Cloud Computing (SoCC)*, 2011.
- [8] Pramod Bhatotia, Pedro Fonseca, Umut A. Acar, Björn B. Brandenburg, and Rodrigo Rodrigues. iThreads: A threading library for parallel incremental computation. In *Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2015.
- [9] Eric Bodden, Andreas Sewe, Jan Sinschek, Hela Oueslati, and Mira Mezini. Taming reflection: Aiding static analysis in the presence of reflection and custom class loaders. In *International Conference on Software Engineering (ICSE)*, 2011.
- [10] Adam Chlipala, Leaf Petersen, and Robert Harper. Strict bidirectional type checking. 2005.
- [11] Ravi Chugh, David Herman, and Ranjit Jhala. Dependent types for JavaScript. In *Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*, 2012.
- [12] Ravi Chugh, Patrick Maxim Rondon, and Ranjit Jhala. Nested refinements: a logic for duck typing. In *Principles of Programming Languages (POPL)*, 2012.
- [13] Olivier Danvy and Karoline Malmkjær. Intensions and extensions in a reflective tower. In *LISP and Functional Programming*, 1988.
- [14] Jim des Rivières and Brian Cantwell Smith. The implementation of procedurally reflective languages. In *LISP and Functional Programming*, 1984.
- [15] Bruno Dufour, Barbara G. Ryder, and Gary Seivitsky. Blended analysis for performance understanding of framework-based applications. In *Software Testing and Analysis (ISSTA)*, 2007.
- [16] Joshua Dunfield and Neelakantan R. Krishnaswami. Sound and complete bidirectional typechecking for higher-rank polymorphism with existentials and indexed types. *CoRR*, abs/1601.05106, 2016.
- [17] Sebastian Erdweg, Oliver Bracevac, Edlira Kuci, Matthias Krebs, and Mira Mezini. A co-contextual formulation of type rules and its application to incremental type checking. In *Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*, 2015.
- [18] Jean-Christophe Filliâtre and Sylvain Conchon. Type-safe modular hash-consing. In *Proceedings of the 2006 Workshop on ML*. ACM, 2006.
- [19] Daniel P. Friedman and Mitchell Wand. Reification: Reflection without metaphysics. In *LISP and Functional Programming*, 1984.
- [20] Michael Furr, Jong hoon (David) An, and Jeffrey S. Foster. Profile-guided static typing for dynamic scripting languages. In *Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*, 2009.
- [21] Philippa Gardner, Sergio Maffeis, and Gareth David Smith. Towards a program logic for JavaScript. In *Principles of Programming Languages (POPL)*, 2012.
- [22] Patrice Godefroid, Nils Klarlund, and Koushik Sen. DART: Directed automated random testing. In *Programming Language Design and Implementation (PLDI)*, 2005.
- [23] Neville Grech, Julian Rathke, and Bernd Fischer. Preemptive type checking in dynamically typed languages. In *Theoretical Aspects of Computing - ICTAC*, 2013.
- [24] Philip J. Guo and Dawson Engler. Using automatic persistent memoization to facilitate data analysis scripting. In *Software Testing and Analysis (ISSTA)*, 2011.
- [25] Matthew A. Hammer, Yit Phang Khoo, Michael Hicks, and Jeffrey S. Foster. Adapton: Composable, demand-driven incremental computation. In *Programming Language Design and Implementation (PLDI)*, 2014.
- [26] Matthew A. Hammer, Joshua Dunfield, Kyle Headley, Nicholas Labich, Jeffrey S. Foster, Michael W. Hicks, and David Van Horn. Incremental computation with names. In *Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*, 2015.
- [27] Simon Holm Jensen, Anders Møller, and Peter Thiemann. Type analysis for JavaScript. In *Static Analysis (SAS)*, 2009.
- [28] Vineeth Kashyap, Kyle Dewey, Ethan A. Kuefner, John Wagner, Kevin Gibbons, John Sarracino, Ben Wiedermann, and Ben Hardekopf. JSAI: a static analysis platform for JavaScript. In *Foundations of Software Engineering (FSE)*, 2014.
- [29] Benjamin S. Lerner, Liam Elbert, Jincheng Li, and Shriram Krishnamurthi. Combining form and function: Static types for jQuery programs. In *Object-Oriented Programming (ECOOP)*, 2013.
- [30] Paul Blain Levy. Call-by-push-value: A subsuming paradigm. In *Typed Lambda Calculi and Applications (TLCA)*, 1999.
- [31] Paul Blain Levy. *Call-by-push-value: A Functional/imperative Synthesis*, volume 2. 2003.
- [32] Changhee Park and Sukyoung Ryu. Scalable and precise static analysis of JavaScript applications via loop-sensitivity. In

- Object-Oriented Programming (ECOOP)*, 2015.
- [33] William Pugh. *Incremental Computation via Function Caching*. PhD thesis, Cornell University, 1988.
 - [34] William Pugh and Tim Teitelbaum. Incremental computation via function caching. In *Principles of Programming Languages (POPL)*, 1989.
 - [35] Max Schäfer, Manu Sridharan, Julian Dolby, and Frank Tip. Dynamic determinacy analysis. In *Programming Language Design and Implementation (PLDI)*, 2013.
 - [36] Jeremy G. Siek and Walid Taha. Gradual typing for functional languages. In *Scheme and Functional Programming Workshop*, 2006.
 - [37] Jeremy G. Siek and Walid Taha. Gradual typing for objects. In *Object-Oriented Programming (ECOOP)*, 2007.
 - [38] Brian Cantwell Smith. Reflection and semantics in Lisp. In *Principles of Programming Languages (POPL)*, 1984.
 - [39] Manu Sridharan, Julian Dolby, Satish Chandra, Max Schäfer, and Frank Tip. Correlation tracking for points-to analysis of JavaScript. In *Object-Oriented Programming (ECOOP)*, 2012.
 - [40] Mitchell Wand and Daniel P. Friedman. The mystery of the tower revealed: A nonreflective description of the reflective tower. *Lisp and Symbolic Computation*, 1(1), 1988.
 - [41] Shiyi Wei and Barbara G. Ryder. Practical blended taint analysis for JavaScript. In *Software Testing and Analysis (ISSTA)*, 2013.