

Practical Optional Types for Clojure

Ambrose Bonnaire-Sergeant¹(✉), Rowan Davies²,
and Sam Tobin-Hochstadt¹

¹ Indiana University, Bloomington, USA
{abonnair,samth}@indiana.edu

² Omnia Team, Commonwealth Bank of Australia, Sydney, Australia
Rowan.Davies@cba.com.au

Abstract. Typed Clojure is an optional type system for Clojure, a dynamic language in the Lisp family that targets the JVM. Typed Clojure enables Clojure programmers to gain greater confidence in the correctness of their code via static type checking while remaining in the Clojure world, and has acquired significant adoption in the Clojure community. Typed Clojure repurposes Typed Racket’s *occurrence typing*, an approach to statically reasoning about predicate tests, and also includes several new type system features to handle existing Clojure idioms.

In this paper, we describe Typed Clojure and present these type system extensions, focusing on three features widely used in Clojure. First, multimethods provide extensible operations, and their Clojure semantics turns out to have a surprising synergy with the underlying occurrence typing framework. Second, Java interoperability is central to Clojure’s mission but introduces challenges such as ubiquitous `null`; Typed Clojure handles Java interoperability while ensuring the absence of null-pointer exceptions in typed programs. Third, Clojure programmers idiomatically use immutable dictionaries for data structures; Typed Clojure handles this with multiple forms of heterogeneous dictionary types. We provide a formal model of the Typed Clojure type system incorporating these and other features, with a proof of soundness. Additionally, Typed Clojure is now in use by numerous corporations and developers working with Clojure, and we present a quantitative analysis on the use of type system features in two substantial code bases.

1 Clojure with Static Typing

The popularity of dynamically-typed languages in software development, combined with a recognition that types often improve programmer productivity, software reliability, and performance, has led to the recent development of a wide variety of optional and gradual type systems aimed at checking existing programs written in existing languages. These include TypeScript [19] and Flow [11] for JavaScript, Hack [10] for PHP, and mypy [15] for Python among the optional

```
(ann pname [(U File String) -> (U nil String)])
(defmulti pname class) ; multimethod dispatching on class of argument
(defmethod pname String [s] (pname (new File s))) ; String case
(defmethod pname File [f] (.getName f)) ; File case, static null check
(pname "STAINS/JELLY") ;=> "JELLY" :- (U nil Str)
```

Fig. 1. A simple typed Clojure program (delimiters: **Java interoperation** (green), **type annotation** (blue), **function invocation** (black), **collection literal** (red), other (gray)) (Color figure online)

systems, and Typed Racket [23], Reticulated Python [25], and GradualTalk [1] among gradually-typed systems.¹

One key lesson of these systems, indeed a lesson known to early developers of optional type systems such as StrongTalk, is that type systems for existing languages must be designed to work with the features and idioms of the target language. Often this takes the form of a core language, be it of functions or classes and objects, together with extensions to handle distinctive language features.

We synthesize these lessons to present *Typed Clojure*, an optional type system for Clojure. Clojure is a dynamically typed language in the Lisp family—built on the Java Virtual Machine (JVM)—which has recently gained popularity as an alternative JVM language. It offers the flexibility of a Lisp dialect, including macros, emphasizes a functional style via immutable data structures, and provides interoperability with existing Java code, allowing programmers to use existing Java libraries without leaving Clojure. Since its initial release in 2007, Clojure has been widely adopted for “backend” development in places where its support for parallelism, functional programming, and Lisp-influenced abstraction is desired on the JVM. As a result, there is an extensive base of existing untyped programs whose developers can benefit from Typed Clojure, an experience we discuss in this paper.

Since Clojure is a language in the Lisp family, we apply the lessons of Typed Racket, an existing gradual type system for Racket, to the core of Typed Clojure, consisting of an extended λ -calculus over a variety of base types shared between all Lisp systems. Furthermore, Typed Racket’s *occurrence typing* has proved necessary for type checking realistic Clojure programs.

However, Clojure goes beyond Racket in many ways, requiring several new type system features which we detail in this paper. Most significantly, Clojure supports, and Clojure developers use, **multimethods** to structure their code in extensible fashion. Furthermore, since Clojure is an untyped language, dispatch within multimethods is determined by application of dynamic predicates to argument values. Fortunately, the dynamic dispatch used by multimethods has surprising symmetry with the conditional dispatch handled by occurrence typing. Typed Clojure is therefore able to effectively handle complex and highly dynamic dispatch as present in existing Clojure programs.

¹ We use “gradual typing” for systems like Typed Racket with sound interoperation between typed and untyped code; Typed Clojure or TypeScript which don’t enforce type invariants we describe as “optionally typed”.

But multimethods are not the only Clojure feature crucial to type checking existing programs. As a language built on the Java Virtual Machine, Clojure provides flexible and transparent access to existing Java libraries, and **Clojure/Java interoperability** is found in almost every significant Clojure code base. Typed Clojure therefore builds in an understanding of the Java type system and handles interoperability appropriately. Notably, `null` is a distinct type in Typed Clojure, designed to automatically rule out null-pointer exceptions.

An example of these features is given in Fig. 1. Here, the `pname` multimethod dispatches on the `class` of the argument—for `Strings`, the first method implementation is called, for `Files`, the second. The `String` method calls a `File` constructor, returning a non-nil `File` instance—the `getName` method on `File` requires a non-nil target, returning a nilable type.

Finally, flexible, high-performance immutable dictionaries are the most common Clojure data structure. Simply treating them as uniformly-typed key-value mappings would be insufficient for existing programs and programming styles. Instead, Typed Clojure provides a flexible **heterogeneous map** type, in which specific entries can be specified.

While these features may seem disparate, they are unified in important ways. First, they leverage the type system mechanisms inherited from Typed Racket—multimethods when using dispatch via predicates, Java interoperability for handling `null` tests, and heterogeneous maps using union types and reasoning about subcomponents of data. Second, they are crucial features for handling Clojure code in practice. Typed Clojure’s use in real Clojure deployments would not be possible without effective handling of these three Clojure features.

Our main contributions are as follows:

1. We motivate and describe Typed Clojure, an optional type system for Clojure that understands existing Clojure idioms.
2. We present a sound formal model for three crucial type system features: multi-methods, Java interoperability, and heterogeneous maps.
3. We evaluate the use of Typed Clojure features on existing Typed Clojure code, including both open source and in-house systems.

The remainder of this paper begins with an example-driven presentation of the main type system features in Sect. 2. We then incrementally present a core calculus for Typed Clojure covering all of these features together in Sect. 3 and prove type soundness (Sect. 4). We then present an empirical analysis of significant code bases written in `core.typed`—the full implementation of Typed Clojure—in Sect. 5. Finally, we discuss related work and conclude.

2 Overview of Typed Clojure

We now begin a tour of the central features of Typed Clojure, beginning with Clojure itself. Our presentation uses the full Typed Clojure system to illustrate key type system ideas,² before studying the core features in detail in Sect. 3.

² Full examples: <https://github.com/typedclojure/esop16>.

2.1 Clojure

Clojure [13] is a Lisp that runs on the Java Virtual Machine with support for concurrent programming and immutable data structures in a mostly-functional style. Clojure provides easy interoperability with existing Java libraries, with Java values being like any other Clojure value. However, this smooth interoperability comes at the cost of pervasive `nil`, which leads to the possibility of null pointer exceptions—a drawback we address in Typed Clojure.

2.2 Typed Clojure

A simple one-argument function `greet` is annotated with `ann` to take and return strings.

```
(ann greet [Str -> Str])
(defn greet [n] (str "Hello," n "!"))
(greet "Grace") ;=> "Hello, Grace!" :- Str
```

Providing `nil` (exactly Java's `null`) is a static type error—`nil` is not a string.

```
(greet nil) ; Type Error: Expected Str, given nil
```

Unions. To allow `nil`, we use *ad-hoc unions* (`nil` and `false` are logically false).

```
(ann greet-nil [(U nil Str) -> Str])
(defn greet-nil [n] (str "Hello" (when n (str ", " n)) "!"))
(greet-nil "Donald") ;=> "Hello, Donald!" :- Str
(greet-nil nil)      ;=> "Hello!"          :- Str
```

Typed Clojure prevents well-typed code from dereferencing `nil`.

Flow Analysis. Occurrence typing [24] models type-based control flow. In `greetings`, a branch ensures `repeat` is never passed `nil`.

```
(ann greetings [Str (U nil Int) -> Str])
(defn greetings [n i]
  (str "Hello," (when i (apply str (repeat i "hello,")))) n "!")
(greetings "Donald" 2) ;=> "Hello, hello, hello, Donald!" :- Str
(greetings "Grace" nil) ;=> "Hello, Grace!"                :- Str
```

Removing the branch is a static type error—`repeat` cannot be passed `nil`.

```
(ann greetings-bad [Str (U nil Int) -> Str])
(defn greetings-bad [n i]
  ; Expected Int, given (U nil Int)
  (str "Hello," (apply str (repeat i "hello,")))) n "!")
```

2.3 Java Interoperability

Clojure can interact with Java constructors, methods, and fields. This program calls the `getParent` on a constructed `File` instance, returning a nullable string.

```
(.getParent (new File "a/b")) ;=> "a" :- (U nil Str) Example 1
```

Typed Clojure can integrate with the Clojure compiler to avoid expensive reflective calls like `getParent`, however if a specific overload cannot be found based on the surrounding static context, a type error is thrown.

```
(fn [f] (.getParent f)) ; Type Error: Unresolved interop: getParent
```

Function arguments default to `Any`, which is similar to a union of all types. Ascribing a parameter type allows Typed Clojure to find a specific method.

```
(ann parent [(U nil File) -> (U nil Str)]) Example 2  
(defn parent [f] (if f (.getParent f) nil))
```

The conditional guards from dereferencing `nil`, and—as before—removing it is a static type error, as typed code could possibly dereference `nil`.

```
(defn parent-bad-in [f :- (U nil File)]  
  (.getParent f)) ; Type Error: Cannot call instance method on nil.
```

Typed Clojure rejects programs that assume methods cannot return `nil`.

```
(defn parent-bad-out [f :- File] :- Str  
  (.getParent f)) ; Type Error: Expected Str, given (U nil Str).
```

Method targets can never be `nil`. Typed Clojure also prevents passing `nil` as Java method or constructor arguments by default—this restriction can be adjusted per method.

In contrast, JVM invariants guarantee constructors return non-null.³

```
(parent (new File s)) Example 3
```

2.4 Multimethods

Multimethods are a kind of extensible function—combining a *dispatch function* with one or more *methods*—widely used to define Clojure operations.

Value-based Dispatch. This simple multimethod takes a keyword (`Kw`) and says hello in different languages.

```
(ann hi [Kw -> Str]) ; multimethod type Example 4  
(defmulti hi identity) ; dispatch function 'identity'  
(defmethod hi :en [_] "hello") ; method for ':en'  
(defmethod hi :fr [_] "bonjour") ; method for ':fr'  
(defmethod hi :default [_] "um...") ; default method
```

When invoked, the arguments are first supplied to the dispatch function—`identity`—yielding a *dispatch value*. A method is then chosen based on the dispatch value, to which the arguments are then passed to return a value.

```
(map hi [:en :fr :bocce]) ;=> ("hello" "bonjour" "um...")
```

³ <http://docs.oracle.com/javase/specs/jls/se7/html/jls-15.html#jls-15.9.4>.

For example, `(hi :en)` evaluates to `"hello"`—it executes the `:en` method because `(= (identity :en) :en)` is true and `(= (identity :en) :fr)` is false.

Dispatching based on literal values enables certain forms of method definition, but this is only part of the story for multimethod dispatch.

Class-based Dispatch. For class values, multimethods can choose methods based on subclassing relationships. Recall the multimethod from Fig. 1. The dispatch function `class` dictates whether the `String` or `File` method is chosen. The multimethod dispatch rules use `isa?`, a hybrid predicate which is both a subclassing check for classes and an equality check for other values.

```
(isa? :en :en)           ;=> true
(isa? String Object) ;=> true
```

The current dispatch value and—in turn—each method’s associated dispatch value is supplied to `isa?`. If exactly one method returns true, it is chosen. For example, the call `(pname "STAINS/JELLY")` picks the `String` method because `(isa? String String)` is true, and `(isa? String File)` is not.

2.5 Heterogeneous Hash-Maps

The most common way to represent compound data in Clojure are immutable hash-maps, typically with keyword keys. Keywords double as functions that look themselves up in a map, or return `nil` if absent.

```
(def breakfast {:en "waffles" :fr "croissants"})
(:en breakfast) ;=> "waffles" :- Str
(:bocce breakfast) ;=> nil :- nil
```

Example 5

HMap types describe the most common usages of keyword-keyed maps.

```
breakfast; : - (HMap :mandatory {:en Str, :fr Str}, :complete? true)
```

This says `:en` and `:fr` are known entries mapped to strings, and the map is fully specified—that is, no other entries exist—by `:complete?` being true.

HMap types default to partial specification, with `'{:en Str :fr Str}` abbreviating `(HMap :mandatory {:en Str, :fr Str})`.

```
(ann lunch ' {:en Str :fr Str})
(def lunch {:en "muffin" :fr "baguette"})
(:bocce lunch) ;=> nil :- Any ; less accurate type
```

Example 6

HMaps in Practice. The next example is extracted from a production system at CircleCI, a company with a large production Typed Clojure system (Sect. 5.2 presents a case study and empirical result from this code base).

```
(defalias RawKeyPair ; extra keys disallowed
  (HMap :mandatory {:pub RawKey, :priv RawKey},
    :complete? true))
(defalias EncKeyPair ; extra keys disallowed
  (HMap :mandatory {:pub RawKey, :enc-priv EncKey}, :complete? true))

(ann enc-keypair [RawKeyPair -> EncKeyPair])
(defn enc-keypair [kp]
  (assoc (dissoc kp :priv) :enc-priv (encrypt (:priv kp))))
```

Example 7

As `EncKeyPair` is fully specified, we remove extra keys like `:priv` via `dissoc`, which returns a new map that is the first argument without the entry named by the second argument. Notice removing `dissoc` causes a type error.

```
(defn enc-keypair-bad [kp] ; Type error: :priv disallowed
  (assoc kp :enc-priv (encrypt (:priv kp))))
```

2.6 HMaps and Multimethods, Joined at the Hip

HMaps and multimethods are the primary ways for representing and dispatching on data respectively, and so are intrinsically linked. As type system designers, we must search for a compositional approach that can anticipate any combination of these features.

Thankfully, occurrence typing, originally designed for reasoning about `if` tests, provides the compositional approach we need. By extending the system with a handful of rules based on HMaps and other functions, we can automatically cover both easy cases and those that compose rules in arbitrary ways.

Futhermore, this approach extends to multimethod dispatch by reusing occurrence typing's approach to conditionals and encoding a small number of rules to handle the `isa?`-based dispatch. In practice, conditional-based control flow typing extends to multimethod dispatch, and vice-versa.

We first demonstrate a very common, simple dispatch style, then move on to deeper structural dispatching where occurrence typing's compositionality shines.

HMaps and Unions. Partially specified HMap's with a common dispatch key combine naturally with ad-hoc unions. An `Order` is one of three kinds of HMaps.

```
(defalias Order "A meal order, tracking dessert quantities."
  (U ' {:Meal ':lunch, :desserts Int} ' {:Meal ':dinner :desserts Int}
    ' {:Meal ':combo :meal1 Order :meal2 Order}))
```

The `:Meal` entry is common to each HMap, always mapped to a known keyword singleton type. It's natural to dispatch on the `class` of an instance—it's similarly natural to dispatch on a known entry like `:Meal`.

```
(ann desserts [Order -> Int])
(defmulti desserts :Meal) ; dispatch on :Meal entry
(defmethod desserts :lunch [o] (:desserts o))
(defmethod desserts :dinner [o] (:desserts o))
```

Example 8

```
(defmethod desserts :combo [o]
  (+ (desserts (:meal1 o)) (desserts (:meal2 o))))

(desserts {:Meal :combo, :meal1 {:Meal :lunch :desserts 1},
          :meal2 {:Meal :dinner :desserts 2}}) ;=> 3
```

The `:combo` method is verified to only structurally recur on `Orders`. This is achieved because we learn the argument `o` must be of type `'{:Meal :combo}` since `(isa? (:Meal o) :combo)` is true. Combining this with the fact that `o` is an `Order` eliminates possibility of `:lunch` and `:dinner` orders, simplifying `o` to `'{:Meal ':combo :meal1 Order :meal2 Order}` which contains appropriate arguments for both recursive calls.

Nested Dispatch. A more exotic dispatch mechanism for `desserts` might be on the `class` of the `:desserts` key. If the result is a number, then we know the `:desserts` key is a number, otherwise the input is a `:combo` meal. We have already seen dispatch on `class` and on keywords in isolation—occurrence typing automatically understands control flow that combines its simple building blocks.

The first method has dispatch value `Long`, a subtype of `Int`, and the second method has `nil`, the sentinel value for a failed map lookup. In practice, `:lunch` and `:dinner` meals will dispatch to the `Long` method, but Typed Clojure infers a slightly more general type due to the definition of `:combo` meals.

```
(ann desserts' [Order -> Int])
(defmulti desserts'
  (fn [o :- Order] (class (:desserts o))))
(defmethod desserts' Long [o]
; o :- (U '{:Meal (U ':dinner ':lunch), :desserts Int}
;
; '{:Meal ':combo, :desserts Int, :meal1 Order, :meal2 Order})
  (:desserts o))
(defmethod desserts' nil [o]
; o :- '{:Meal ':combo, :meal1 Order, :meal2 Order}
  (+ (desserts' (:meal1 o)) (desserts' (:meal2 o))))
```

Example 9

In the `Long` method, Typed Clojure learns that its argument is at least of type `'{:desserts Long}`—since `(isa? (class (:desserts o)) Long)` must be true. Here the `:desserts` entry *must* be present and mapped to a `Long`—even in a `:combo` meal, which does not specify `:desserts` as present or absent.

In the `nil` method, `(isa? (class (:desserts o)) nil)` must be true—which implies `(class (:desserts o))` is `nil`. Since lookups on missing keys return `nil`, either

- `o` has a `:desserts` entry to `nil`, like `:desserts nil`, or
- `o` is missing a `:desserts` entry.

We can express this type with the `:absent-keys` `HMap` option

```
(U '{:desserts nil} (HMap :absent-keys #{:desserts}))
```


This eliminates non-`:combo` meals since their `'{:desserts Int}` type does not agree with this new information (because `:desserts` is neither `nil` or `absent`).

From Multiple to Arbitrary Dispatch. Clojure multimethod dispatch, and Typed Clojure’s handling of it, goes even further, supporting dispatch on multiple arguments via vectors. Dispatch on multiple arguments is beyond the scope of this paper, but the same intuition applies—adding support for multiple dispatch admits arbitrary combinations and nestings of it and previous dispatch rules.

3 A Formal Model of λ_{TC}

After demonstrating the core features of Typed Clojure, we link them together in a formal model called λ_{TC} . Building on occurrence typing, we incrementally add each novel feature of Typed Clojure to the formalism, interleaving presentation of syntax, typing rules, operational semantics, and subtyping.

3.1 Core Type System

We start with a review of occurrence typing [24], the foundation of λ_{TC} .

Expressions. Syntax is given in Fig. 2. Expressions e include variables x , values v , applications, abstractions, conditionals, and let expressions. All binding forms introduce fresh variables—a subtle but important point since our type environments are not simply dictionaries. Values include booleans b , `nil`, class literals C , keywords k , integers n , constants c , and strings s . Lexical closures $[\rho, \lambda x^\tau. e]_c$ close value environments ρ —which map bindings to values—over functions.

Types. Types σ or τ include the top type \top , *untagged* unions ($\bigcup \vec{\tau}$), singletons (**Val** l), and class instances C . We abbreviate the classes **Boolean** to **B**, **Keyword** to **K**, **Nat** to **N**, **String** to **S**, and **File** to **F**. We also abbreviate the types (\bigcup) to \perp , (**Val** `nil`) to **nil**, (**Val** `true`) to **true**, and (**Val** `false`) to **false**. The difference between the types (**Val** C) and C is subtle. The former is inhabited by class literals like **K** and the result of (`class :a`)—the latter by *instances* of classes, like a keyword literal `:a`, an instance of the type **K**. Function types $x:\sigma \xrightarrow[o]{\psi|\psi} \tau$ contain *latent* (terminology from [17]) propositions ψ , object o , and return type τ , which may refer to the function argument x . They are instantiated with the actual object of the argument in applications.

Objects. Each expression is associated with a symbolic representation called an *object*. For example, variable m has object m ; (**class** `(:lunch m)`) has object **class**(**key**_{·lunch}(m)); and `42` has the *empty* object \emptyset since it is unimportant in our system. Figure 2 gives the syntax for objects o —non-empty objects $\pi(x)$ combine of a root variable x and a *path* π , which consists of a possibly-empty sequence of *path elements* (pe) applied right-to-left from the root variable. We use two path elements—**class** and **key** _{k} —representing the results of calling *class* and looking up a keyword k , respectively.

$e ::= x \mid v \mid (e e) \mid \lambda x^\tau. e \mid (\text{if } e e e) \mid (\text{let } [x e] e)$	Expressions
$v ::= l \mid n \mid c \mid s \mid [\rho, \lambda x^\tau. e]_c$	Values
$c ::= \text{class} \mid n?$	Constants
$\sigma, \tau ::= \top \mid (\bigcup \vec{\tau}) \mid x : \tau \xrightarrow{\psi \psi} \tau \mid (\mathbf{Val} l) \mid C$	Types
$l ::= k \mid C \mid \text{nil} \mid b$	Value types
$b ::= \text{true} \mid \text{false}$	Boolean values
$\psi ::= \tau_{\pi(x)} \mid \bar{\tau}_{\pi(x)} \mid \psi \supset \psi \mid \psi \wedge \psi \mid \psi \vee \psi \mid \text{tt} \mid \text{ff}$	Propositions
$o ::= \pi(x) \mid \emptyset$	Objects
$\pi ::= \vec{p}\bar{e}$	Paths
$pe ::= \text{class} \mid \text{key}_k$	Path elements
$\Gamma ::= \vec{\psi}$	Proposition environments
$\rho ::= \{x \mapsto \vec{v}\}$	Value environments

Fig. 2. Syntax of terms, types, propositions and objects

Propositions with a Logical System. In standard type systems, association lists often track the types of variables, like in LC-Let and LC-Local.

$$\frac{\text{LC-LET} \quad \Gamma \vdash e_1 : \sigma \quad \Gamma, x \mapsto \sigma \vdash e_2 : \tau}{\Gamma \vdash (\text{let } [x e_1] e_2) : \tau} \quad \text{LC-LOCAL} \quad \frac{\Gamma(x) = \tau}{\Gamma \vdash x : \tau}$$

Occurrence typing instead pairs *logical formulas*, that can reason about arbitrary non-empty objects, with a *proof system*. The logical statement σ_x says variable x is of type σ .

$$\frac{\text{T0-LET} \quad \Gamma \vdash e_1 : \sigma \quad \Gamma, \sigma_x \vdash e_2 : \tau}{\Gamma \vdash (\text{let } [x e_1] e_2) : \tau} \quad \text{T0-LOCAL} \quad \frac{\Gamma \vdash \tau_x}{\Gamma \vdash x : \tau}$$

In T0-Local, $\Gamma \vdash \tau_x$ appeals to the proof system to solve for τ .

We further extend logical statements to *propositional logic*. Figure 2 describes the syntax for propositions ψ , consisting of positive and negative *type propositions* about non-empty objects— $\tau_{\pi(x)}$ and $\bar{\tau}_{\pi(x)}$ respectively—the latter pronounced “the object $\pi(x)$ is *not* of type τ ”. The other propositions are standard logical connectives: implications, conjunctions, disjunctions, and the trivial (tt) and impossible (ff) propositions. The full proof system judgement $\Gamma \vdash \psi$ says *proposition environment* Γ proves proposition ψ .

Each expression is associated with two propositions—when expression e_1 is in test position like $(\text{if } e_1 e_2 e_3)$, the type system extracts e_1 ’s ‘then’ and ‘else’ proposition to check e_2 and e_3 respectively. For example, in $(\text{if } o e_2 e_3)$ we learn variable o is true in e_2 via o ’s ‘then’ proposition $(\bigcup \text{nil false})_o$, and that o is false in e_3 via o ’s ‘else’ proposition $(\bigcup \text{nil false})_o$.

To illustrate, recall Example 8. The parameter o is of type **Order**, written Order_o as a proposition. In the `:combo` method, we know $(:\text{Meal } o) : \text{combo}$,

	T-ABS	T-IF
T-LOCAL	$\frac{\Gamma \vdash \tau_x \quad \Gamma, \sigma_x \vdash e \Rightarrow e' : \sigma' ; \psi_+ \psi_- ; o}{\sigma = (\bigcup \text{nil false}) \quad \tau = x : \sigma \xrightarrow{\psi_+ \psi_-} \sigma'}}$	$\frac{\Gamma \vdash e_1 \Rightarrow e'_1 : \tau_1 ; \psi_{1+} \psi_{1-} ; o_1 \quad \Gamma, \psi_{1+} \vdash e_2 \Rightarrow e'_2 : \tau ; \psi_+ \psi_- ; o \quad \Gamma, \psi_{1-} \vdash e_3 \Rightarrow e'_3 : \tau ; \psi_+ \psi_- ; o \quad e' = (\text{if } e'_1 \ e'_2 \ e'_3)}{\Gamma \vdash (\text{if } e_1 \ e_2 \ e_3) \Rightarrow e' : \tau ; \psi_+ \psi_- ; o}}$
$\Gamma \vdash x : \tau ; \bar{\sigma}_x \sigma_x ; x \quad \Gamma \vdash \lambda x^\sigma . e \Rightarrow \lambda x^\sigma . e' : \tau ; \text{tt} \text{ff} ; \emptyset \quad \Gamma \vdash (\text{if } e_1 \ e_2 \ e_3) \Rightarrow e' : \tau ; \psi_+ \psi_- ; o$		
T-KW	T-NIL	T-STR
$\Gamma \vdash k : (\mathbf{Val} \ k) ; \text{tt} \text{ff} ; \emptyset$	$\Gamma \vdash \text{nil} : \mathbf{nil} ; \text{ff} \text{tt} ; \emptyset$	$\Gamma \vdash s : \mathbf{S} ; \text{tt} \text{ff} ; \emptyset$
T-NUM	T-FALSE	T-CLASS
$\Gamma \vdash n : \mathbf{N} ; \text{tt} \text{ff} ; \emptyset$	$\Gamma \vdash \text{false} : \mathbf{false} ; \text{ff} \text{tt} ; \emptyset$	$\Gamma \vdash C : (\mathbf{Val} \ C) ; \text{tt} \text{ff} ; \emptyset$
	T-CONST	T-TRUE
	$\Gamma \vdash c : \delta_r(c) ; \text{tt} \text{ff} ; \emptyset$	$\Gamma \vdash \text{true} : \mathbf{true} ; \text{tt} \text{ff} ; \emptyset$
$\text{T-LET} \quad \frac{\Gamma \vdash e_1 \Rightarrow e'_1 : \sigma ; \psi_{1+} \psi_{1-} ; o_1 \quad \psi' = (\bigcup \text{nil false})_x \supset \psi_{1+} \quad \psi'' = (\bigcup \text{nil false})_x \supset \psi_{1-} \quad \Gamma, \sigma_x, \psi', \psi'' \vdash e_2 \Rightarrow e'_2 : \tau ; \psi_+ \psi_- ; o}{\Gamma \vdash (\text{let } [x \ e_1] \ e_2) \Rightarrow (\text{let } [x \ e'_1] \ e'_2) : \tau[o_1/x] ; \psi_+ \psi_-[o_1/x] ; o[o_1/x]}}$		
T-APP		T-SUBSUME
$\Gamma \vdash e \Rightarrow e_1 : x : \sigma \xrightarrow{\psi_{f+} \psi_{f-}} \tau ; \psi_+ \psi_- ; o$		$\Gamma \vdash e \Rightarrow e' : \tau ; \psi_+ \psi_- ; o$
$\Gamma \vdash e' \Rightarrow e'_1 : \sigma ; \psi'_{+} \psi'_{-} ; o'$		$\Gamma, \psi_+ \vdash \psi'_{+} \quad \Gamma, \psi_- \vdash \psi'_{-}$
$\Gamma \vdash (e \ e') \Rightarrow (e_1 \ e'_1) : \tau[o'/x] ; \psi_{f+} \psi_{f-}[o'/x] ; o_f[o'/x]$		$\vdash \tau <: \tau' \quad \vdash o <: o'$
$\Gamma \vdash (e \ e') \Rightarrow (e_1 \ e'_1) : \tau[o'/x] ; \psi_{f+} \psi_{f-}[o'/x] ; o_f[o'/x]$		$\Gamma \vdash e \Rightarrow e' : \tau' ; \psi'_{+} \psi'_{-} ; o'$

Fig. 3. Core typing rules

based on multimethod dispatch rules. This is written $(\mathbf{Val} : \text{combo})_{\text{key} : \text{Meal}(o)}$, pronounced “the :Meal path of variable o is of type $(\mathbf{Val} : \text{combo})$ ”.

To attain the type of o , we must solve for τ in $\Gamma \vdash \tau_o$, under proposition environment $\Gamma = \mathbf{Order}_o$, $(\mathbf{Val} : \text{combo})_{\text{key} : \text{Meal}(o)}$ which deduces τ to be a :combo meal. The logical system *combines* pieces of type information to deduce more accurate types for lexical bindings—this is explained in Sect. 3.6.

Typing Judgment. We formalize our system following Tobin-Hochstadt and Felleisen [24]. The typing judgment $\Gamma \vdash e \Rightarrow e' : \tau ; \psi_+ | \psi_- ; o$ says expression e rewrites to e' , which is of type τ in the proposition environment Γ , with ‘then’ proposition ψ_+ , ‘else’ proposition ψ_- and object o .

We write $\Gamma \vdash e \Rightarrow e' : \tau$ to mean $\Gamma \vdash e \Rightarrow e' : \tau ; \psi'_{+} | \psi'_{-} ; o'$ for some ψ'_{+} , ψ'_{-} and o' , and abbreviate self rewriting judgements $\Gamma \vdash e \Rightarrow e : \tau ; \psi_+ | \psi_- ; o$ to $\Gamma \vdash e : \tau ; \psi_+ | \psi_- ; o$.

Typing Rules. The core typing rules are given as Fig. 3. We introduce the interesting rules with the complement number predicate as a running example.

$$\lambda d^\top . (\text{if } (n? \ d) \ \text{false} \ \text{true}) \quad (1)$$

The lambda rule T-Abs introduces $\sigma_x = \top_d$ to check the body. With $\Gamma = \top_d$, T-If first checks the test $e_1 = (n? \ d)$ via the T-App rule, with three steps.

$$\begin{array}{c}
 \text{S-UNIONSUPER} \quad \text{S-UNIONSUB} \quad \text{S-FUNMONO} \quad \text{S-OBJECT} \\
 \frac{\exists i. \vdash \tau <: \sigma_i}{\vdash \tau <: (\bigcup \overline{\sigma^i})} \quad \frac{\text{S-UNIONSUB}}{\vdash \tau_i <: \sigma} \quad \frac{\text{S-FUNMONO}}{\vdash x : \sigma \xrightarrow[o]{\psi_+ | \psi_-} \tau <: \mathbf{Fn}} \quad \begin{array}{l} \vdash C <: \mathbf{Object} \\ \text{S-SCCLASS} \\ \vdash (\mathbf{Val } C) <: \mathbf{Class} \\ \text{S-SBOOL} \\ \vdash (\mathbf{Val } b) <: \mathbf{B} \end{array} \\
 \\
 \text{S-FUN} \quad \text{S-REFL} \quad \text{S-TOP} \\
 \frac{\vdash \sigma' <: \sigma \quad \vdash \tau <: \tau' \quad \psi_+ \vdash \psi'_+ \quad \psi_- \vdash \psi'_- \quad \vdash o <: o'}{\vdash x : \sigma \xrightarrow[o]{\psi_+ | \psi_-} \tau <: x : \sigma' \quad \frac{\psi'_+ | \psi'_-}{o'} \rightarrow \tau'} \quad \frac{\text{S-REFL} \quad \text{S-TOP}}{\vdash \tau <: \tau \quad \vdash \tau <: \top} \\
 \text{S-SKW} \\
 \vdash (\mathbf{Val } k) <: \mathbf{K}
 \end{array}$$

Fig. 4. Core subtyping rules

$$\begin{array}{c}
 \text{B-IFTRUE} \quad \text{B-IFFALSE} \\
 \frac{\rho \vdash e_1 \Downarrow v_1 \quad v_1 \neq \mathbf{false} \quad v_1 \neq \mathbf{nil} \quad \rho \vdash e_2 \Downarrow v}{\rho \vdash (\text{if } e_1 \ e_2 \ e_3) \Downarrow v} \quad \frac{\rho \vdash e_1 \Downarrow \mathbf{false} \text{ or } \rho \vdash e_1 \Downarrow \mathbf{nil} \quad \rho \vdash e_3 \Downarrow v}{\rho \vdash (\text{if } e_1 \ e_2 \ e_3) \Downarrow v}
 \end{array}$$

Fig. 5. Select core semantics

First, in T-App the operator $e = n?$ is checked with T-Const, which uses δ_τ (Fig. 7, dynamic semantics in the supplemental material) to type constants. $n?$ is a predicate over numbers, and *class* returns its argument’s class.

Resuming $(n? d)$, in T-App the operand $e' = d$ is checked with T-Local as

$$\Gamma \vdash d : \top ; \overline{(\bigcup \mathbf{nil } \mathbf{false})}_d | (\bigcup \mathbf{nil } \mathbf{false})_d ; d \quad (2)$$

which encodes the type, proposition, and object information about variables. The proposition $\overline{(\bigcup \mathbf{nil } \mathbf{false})}_d$ says “it is not the case that variable d is of type $(\bigcup \mathbf{nil } \mathbf{false})$ ”; $(\bigcup \mathbf{nil } \mathbf{false})_d$ says “ d is of type $(\bigcup \mathbf{nil } \mathbf{false})$ ”.

Finally, the T-App rule substitutes the operand’s object o' for the parameter x in the latent type, propositions, and object. The proposition \mathbf{N}_d says “ d is of type \mathbf{N} ”; $\overline{\mathbf{N}}_d$ says “it is not the case that d is of type \mathbf{N} ”. The object d is the symbolic representation of what the expression d evaluates to.

$$\Gamma \vdash (n? d) : \mathbf{B} ; \mathbf{N}_d | \overline{\mathbf{N}}_d ; \emptyset \quad (3)$$

To demonstrate, the ‘then’ proposition—in T-App $\psi_+[o'/x]$ —substitutes the latent ‘then’ proposition of $\delta_\tau(n?)$ with d , giving $\mathbf{N}_x[d/x] = \mathbf{N}_d$.

To check the branches of $(\text{if } (n? d) \ \mathbf{false} \ \mathbf{true})$, T-If introduces $\psi_{1+} = \mathbf{N}_d$ to check $e_2 = \mathbf{false}$, and $\psi_{1-} = \overline{\mathbf{N}}_d$ to check $e_3 = \mathbf{true}$. The branches are first checked with T-False and T-True respectively, the T-Subsume premises $\Gamma, \psi_+ \vdash \psi'_+$ and $\Gamma, \psi_- \vdash \psi'_-$ allow us to pick compatible propositions for both branches.

$$\begin{array}{l}
 \Gamma, \mathbf{N}_d \vdash \mathbf{false} : \mathbf{B} ; \overline{\mathbf{N}}_d | \mathbf{N}_d ; \emptyset \\
 \Gamma, \overline{\mathbf{N}}_d \vdash \mathbf{true} : \mathbf{B} ; \overline{\mathbf{N}}_d | \mathbf{N}_d ; \emptyset
 \end{array}$$

Finally T-Abs assigns a type to the overall function:

$$\vdash \lambda d^\top. (\text{if } (n? d) \text{ false true}) : d : \top \xrightarrow[\emptyset]{\overline{\mathbf{N}}_d | \mathbf{N}_d} \mathbf{B} ; \text{tt|fff} ; \emptyset$$

Subtyping. Figure 4 presents subtyping as a reflexive and transitive relation with top type \top . Singleton types are instances of their respective classes—boolean singleton types are of type \mathbf{B} , class literals are instances of \mathbf{Class} and keywords are instances of \mathbf{K} . Instances of classes C are subtypes of \mathbf{Object} . Function types are subtypes of \mathbf{Fn} . All types except for \mathbf{nil} are subtypes of \mathbf{Object} , so \top is similar to $(\bigcup \mathbf{nil} \mathbf{Object})$. Function subtyping is contravariant left of the arrow—latent propositions, object and result type are covariant. Subtyping for untagged unions is standard.

Operational Semantics. We define the dynamic semantics for λ_{TC} in a big-step style using an environment, following [24]. We include both errors and a *wrong* value, which is provably ruled out by the type system. The main judgment is $\rho \vdash e \Downarrow \alpha$ which states that e evaluates to answer α in environment ρ . We chose to omit the core rules (included in supplemental material) however a notable difference is \mathbf{nil} is a false value, which affects the semantics of if (Fig. 5).

3.2 Java Interoperability

We present Java interoperability in a restricted setting without class inheritance, overloading or Java Generics. We extend the syntax in Fig. 6 with Java field lookups and calls to methods and constructors. To prevent ambiguity between zero-argument methods and fields, we use Clojure’s primitive “dot” syntax: field accesses are written $(. e fld)$ and method calls $(. e (mth \vec{e}))$.

In Example 1, $(. \text{getParent } (\text{new File } "a/b"))$ translates to

$$(. (\text{new } \mathbf{F} \text{ "a/b"}) (\text{getParent})) \quad (4)$$

But both the constructor and method are unresolved. We introduce *non-reflective* expressions for specifying exact Java overloads.

$$(. (\text{new}_{[\mathbf{S}]} \mathbf{F} \text{ "a/b"}) (\text{getParent}_{[\mathbf{S}]}^{\mathbf{F}})) \quad (5)$$

From the left, the one-argument constructor for \mathbf{F} takes a \mathbf{S} , and the getParent method of \mathbf{F} takes zero arguments and returns a \mathbf{S} .

We now walk through this conversion.

Constructors. First we check and convert $(\text{new } \mathbf{F} \text{ "a/b"})$ to $(\text{new}_{[\mathbf{S}]} \mathbf{F} \text{ "a/b"})$. The T-New typing rule checks and rewrites constructors. To check $(\text{new } \mathbf{F} \text{ "a/b"})$ we first resolve the constructor overload in the class table—there is at most one to simplify presentation. With $C_1 = \mathbf{S}$, we convert to a nilable type the argument with $\tau_1 = (\bigcup \mathbf{nil} \mathbf{S})$ and type check “a/b” against τ_1 . Typed Clojure defaults to allowing non-nilable arguments, but this can be overridden, so we model the more

$ \begin{aligned} e &::= \dots (. e fld) \mid (. e (mth \vec{e})) \mid (new C \vec{e}) \\ &\quad \mid (. e fld_C^C) \mid (. e (mth_{[[\vec{C}], C]}^C \vec{e})) \mid (new_{[\vec{C}]} C \vec{e}) \\ v &::= \dots \mid C \{fld : v\} \\ ce &::= \{m \mapsto \{mth \mapsto [[\vec{C}], C]\}, f \mapsto \{fld \mapsto \vec{C}\}, c \mapsto \{[\vec{C}]\}\} \\ \mathcal{CT} &::= \{C \mapsto ce\} \end{aligned} $	Expressions Non-reflective Expressions Values Class descriptors Class Table
---	---

$$\text{T-NEW} \quad \frac{\overline{[\vec{C}_i] \in \mathcal{CT}[C][c]} \quad \overline{JT_{\text{nil}}(C_i) = \tau_i} \quad \overline{\Gamma \vdash e_i \Rightarrow e'_i : \tau_i} \quad \overline{JT(C) = \tau}}{\Gamma \vdash (new C \vec{e}_i) \Rightarrow (new_{[\vec{C}_i]} C \vec{e}'_i) : \tau ; \text{\texttt{}} \text{\texttt{}} \text{\texttt{}} ; \emptyset}$$

$$\text{T-METHOD} \quad \frac{\overline{\Gamma \vdash e \Rightarrow e' : \sigma} \quad \overline{TJ(\sigma) = C_1} \quad \overline{mth \mapsto [[\vec{C}_i], C_2] \in \mathcal{CT}[C_1][m]} \quad \overline{JT_{\text{nil}}(C_i) = \tau_i} \quad \overline{\Gamma \vdash e_i \Rightarrow e'_i : \tau_i} \quad \overline{JT_{\text{nil}}(C_2) = \tau} \quad \vdash \sigma <: \mathbf{Object}}{\Gamma \vdash (. e (mth \vec{e}_i)) \Rightarrow (. e' (mth_{[[\vec{C}_i], C_2]}^{C_1} \vec{e}'_i)) : \tau ; \text{\texttt{}} \text{\texttt{}} \text{\texttt{}} ; \emptyset}$$

$$\text{T-FIELD} \quad \frac{\overline{\Gamma \vdash e \Rightarrow e' : \sigma} \quad \vdash \sigma <: \mathbf{Object} \quad \overline{TJ(\sigma) = C_1} \quad \overline{fld \mapsto C_2 \in \mathcal{CT}[C_1][f]} \quad \overline{JT_{\text{nil}}(C_2) = \tau}}{\Gamma \vdash (. e fld) \Rightarrow (. e' fld_{C_2}^{C_1}) : \tau ; \text{\texttt{}} \text{\texttt{}} \text{\texttt{}} ; \emptyset}$$

$$\begin{array}{llll}
 JT_{\text{nil}}(\mathbf{Void}) = \mathbf{nil} & JT(\mathbf{Void}) = \mathbf{nil} & TJ(\tau) = C & \text{if } \vdash \tau <: JT_{\text{nil}}(C) \\
 JT_{\text{nil}}(C) = (\bigcup \mathbf{nil} C) & JT(C) = C & &
 \end{array}$$

$$\text{B-FIELD} \quad \frac{\rho \vdash e \Downarrow v \quad \overline{JVM_{\text{getstatic}}[C_1, v_1, fld, C_2] = v} \quad \overline{\rho \vdash e_i \Downarrow v_i} \quad \overline{JVM_{\text{new}}[C_1, [\vec{C}_i], [\vec{v}_i]] = v}}{\rho \vdash (. e fld_{C_2}^{C_1}) \Downarrow v} \quad \overline{\rho \vdash (new_{[\vec{C}_i]} C \vec{e}_i) \Downarrow v}$$

$$\text{B-METHOD} \quad \frac{\rho \vdash e_m \Downarrow v_m \quad \overline{\rho \vdash e_a \Downarrow v_a} \quad \overline{JVM_{\text{invokestatic}}[C_1, v_m, mth, [\vec{C}_a], [\vec{v}_a], C_2] = v}}{\rho \vdash (. e_m (mth_{[[\vec{C}_a], C_2]}^{C_1} \vec{e}_a)) \Downarrow v}$$

Fig. 6. Java interoperability syntax, typing and operational semantics

$$\begin{aligned}
 \delta_\tau(\text{class}) &= x : \top \xrightarrow[\text{class}(x)]{\text{\texttt{}} \text{\texttt{}} \text{\texttt{}}} (\bigcup \mathbf{nil} \mathbf{Class}) \\
 \delta_\tau(n?) &= x : \top \xrightarrow[\emptyset]{\mathbf{N}_x | \overline{\mathbf{N}}_x} \mathbf{B}
 \end{aligned}$$

Fig. 7. Constant typing

general case. The return Java type \mathbf{F} is converted to a non-nil Typed Clojure type $\tau = \mathbf{F}$ for the return type, and the propositions say constructors can never be false—constructors can never produce the internal boolean value that Clojure uses for false, or nil. Finally, the constructor rewrites to $(new_{[S]} \mathbf{F} \text{ “a/b”})$.

Methods. Next we convert $(. (new_{[S]} \mathbf{F} \text{ “a/b”}) (\text{getParent}))$ to the non-reflective expression $(. (new_{[S]} \mathbf{F} \text{ “a/b”}) (\text{getParent}_{[[\emptyset, S]]}^{\mathbf{F}}))$. The T-Method rule for unresolved methods checks $(. (new_{[S]} \mathbf{F} \text{ “a/b”}) (\text{getParent}))$. We verify the target

type $\sigma = \mathbf{F}$ is non-nil by T-New. The overload is chosen from the class table based on $C_1 = \mathbf{F}$ —there is at most one. The nilable return type $\tau = (\bigcup \mathbf{nil} \mathbf{S})$ is given, and the entire expression rewrites to expression 5.

The T-Field rule (Fig. 6) is like T-Method, but without arguments.

The evaluation rules B-Field, B-New and B-Method (Fig. 6) simply evaluate their arguments and call the relevant JVM operation, which we do not model—Sect. 4 states our exact assumptions. There are no evaluation rules for reflective Java interoperability, since there are no typing rules that rewrite to reflective calls.

3.3 Multimethod Preliminaries: isa?

We now consider the `isa?` operation, a core part of the multimethod dispatch mechanism. Recalling the examples in Sect. 2.4, `isa?` is a subclassing test for classes, but otherwise is an equality test. The T-IsA rule uses `IsAProps` (Fig. 8), a metafunction which produces the propositions for `isa?` expressions.

To demonstrate the first `IsAProps` case, the expression `(isa? (class x) K)` is true if x is a keyword, otherwise false. When checked with T-IsA, the object of the left subexpression $o = \mathbf{class}(x)$ (which starts with the `class` path element) and the type of the right subexpression $\tau = (\mathbf{Val} \mathbf{K})$ (a singleton class type) together trigger the first `IsAProps` case $\mathbf{IsAProps}(\mathbf{class}(x), (\mathbf{Val} \mathbf{K})) = \mathbf{K}_x | \overline{\mathbf{K}}_x$, giving propositions that correspond to our informal description $\psi_+ | \psi_- = \mathbf{K}_x | \overline{\mathbf{K}}_x$.

The second `IsAProps` case captures the simple equality mode for non-class singleton types. For example, the expression `(isa? x :en)` produces true when x evaluates to `:en`, otherwise it produces false. Using T-IsA, it has the propositions $\psi_+ | \psi_- = \mathbf{IsAProps}(x, (\mathbf{Val} :en)) = (\mathbf{Val} :en)_x | (\overline{\mathbf{Val} :en})_x$ since $o = x$ and $\tau = (\mathbf{Val} :en)$. The side condition on the second `IsAProps` case ensures we are in equality mode—if x can possibly be a class in `(isa? x Object)`, `IsAProps` uses its conservative default case, since if x is a class literal, subclassing mode could be triggered. Capture-avoiding substitution of objects $[o/x]$ used in this case erases propositions that would otherwise have \emptyset substituted in for their objects—it is defined in the appendix.

The operational behavior of `isa?` is given by B-IsA (Fig. 8). `IsA` explicitly handles classes in the second case.

3.4 Multimethods

Figure 8 presents *immutable* multimethods without default methods to ease presentation. Figure 9 translates the mutable Example 4 to λ_{TC} .

To check `(defmulti x:K \rightarrow S $\lambda x^{\mathbf{K}}.x$)`, we note `(defmulti σ e)` creates a multimethod with *interface type* σ , and dispatch function e of type σ' , producing a value of type `(Multi σ σ')`. The T-DefMulti typing rule checks the dispatch function, and verifies both the interface and dispatch type's domain agree. Our example checks with $\tau = \mathbf{K}$, interface type $\sigma = x:\mathbf{K} \rightarrow \mathbf{S}$, dispatch function type $\sigma' = x:\mathbf{K} \xrightarrow{\text{tt|tt}} \mathbf{K}$, and overall type `(Multi $x:\mathbf{K} \rightarrow \mathbf{S} x:\mathbf{K} \xrightarrow{\text{tt|tt}} \mathbf{K}$)`.

e	::= ... (defmulti τ e) (defmethod e e e) (isa? e e)	Expressions
v	::= ... $[v, t]_m$	Values
t	::= $\{\overrightarrow{v \mapsto \tilde{v}}\}$	Dispatch tables
σ, τ	::= ... (Multi τ τ)	Types

$$\frac{\text{T-DEFMULTI} \quad \sigma = x:\tau \xrightarrow{o}^{\psi_+|\psi_-} \tau' \quad \sigma' = x:\tau \xrightarrow{o'}^{\psi'_+|\psi'_-} \tau'' \quad \Gamma \vdash e \Rightarrow e' : \sigma'}{\Gamma \vdash (\text{defmulti } \sigma \ e) \Rightarrow (\text{defmulti } \sigma \ e') : (\mathbf{Multi} \ \sigma \ \sigma') ; \text{tt}|\text{ff} ; \emptyset}$$

T-DEFMETHOD

$$\frac{\begin{array}{c} \tau_m = x:\tau \xrightarrow{o}^{\psi_+|\psi_-} \sigma \quad \tau_d = x:\tau \xrightarrow{o'}^{\psi'_+|\psi'_-} \sigma' \quad \Gamma \vdash e_m \Rightarrow e'_m : (\mathbf{Multi} \ \tau_m \ \tau_d) \\ \Gamma \vdash e_v \Rightarrow e'_v : \tau_v \quad \text{IsAProps}(o', \tau_v) = \psi''_+|\psi''_- \\ \Gamma, \tau_x, \psi''_+ \vdash e_b \Rightarrow e'_b : \sigma ; \psi_+|\psi_- ; o \quad e' = (\text{defmethod } e'_m \ e'_v \ \lambda x^\tau. e'_b) \end{array}}{\Gamma \vdash (\text{defmethod } e_m \ e_v \ \lambda x^\tau. e_b) \Rightarrow e' : (\mathbf{Multi} \ \tau_m \ \tau_d) ; \text{tt}|\text{ff} ; \emptyset}$$

T-ISA

$$\frac{\Gamma \vdash e \Rightarrow e_1 : \sigma ; \psi'_+|\psi'_- ; o \quad \Gamma \vdash e' \Rightarrow e'_1 : \tau \quad \text{IsAProps}(o, \tau) = \psi_+|\psi_-}{\Gamma \vdash (\text{isa? } e \ e') \Rightarrow (\text{isa? } e_1 \ e'_1) : \mathbf{B} ; \psi_+|\psi_- ; \emptyset}$$

$$\begin{array}{ll} \text{IsAProps}(\text{class}(\pi(x)), (\mathbf{Val} \ C)) = C_{\pi(x)}|\overline{C}_{\pi(x)} & \\ \text{IsAProps}(o, (\mathbf{Val} \ l)) = ((\mathbf{Val} \ l)_x|\overline{(\mathbf{Val} \ l)_x})[o/x] \text{ if } l \neq C & \\ \text{IsAProps}(o, \tau) = \text{tt}|\text{tt} \text{ otherwise} & \end{array}$$

S-PMULTIFN

$$\frac{\vdash \sigma_t <: x:\sigma \xrightarrow{o}^{\psi_+|\psi_-} \tau \quad \vdash \sigma_d <: x:\sigma \xrightarrow{o'}^{\psi'_+|\psi'_-} \tau'}{\vdash (\mathbf{Multi} \ \sigma_t \ \sigma_d) <: x:\sigma \xrightarrow{o}^{\psi_+|\psi_-} \tau}$$

S-PMULTI

$$\frac{\vdash \sigma <: \sigma' \quad \vdash \tau <: \tau'}{\vdash (\mathbf{Multi} \ \sigma \ \tau) <: (\mathbf{Multi} \ \sigma' \ \tau')}$$

S-MULTIMONO

$$\vdash (\mathbf{Multi} \ x:\sigma \xrightarrow{o}^{\psi_+|\psi_-} \tau \ x:\sigma \xrightarrow{o'}^{\psi'_+|\psi'_-} \tau') <: \mathbf{Multi}$$

B-DEFMULTI

$$\frac{\rho \vdash e \Downarrow v_d \quad v = [v_d, \{\}]_m}{\rho \vdash (\text{defmulti } \tau \ e) \Downarrow v}$$

B-DEFMETHOD

$$\frac{\rho \vdash e \Downarrow [v_d, t]_m \quad \rho \vdash e' \Downarrow v_v \quad \rho \vdash e_f \Downarrow v_f \quad v = [v_d, t[v_v \mapsto v_f]]_m}{\rho \vdash (\text{defmethod } e \ e' \ e_f) \Downarrow v}$$

GM(t, v_e) = v_f if $\overrightarrow{v_f s} = \{v_f\}$ where $\overrightarrow{v_f s} = \{v_f|v_k \mapsto v_f \in t \text{ and } \text{IsA}(v_e, v_k) = \text{true}\}$
 GM(t, v_e) = err otherwise

B-ISA

$$\frac{\rho \vdash e_1 \Downarrow v_1 \quad \rho \vdash e_2 \Downarrow v_2 \quad \text{IsA}(v_1, v_2) = v \quad \begin{array}{l} \text{IsA}(v, v) = \text{true} \quad v \neq C \\ \text{IsA}(C, C') = \text{true} \vdash C <: C' \\ \text{IsA}(v, v') = \text{false} \text{ otherwise} \end{array}}{\rho \vdash (\text{isa? } e_1 \ e_2) \Downarrow v}$$

B-BETAMULTI

$$\frac{\rho \vdash e \Downarrow [v_d, t]_m \quad \rho \vdash e' \Downarrow v' \quad \rho \vdash (v_d \ v') \Downarrow v_e \quad \text{GM}(t, v_e) = v_f \quad \rho \vdash (v_f \ v') \Downarrow v}{\rho \vdash (e \ e') \Downarrow v}$$

Fig. 8. Multimethod syntax, typing and operational semantics

$$\begin{aligned}
& (\text{let } [hi_0 \text{ (defmulti } x : \mathbf{K} \xrightarrow[\emptyset]{\text{tt}|\text{tt}} \mathbf{S} \lambda x^{\mathbf{K}}.x)] \\
& \quad (\text{let } [hi_1 \text{ (defmethod } hi_0 : \text{en } \lambda x^{\mathbf{K}}.\text{“hello”})] \\
& \quad \quad (\text{let } [hi_2 \text{ (defmethod } hi_1 : \text{fr } \lambda x^{\mathbf{K}}.\text{“bonjour”})] \\
& \quad \quad \quad (hi_2 : \text{en}))))
\end{aligned}$$

Fig. 9. Multimethod example

Next, we show how to check $(\text{defmethod } hi_0 : \text{en } \lambda x^{\mathbf{K}}.\text{“hello”})$. The expression $(\text{defmethod } e_m \ e_v \ e_f)$ creates a new multimethod that extends multimethod e_m 's dispatch table, mapping dispatch value e_v to method e_f . The T-DefMulti typing rule checks e_m is a multimethod with dispatch function type τ_d , then calculates the extra information we know based on the current dispatch value ψ''_+ , which is assumed when checking the method body. Our example checks with e_m being of type $(\text{Multi } x : \mathbf{K} \rightarrow \mathbf{S} \ x : \mathbf{K} \xrightarrow[x]{\text{tt}|\text{tt}} \mathbf{K})$ with $o' = x$ (from below the arrow on the right argument of the previous type) and $\tau_v = (\mathbf{Val} : \text{en})$. Then $\psi''_+ = (\mathbf{Val} : \text{en})_x$ from $\text{lsAProps}(x, (\mathbf{Val} : \text{en})) = (\mathbf{Val} : \text{en})_x | (\mathbf{Val} : \text{en})_x$ (see Sect. 3.3). Since $\tau = \mathbf{K}$, we check the method body with $\mathbf{K}_x, (\mathbf{Val} : \text{en})_x \vdash \text{“hello”} : \mathbf{S} ; \text{tt}|\text{tt} ; \emptyset$. Finally from the interface type τ_m , we know $\psi_+ = \psi_- = \text{tt}$, and $o = \emptyset$, which also agrees with the method body, above. Notice the overall type of a `defmethod` is the same as its first subexpression e_m .

It is worth noting the lack of special typing rules for overlapping methods—each method is checked independently based on local type information.

Subtyping. Multimethods are functions, via S-PMultiFn, which says a multimethod can be upcast to its interface type. Multimethod call sites are then handled by T-App via T-Subsume. Other rules are given in Fig. 8.

Semantics. Multimethod definition semantics are also given in Fig. 8. B-DefMulti creates a multimethod with the given dispatch function and an empty dispatch table. B-DefMethod produces a new multimethod with an extended dispatch table.

The overall dispatch mechanism is summarised by B-BetaMulti. First the dispatch function v_d is applied to the argument v' to obtain the dispatch value v_e . Based on v_e , the GM metafunction (Fig. 8) extracts a method v_f from the method table t and applies it to the original argument for the final result.

3.5 Precise Types for Heterogeneous Maps

Figure 10 presents heterogeneous map types. The type $(\mathbf{HMap}^\varepsilon \mathcal{M} \ \mathcal{A})$ contains \mathcal{M} , a map of *present* entries (mapping keywords to types), \mathcal{A} , a set of keyword keys that are known to be *absent* and tag \mathcal{E} which is either \mathcal{C} (“complete”) if the map is fully specified by \mathcal{M} , and \mathcal{P} (“partial”) if there are *unknown* entries. The partially specified map of `lunch` in Example 6 is written $(\mathbf{HMap}^\mathcal{P} \{(\mathbf{Val} : \text{en}) \ \mathbf{S}, (\mathbf{Val} : \text{fr}) \ \mathbf{S}\} \})$ (abbreviated **Lu**). The type of the

$e ::= \dots \mid (\text{get } e \ e) \mid (\text{assoc } e \ e \ e)$	Expressions
$v ::= \dots \mid \{\}$	Values
$\tau ::= \dots \mid (\mathbf{HMap}^\mathcal{E} \mathcal{M} \ \mathcal{A})$	Types
$\mathcal{M} ::= \{k \mapsto \vec{\tau}\}$	HMap mandatory entries
$\mathcal{A} ::= \{\vec{k}\}$	HMap absent entries
$\mathcal{E} ::= \mathcal{C} \mid \mathcal{P}$	HMap completeness tags

T-ASSOCHMAP	
$\Gamma \vdash e \Rightarrow (\text{assoc } e' \ e'_k \ e'_v) : (\mathbf{HMap}^\mathcal{E} \mathcal{M} \ \mathcal{A})$	$\Gamma \vdash e_k \Rightarrow e'_k : (\mathbf{Val} \ k) \quad \Gamma \vdash e_v \Rightarrow e'_v : \tau \quad k \notin \mathcal{A}$
$\Gamma \vdash (\text{assoc } e \ e_k \ e_v) \Rightarrow (\text{assoc } e' \ e'_k \ e'_v) : (\mathbf{HMap}^\mathcal{E} \mathcal{M}[k \mapsto \tau] \ \mathcal{A}) ; \text{tt} \text{fff} ; \emptyset$	

T-GETHMAP	
$\Gamma \vdash e \Rightarrow e' : (\bigcup (\mathbf{HMap}^\mathcal{E} \mathcal{M} \ \mathcal{A}) \xrightarrow{i}) ; \psi_{1+} \psi_{1-} ; o$	$\Gamma \vdash e_k \Rightarrow e'_k : (\mathbf{Val} \ k) \quad \overline{\mathcal{M}[k] = \tau} \xrightarrow{i}$
$\Gamma \vdash (\text{get } e \ e_k) \Rightarrow (\text{get } e' \ e'_k) : (\bigcup \overline{\tau}^i) ; \text{tt} \text{tt} ; \mathbf{key}_k(x)[o/x]$	

T-GETHMAPABSENT	
$\Gamma \vdash e \Rightarrow e' : (\mathbf{HMap}^\mathcal{E} \mathcal{M} \ \mathcal{A}) ; \psi_{1+} \psi_{1-} ; o$	$\Gamma \vdash e_k \Rightarrow e'_k : (\mathbf{Val} \ k) \quad k \in \mathcal{A}$
$\Gamma \vdash (\text{get } e \ e_k) \Rightarrow (\text{get } e' \ e'_k) : \mathbf{nil} ; \text{tt} \text{tt} ; \mathbf{key}_k(x)[o/x]$	

T-GETHMAPPARTIALDEFAULT	
$\Gamma \vdash e \Rightarrow e' : (\mathbf{HMap}^\mathcal{P} \mathcal{M} \ \mathcal{A}) ; \psi_{1+} \psi_{1-} ; o$	
$\Gamma \vdash e_k \Rightarrow e'_k : (\mathbf{Val} \ k) \quad k \notin \text{dom}(\mathcal{M}) \quad k \notin \mathcal{A}$	S-HMAPMONO
$\Gamma \vdash (\text{get } e \ e_k) \Rightarrow (\text{get } e' \ e'_k) : \top ; \text{tt} \text{tt} ; \mathbf{key}_k(x)[o/x]$	
$\vdash (\mathbf{HMap}^\mathcal{E} \mathcal{M} \ \mathcal{A}) <: (\mathbf{HMap}^\mathcal{E} \mathcal{M} \ \mathcal{A}) <: \mathbf{Map}$	

S-HMAPP	S-HMAP
$\forall i. \mathcal{M}[k_i] = \sigma_i \text{ and } \vdash \sigma_i <: \tau_i$	$\forall i. \mathcal{M}[k_i] = \sigma_i \text{ and } \vdash \sigma_i <: \tau_i \quad \mathcal{A}_1 \supseteq \mathcal{A}_2$
$\vdash (\mathbf{HMap}^\mathcal{C} \mathcal{M} \ \mathcal{A}') <: (\mathbf{HMap}^\mathcal{P} \{k \mapsto \vec{\tau}\}^i \ \mathcal{A}) \quad \vdash (\mathbf{HMap}^\mathcal{E} \mathcal{M} \ \mathcal{A}_1) <: (\mathbf{HMap}^\mathcal{E} \{k \mapsto \vec{\tau}\}^i \ \mathcal{A}_2)$	

B-ASSOC	B-GET	B-GETMISSING
$\rho \vdash e \Downarrow m \quad \rho \vdash e_k \Downarrow k$	$\rho \vdash e \Downarrow m \quad \rho \vdash e' \Downarrow k$	$\rho \vdash e \Downarrow m$
$\rho \vdash e_v \Downarrow v_v$	$k \in \text{dom}(m)$	$\rho \vdash e' \Downarrow k \quad k \notin \text{dom}(m)$
$\rho \vdash (\text{assoc } e \ e_k \ e_v) \Downarrow m[k \mapsto v_v]$	$\rho \vdash (\text{get } e \ e') \Downarrow m[k]$	$\rho \vdash (\text{get } e \ e') \Downarrow \mathbf{nil}$

Fig. 10. HMap syntax, typing and operational semantics

$\text{restrict}(\tau, \sigma) = \perp$ if $\not\exists v. \vdash v : \tau ; \psi ; o$ and $\vdash v : \sigma ; \psi' ; o'$	$\text{remove}(\tau, \sigma) = \perp$ if $\not\vdash \tau <: \sigma$
$\text{restrict}(\tau, \sigma) = \tau$ if $\vdash \tau <: \sigma$	$\text{remove}(\tau, \sigma) = \tau$ otherwise
$\text{restrict}(\tau, \sigma) = \sigma$ otherwise	

Fig. 11. Restrict and remove

fully specified map `breakfast` in Example 5 elides the absent entries, written $(\mathbf{HMap}^\mathcal{C} \{(\mathbf{Val} : \text{en}) \ \mathbf{S}, (\mathbf{Val} : \text{fr}) \ \mathbf{S}\})$ (abbreviated **Bf**). To ease presentation, if an HMap has completeness tag \mathcal{C} then \mathcal{A} is elided and implicitly contains all keywords not in the domain of \mathcal{M} —dissociating keys is not modelled, so the set of absent entries otherwise never grows. Keys cannot be both present and absent.

The metavariable m ranges over the runtime value of maps $\{\overline{k \mapsto v}\}$, usually written $\{k \ v\}$. We only provide syntax for the empty map literal, however

when convenient we abbreviate non-empty map literals to be a series of `assoc` operations on the empty map. We restrict lookup and extension to keyword keys.

How to Check. A mandatory lookup is checked by `T-GetHMap`.

$$\lambda b^{\mathbf{Bf}}.(\text{get } b : \text{en})$$

The result type is \mathbf{S} , and the return object is $\mathbf{key}_{:\text{en}}(b)$. The object $\mathbf{key}_k(x)[o/x]$ is a symbolic representation for a keyword lookup of k in o . The substitution for x handles the case where o is empty.

$$\mathbf{key}_k(x)[y/x] = \mathbf{key}_k(y) \quad \mathbf{key}_k(x)[\emptyset/x] = \emptyset$$

An absent lookup is checked by `T-GetHMapAbsent`.

$$\lambda b^{\mathbf{Bf}}.(\text{get } b : \text{bocce})$$

The result type is `nil`—since \mathbf{Bf} is fully specified—with return object $\mathbf{key}_{:\text{bocce}}(b)$.

A lookup that is not present or absent is checked by `T-GetHMapPartial-Default`.

$$\lambda u^{\mathbf{Lu}}.(\text{get } u : \text{bocce})$$

The result type is \top —since \mathbf{Lu} has an unknown `:bocce` entry—with return object $\mathbf{key}_{:\text{bocce}}(u)$. Notice propositions are erased once they enter a `HMap` type.

For presentational reasons, lookups on unions of `HMaps` are only supported in `T-GetHMap` and each element of the union must contain the relevant key.

$$\lambda u^{(\cup \mathbf{Bf} \mathbf{Lu})}.(\text{get } u : \text{en})$$

The result type is \mathbf{S} , and the return object is $\mathbf{key}_{:\text{en}}(u)$. However, lookups of `:bocce` on $(\cup \mathbf{Bf} \mathbf{Lu})$ maps are unsupported. This restriction still allows us to check many of the examples in Sect. 2—in particular we can check Example 8, as `:Meal` is in common with both `HMaps`, but cannot check Example 9 because a `:combo` meal lacks a `:desserts` entry. Adding a rule to handle Example 9 is otherwise straightforward.

Extending a map with `T-AssocHMap` preserves its completeness.

$$\lambda b^{\mathbf{Bf}}.(\text{assoc } b : \text{au "beans"})$$

The result type is $(\mathbf{HMap}^c \{(\mathbf{Val} : \text{en}) \mathbf{S}, (\mathbf{Val} : \text{fr}) \mathbf{S}, (\mathbf{Val} : \text{au}) \mathbf{S}\})$, a complete map. `T-AssocHMap` also enforces $k \notin \mathcal{A}$ to prevent badly formed types.

Subtyping. Subtyping for `HMaps` designate \mathbf{Map} as a common supertype for all `HMaps`. `S-HMap` says that `HMaps` are subtypes if they agree on \mathcal{E} , agree on mandatory entries with subtyping and at least cover the absent keys of the supertype. Complete maps are subtypes of partial maps as long as they agree on the mandatory entries of the partial map via subtyping (`S-HMapP`).

The semantics for `get` and `assoc` are straightforward.

$$\begin{aligned}
 \text{update}(\left(\bigcup \bar{\tau}\right), \nu, \pi) &= \overrightarrow{\left(\bigcup \text{update}(\tau, \nu, \pi)\right)} \\
 \text{update}(\tau, (\mathbf{Val} C), \pi :: \mathbf{class}) &= \text{update}(\tau, C, \pi) \\
 \text{update}(\tau, \nu, \pi :: \mathbf{class}) &= \tau \\
 \text{update}((\mathbf{HMap}^{\mathcal{E}} \mathcal{M} \mathcal{A}), \nu, \pi :: \mathbf{key}_k) &= (\mathbf{HMap}^{\mathcal{E}} \mathcal{M}[k \mapsto \text{update}(\tau, \nu, \pi)] \mathcal{A}) \\
 &\quad \text{if } \mathcal{M}[k] = \tau \\
 \text{update}((\mathbf{HMap}^{\mathcal{E}} \mathcal{M} \mathcal{A}), \nu, \pi :: \mathbf{key}_k) &= \perp \quad \text{if } \vdash \mathbf{nil} \not\prec: \nu \text{ and } k \in \mathcal{A} \\
 \text{update}((\mathbf{HMap}^{\mathcal{P}} \mathcal{M} \mathcal{A}), \tau, \pi :: \mathbf{key}_k) &= (\cup (\mathbf{HMap}^{\mathcal{P}} \mathcal{M}[k \mapsto \tau] \mathcal{A}) \\
 &\quad (\mathbf{HMap}^{\mathcal{P}} \mathcal{M} (\mathcal{A} \cup \{k\}))) \\
 &\quad \text{if } \vdash \mathbf{nil} \prec: \tau, k \notin \text{dom}(\mathcal{M}) \text{ and } k \notin \mathcal{A} \\
 \text{update}((\mathbf{HMap}^{\mathcal{P}} \mathcal{M} \mathcal{A}), \nu, \pi :: \mathbf{key}_k) &= (\mathbf{HMap}^{\mathcal{P}} \mathcal{M}[k \mapsto \text{update}(\tau, \nu, \pi)] \mathcal{A}) \\
 &\quad \text{if } \vdash \mathbf{nil} \not\prec: \nu, k \notin \text{dom}(\mathcal{M}) \text{ and } k \notin \mathcal{A} \\
 \text{update}(\tau, \nu, \pi :: \mathbf{key}_k) &= \tau \\
 \text{update}(\tau, \sigma, \epsilon) &= \text{restrict}(\tau, \sigma) \\
 \text{update}(\tau, \bar{\sigma}, \epsilon) &= \text{remove}(\tau, \sigma)
 \end{aligned}$$

Fig. 12. Type update (the metavariable ν ranges over τ and $\bar{\tau}$ (without variables), $\vdash \mathbf{nil} \not\prec: \bar{\tau}$ when $\vdash \mathbf{nil} \prec: \tau$, see Fig. 11 for restrict and remove.)

3.6 Proof System

The occurrence typing proof system uses standard propositional logic, except for where nested information is combined. This is handled by L-Update:

$$\text{L-UPDATE} \quad \frac{\Gamma \vdash \tau_{\pi'(x)} \quad \Gamma \vdash \nu_{\pi(\pi'(x))}}{\Gamma \vdash \text{update}(\tau, \nu, \pi)_{\pi'(x)}}$$

It says under Γ , if object $\pi'(x)$ is of type τ , and an extension $\pi(\pi'(x))$ is of possibly-negative type ν , then $\text{update}(\tau, \nu, \pi)$ is $\pi'(x)$'s type under Γ .

Recall Example 8. Solving \mathbf{Order}_o , $(\mathbf{Val} : \text{combo})_{\mathbf{key} : \text{Meal}(o)} \vdash \tau_o$ uses L-Update, where $\pi = \epsilon$ and $\pi' = [\mathbf{key} : \text{Meal}]$.

$$\Gamma \vdash \text{update}(\mathbf{Order}, (\mathbf{Val} : \text{combo}), [\mathbf{key} : \text{Meal}])_o$$

Since \mathbf{Order} is a union of HMaps, we structurally recur on the first case of update (Fig. 12), which preserves π . Each initial recursion hits the first HMap case, since there is some τ such that $\mathcal{M}[k] = \tau$ and \mathcal{E} accepts partial maps \mathcal{P} .

To demonstrate, `:lunch` meals are handled by the first HMap case and update to $(\mathbf{HMap}^{\mathcal{P}} \mathcal{M}[(\mathbf{Val} : \text{Meal}) \mapsto \sigma'] \{\})$ where $\sigma' = \text{update}((\mathbf{Val} : \text{lunch}), (\mathbf{Val} : \text{combo}), \epsilon)$ and $\mathcal{M} = \{(\mathbf{Val} : \text{Meal}) \mapsto (\mathbf{Val} : \text{lunch}), (\mathbf{Val} : \text{desserts}) \mapsto \mathbf{N}\}$. σ' updates to \perp via the penultimate update case, because $\text{restrict}((\mathbf{Val} : \text{lunch}), (\mathbf{Val} : \text{combo})) = \perp$ by the first restrict case. The same happens to `:dinner` meals, leaving just the `:combo` HMap.

In Example 9, $\Gamma \vdash \text{update}(\mathbf{Order}, \mathbf{Long}, [\mathbf{class}, \mathbf{key} : \text{desserts}])_o$ updates the argument in the \mathbf{Long} method. This recurs twice for each meal to handle the \mathbf{class} path element.

We describe the other **update** cases. The first **class** case updates to C if *class* returns (**Val** C). The second **key_k** case detects contradictions in absent keys. The third **key_k** case updates unknown entries to be mapped to τ or absent. The fourth **key_k** case updates unknown entries to be *present* when they do not overlap with **nil**.

4 Metatheory

We prove type soundness following Tobin-Hochstadt and Felleisen [24]. Our model is extended to include errors **err** and a *wrong* value, and we prove well-typed programs do not go wrong; this is therefore a stronger theorem than proved by Tobin-Hochstadt and Felleisen [24]. Errors behave like Java exceptions—they can be thrown and propagate “upwards” in the evaluation rules (**err** rules are deferred to the appendix).

Rather than modeling Java’s dynamic semantics, a task of daunting complexity, we instead make our assumptions about Java explicit. We concede that method and constructor calls may diverge or error, but assume they can never go wrong (other assumptions given in the supplemental material).

Assumption 1 (JVM_{new}). *If $\forall i. v_i = C_i \{\overrightarrow{fld_j : v_j}\}$ or $v_i = \text{nil}$ and v_i is consistent with ρ then either*

- $\text{JVM}_{\text{new}}[C, [\overrightarrow{C_i}], [\overrightarrow{v_i}]] = C \{\overrightarrow{fld_k : v_k}\}$ which is consistent with ρ ,
- $\text{JVM}_{\text{new}}[C, [\overrightarrow{C_i}], [\overrightarrow{v_i}]] = \text{err}$, or
- $\text{JVM}_{\text{new}}[C, [\overrightarrow{C_i}], [\overrightarrow{v_i}]]$ is undefined.

For the purposes of our soundness proof, we require that all values are *consistent*. Consistency (defined in the supplemental material) states that the types of closures are well-scoped—they do not claim propositions about variables hidden in their closures.

We can now state our main lemma and soundness theorem. The metavariable α ranges over v , **err** and *wrong*. Proofs are deferred to the supplemental material.

Lemma 1. *If $\Gamma \vdash e' \Rightarrow e : \tau$; $\psi_+ | \psi_-$; $o, \rho \models \Gamma$, ρ is consistent, and $\rho \vdash e \Downarrow \alpha$ then either*

- $\rho \vdash e \Downarrow v$ and all of the following hold:
 1. either $o = \emptyset$ or $\rho(o) = v$,
 2. either $\text{TrueVal}(v)$ and $\rho \models \psi_+$ or $\text{FalseVal}(v)$ and $\rho \models \psi_-$,
 3. $\vdash v \Rightarrow v : \tau$; $\psi'_+ | \psi'_-$; o' for some ψ'_+ , ψ'_- and o' , and
 4. v is consistent with ρ , or
- $\rho \vdash e \Downarrow \text{err}$.

Theorem 1 (Type Soundness). *If $\Gamma \vdash e' \Rightarrow e : \tau$; $\psi_+ | \psi_-$; o and $\rho \vdash e \Downarrow v$ then $\vdash v \Rightarrow v : \tau$; $\psi'_+ | \psi'_-$; o' for some ψ'_+ , ψ'_- and o' .*

5 Experience

Typed Clojure is implemented as `core.typed` [2], which has seen wide usage.

5.1 Implementation

`core.typed` provides preliminary integration with the Clojure compilation pipeline, primarily to resolve Java interoperability.

The `core.typed` implementation extends this paper in several key areas to handle checking real Clojure code, including an implementation of Typed Racket’s variable-arity polymorphism [22], and support for other Clojure idioms like datatypes and protocols. There is no integration with Java Generics, so only Java 1.4-style erased types are “trusted” by `core.typed`. Casts are needed to recover the discarded information, which—for collections—are then tracked via Clojure’s universal sequence interface [14].

	feeds2imap	CircleCI
Total number of typed namespaces	11 (825 LOC)	87 (19,000 LOC)
Total number of <code>def</code> expressions	93	1834
• checked	52 (56%)	407 (22%)
• unchecked	41 (44%)	1427 (78%)
Total number of Java interactions	32	105
• static methods	5 (16%)	26 (25%)
• instance methods	20 (62%)	36 (34%)
• constructors	6 (19%)	38 (36%)
• static fields	1 (3%)	5 (5%)
Methods overridden to return non-nil	0	35
Methods overridden to accept nil arguments	0	1
Total HMap lookups	27	328
• resolved to mandatory key	20 (74%)	208 (64%)
• resolved to optional key	6 (22%)	70 (21%)
• resolved of absent key	0 (0%)	20 (6%)
• unresolved key	1 (4%)	30 (9%)
Total number of <code>defalias</code> expressions	18	95
• contained HMap or union of HMap type	7 (39%)	62 (65%)
Total number of checked <code>defmulti</code> expressions	0	11
Total number of checked <code>defmethod</code> expressions	0	89

Fig. 13. Typed Clojure features used in practice

5.2 Evaluation

Throughout this paper, we have focused on three interrelated type system features: heterogenous maps, Java interoperability, and multimethods. Our hypothesis is that these features are widely used in existing Clojure programs

in interconnecting ways, and that handling them as we have done is required to type check realistic Clojure programs.

To evaluate this hypothesis, we analyzed two existing `core.typed` code bases, one from the open-source community, and one from a company that uses `core.typed` in production. For our data gathering, we instrumented the `core.typed` type checker to record how often various features were used (summarized in Fig. 13).

feeds2imap. `feeds2imap`⁴ is an open source library written in Typed Clojure. It provides an RSS reader using the *javax.mail* framework.

Of 11 typed namespaces containing 825 lines of code, there are 32 Java interactions. The majority are method calls, consisting of 20 (62%) instance methods and 5 (16%) static methods. The rest consists of 1 (3%) static field access, and 6 (19%) constructor calls—there are no instance field accesses.

There are 27 lookup operations on HMap types, of which 20 (74%) resolve to mandatory entries, 6 (22%) to optional entries, and 1 (4%) is an unresolved lookup. No lookups involved fully specified maps.

From 93 `def` expressions in typed code, 52 (56%) are checked, with a rate of 1 Java interaction for 1.6 checked top-level definitions, and 1 HMap lookup to 1.9 checked top-level definitions. That leaves 41 (44%) unchecked vars, mainly due to partially complete porting to Typed Clojure, but in some cases due to unannotated third-party libraries.

No typed multimethods are defined or used. Of 18 total type aliases, 7 (39%) contained one HMap type, and none contained unions of HMaps—on further inspection there was no HMap entry used to dictate control flow, often handled by multimethods. This is unusual in our experience, and is perhaps explained by *feeds2imap* mainly wrapping existing *javax.mail* functionality.

CircleCI. CircleCI [7] provides continuous integration services built with a mixture of open- and closed-source software. Typed Clojure was used at CircleCI in production systems for two years [8], maintaining 87 namespaces and 19,000 lines of code, an experience we summarise in Sect. 5.3.

The CircleCI code base contains 11 checked multimethods. All 11 dispatch functions are on a HMap key containing a keyword, in a similar style to Example 8. Correspondingly, all 89 methods are associated with a keyword dispatch value. The argument type was in all cases a single HMap type, however, rather than a union type. In our experience from porting other libraries, this is unusual.

Of 328 lookup operations on HMaps, 208 (64%) resolve to mandatory keys, 70 (21%) to optional keys, 20 (6%) to absent keys, and 30 (9%) lookups are unresolved. Of 95 total type aliases defined with `defalias`, 62 (65%) involved one or more HMap types. Out of 105 Java interactions, 26 (25%) are static methods, 36 (34%) are instance methods, 38 (36%) are constructors, and 5 (5%) are static fields. 35 methods are overridden to return non-nil, and 1 method overridden to accept nil—suggesting that `core.typed` disallowing nil as a method argument by default is justified.

⁴ <https://github.com/frenchy64/feeds2imap.clj>.

Of 464 checked top-level definitions (which consists of 57 `defmethod` calls and 407 `def` expressions), 1 HMap lookup occurs per 1.4 top-level definitions, and 1 Java interaction occurs every 4.4 top-level definitions.

From 1834 `def` expressions in typed code, only 407 (22%) were checked. That leaves 1427 (78%) which have unchecked definitions, either by an explicit `:no-check` annotation or `tc-ignore` to suppress type checking, or the `warn-on-unannotated-vars` option, which skips `def` expressions that lack expected types via `ann`. From a brief investigation, reasons include unannotated third-party libraries, work-in-progress conversions to Typed Clojure, unsupported Clojure idioms, and hard-to-check code.

Lessons. Based on our empirical survey, HMaps and Java interoperability support are vital features used on average more than once per typed function. Multimethods are less common in our case studies. The CircleCI code base contains only 26 multimethods total in 55,000 lines of mixed untyped-typed Clojure code, a low number in our experience.

5.3 Further Challenges

After a 2 year trial, the second case study decided to disabled type checking [9]. They were supportive of the fundamental ideas presented in this paper, but primarily cited issues with the checker implementation in practice and would reconsider type checking if they were resolved. This is also supported by Fig. 13, where 78% of `def` expressions are unchecked.

Performance. Rechecking files with transitive dependencies is expensive since all dependencies must be rechecked. We conjecture caching type state will significantly improve re-checking performance, though preserving static soundness in the context of arbitrary code reloading is a largely unexplored area.

Library Annotations. Annotations for external code are rarely available, so a large part of the untyped-typed porting process is reverse engineering libraries.

Unsupported Idioms. While the current set of features is vital to checking Clojure code, there is still much work to do. For example, common Clojure functions are often too polymorphic for the current implementation or theory to account for. The post-mortem [9] contains more details.

6 Related Work

Multimethods. Millstein [20] and collaborators present a sequence of systems [4, 5, 20] with statically-typed multimethods and modular type checking. In contrast to Typed Clojure, in these system methods declare the types of arguments that they expect which corresponds to exclusively using `class` as the dispatch function in Typed Clojure. However, Typed Clojure does not attempt to rule out failed dispatches.

Record Types. Row polymorphism [3, 12, 26], used in systems such as the OCaml object system, provides many of the features of HMap types, but defined using universally-quantified row variables. HMaps in Typed Clojure are instead designed to be used with subtyping, but nonetheless provide similar expressiveness, including the ability to require presence and absence of certain keys.

Dependent JavaScript [6] can track similar invariants as HMaps with types for JS objects. They must deal with mutable objects, they feature refinement types and strong updates to the heap to track changes to objects.

TeJaS [16], another type system for JavaScript, also supports similar HMaps, with the ability to record the presence and absence of entries, but lacks a compositional flow-checking approach like occurrence typing.

Typed Lua [18] has *table types* which track entries in a mutable Lua table. Typed Lua changes the dynamic semantics of Lua to accommodate mutability: Typed Lua raises a runtime error for lookups on missing keys—HMaps consider lookups on missing keys normal.

Java Interoperability in Statically Typed Languages. Scala [21] has nullable references for compatibility with Java. Programmers must manually check for `null` as in Java to avoid null-pointer exceptions.

Other Optional and Gradual Type Systems. Several other gradual type systems have been developed for existing dynamically-typed languages. Reticulated Python [25] is an experimental gradually typed system for Python, implemented as a source-to-source translation that inserts dynamic checks at language boundaries and supporting Python’s first-class object system. Clojure’s nominal classes avoids the need to support first-class object system in Typed Clojure, however HMaps offer an alternative to the structural objects offered by Reticulated. Similarly, Gradualtalk [1] offers gradual typing for Smalltalk, with nominal classes.

Optional types have been adopted in industry, including Hack [10], and Flow [11] and TypeScript [19], two extensions of JavaScript. These systems support limited forms of occurrence typing, and do not include the other features we present.

7 Conclusion

Optional type systems must be designed with close attention to the language that they are intended to work for. We have therefore designed Typed Clojure, an optionally-typed version of Clojure, with a type system that works with a wide variety of distinctive Clojure idioms and features. Although based on the foundation of Typed Racket’s occurrence typing approach, Typed Clojure both extends the fundamental control-flow based reasoning as well as applying it to handle seemingly unrelated features such as multi-methods. In addition, Typed Clojure supports crucial features such as heterogeneous maps and Java interoperability while integrating these features into the core type system. Not only are each of these features important in isolation to Clojure and Typed

Clojure programmers, but they must fit together smoothly to ensure that existing untyped programs are easy to convert to Typed Clojure.

The result is a sound, expressive, and useful type system which, as implemented in `core.typed` with appropriate extensions, is suitable for typechecking a significant amount of existing Clojure programs. As a result, Typed Clojure is already successful: it is used in the Clojure community among both enthusiasts and professional programmers.

Our empirical analysis of existing Typed Clojure programs bears out our design choices. Multimethods, Java interoperation, and heterogeneous maps are indeed common in both Clojure and Typed Clojure, meaning that our type system must accommodate them. Furthermore, they are commonly used together, and the features of each are mutually reinforcing. Additionally, the choice to make Java's `null` explicit in the type system is validated by the many Typed Clojure programs that specify non-nullable types.

References

1. Allende, E., Callau, O., Fabry, J., Tanter, É., Denker, M.: Gradual typing for Smalltalk. *Sci. Comput. Program.* **96**, 52–69 (2014)
2. Bonnaire-Sergeant, A.: contributors, `core.typed`. <https://github.com/clojure/core.typed>
3. Cardelli, L., Mitchell, J.C.: Operations on records. *Math. Struct. Comput. Sci.* **1**, 3–48 (1991)
4. Chambers, C.: Object-oriented multi-methods in Cecil. In: Proceedings of the ECOOP (1992)
5. Chambers, C., Leavens, G.T.: Typechecking and modules for multi-methods. In: Proc. OOPSLA (1994)
6. Chugh, R., Herman, D., Jhala, R.: Dependent types for JavaScript. In: Proceedings of the OOPSLA (2012)
7. CircleCI: CircleCI. <https://circleci.com>
8. CircleCI: Why we're supporting Typed Clojure, and you should too! (September 2013). <http://blog.circleci.com/supporting-typed-clojure/>
9. CircleCI; O'Morain, M.: Why we're no longer using `core.typed` (September 2015). <http://blog.circleci.com/why-were-no-longer-using-core-typed/>
10. Facebook: Hack language specification. Technical report, Facebook (2014)
11. Facebook: Flow language specification. Technical report, Facebook (2015)
12. Harper, R., Pierce, B.: A record calculus based on symmetric concatenation. In: Proceedings of the POPL (1991)
13. Hickey, R.: The Clojure programming language. In: Proceedings of the DLS (2008)
14. Hickey, R.: Clojure sequence documentation (February 2015). <http://clojure.org/sequences>
15. Lehtosalo, J.: mypy. <http://mypy-lang.org/>
16. Lerner, B.S., Politz, J.G., Guha, A., Krishnamurthi, S.: TeJaS: retrofitting type systems for JavaScript. In: DLS 2013, Proceedings of the 9th Symposium on Dynamic Languages, pp. 1–16. ACM, New York, NY, USA (2013). <http://doi.acm.org/10.1145/2508168.2508170>
17. Lucassen, J.M., Gifford, D.K.: Polymorphic effect systems. In: Proceedings of the POPL (1988)

18. Maidl, A.M., Mascarenhas, F., Ierusalimschy, R.: Typed Lua: an optional type system for Lua. In: *Proceedings of the Dyla (2014)*
19. Microsoft: Typescript language specification. Technical report Version 1.4, Microsoft (2014)
20. Millstein, T., Chambers, C.: Modular statically typed multimethods. In: *Information and Computation*. pp. 279–303. Springer-Verlag (2002)
21. Odersky, M., Cremet, V., Dragos, I., Dubochet, G., Emir, B., McDirmid, S., Micheloud, S., Mihaylov, N., Schinz, M., Stenman, E., Spoon, L., Zenger, M., et al.: An overview of the Scala programming language, 2nd edn, Technical report, EPFL Lausanne, Switzerland (2006)
22. Strickland, T.S., Tobin-Hochstadt, S., Felleisen, M.: Practical variable-arity polymorphism. In: Castagna, G. (ed.) *ESOP 2009*. LNCS, vol. 5502, pp. 32–46. Springer, Heidelberg (2009)
23. Tobin-Hochstadt, S., Felleisen, M.: The design and implementation of Typed Scheme. In: *Proceedings of the POPL (2008)*
24. Tobin-Hochstadt, S., Felleisen, M.: Logical types for untyped languages. In: *ICFP 2010, Proceedings of the ICFP (2010)*
25. Vitousek, M.M., Kent, A.M., Siek, J.G., Baker, J.: Design and evaluation of gradual typing for Python. In: *Proc. DLS (2014)*
26. Wand, M.: Type inference for record concatenation and multiple inheritance (1989)