

# Squash the work!

## Inferring Recursive Type Annotations from Plain Data for Optional Type Systems

Ambrose Bonnaire-Sergeant

Sam Tobin-Hochstadt

### Abstract

To type check a dynamically-typed program with an optional type system, type annotations must be added. This burden is sometimes large, and has put off real users attempting to migrate to existing optional type systems. When not discouraged, programmers often annotate tens of thousands of lines of code without assistance.

We present an approach to lighten the load on programmers moving to optional types. Our only requirement of existing programs is that they be runnable, with a suite of tests or examples. Given a running program, we instrument the execution, record type information, summarize it, and annotate the existing program with the recovered types.

We apply our approach to Clojure, a dynamically typed language with a culture of unit testing as well as both an existing optional type system and a contract system. Given a component under consideration, we instrument the source and analyze the behavior of the program while running unit tests. Equipped with this information, we summarize it by generating compact type specifications for all the functions in the component, including well-named type definitions. Our tool can also automatically generate contracts using the Clojure spec tool. Since Clojure relies heavily on ad-hoc data structures in the Lisp tradition, we describe an algorithm for automatically inferring recursive structural types from data examples, a challenge not considered in prior work.

Our approach, as must be the case for a testing-driven tool, is incomplete—programs may have too few unit tests, and untested execution paths can have differing type behavior. We therefore evaluate our tool by running it on real Clojure programs and then completing the porting to Typed Clojure. We find that while some changes are always needed, the generated types are valuable and the effort reduction is substantial.

### 1 Introduction

*It is better to have 100 functions operate on one data structure than 10 functions on 10 data structures.*—Alan Perlis

Optional type systems [6] extend existing untyped languages with type checking. For example, TypeScript [3] and Flow [1] extend JavaScript, Hack [2] extends PHP, and mypy [16] extends Python. They typically support existing syntax and idioms of their target languages.

Transitioning to an optional type system requires adding type annotations to existing code, a significant manual burden. This overhead has sparked interest in tooling to help create [4, 7, 10, 12, 20] and evolve [15] these annotations.

```
(defn nodes
  "Returns the number of nodes in a binary tree."
  [t] (case (:op t)
        :leaf 1
        :node (+ 1 (nodes (:left t))
                (nodes (:right t)))))

(assert (= 3 (nodes
             {:op :node,
              :left {:op :leaf, :val 2},
              :right {:op :leaf, :val 3}))))
```

**Figure 1.** A typical use of maps in Clojure to represent records, that requires type annotations to check with Typed Clojure. Our tool will automatically annotate this program with a useful recursive type (Figure 3).

Clojure [14] is an untyped language that compiles to the Java Virtual Machine. Compared to languages already mentioned, it strongly encourages programming with plain data structures, and is a good example of implementing Perlis’ advise in our opening quote. Clojure provides many functions and idioms around persistent, immutable hash-maps, including literal map syntax `{k v ...}`, interned *keywords* suitable both for map keys (e.g., `:a`, `:b`) and functions that look themselves up in a map (e.g., `(:a {:a 1}) => 1`), and a suite of functions to deeply transform, manipulate, and validate maps, with multimethods providing open extension. Typed Clojure [5] is an optional type system for Clojure designed to recognize these idioms given sufficient type annotations—which our tool assists the programmer in writing.

Maps in Clojure often replace records or objects, demonstrated in Figure 1: instead of representing a binary tree with `Node` and `Leaf` classes, they are encoded in maps with an explicit keyword *dispatch entry* (e.g., `:op`) to distinguish cases—`{:op :leaf, :val ...}` for instances of `Leaf`, and `{:op :node, :left ..., :right ...}` for `Node`.

This emphasis on maps has far-reaching implications for Typed Clojure. Types for maps (written `'{:op ':leaf, :val ...}` and `'{:op ':node, :left ..., :right ...}` for our examples) combine with ad-hoc union types and equirecursive type aliases in the type:

```
(defalias Tree
  (U '{:op ':node, :left Tree, :right Tree}
     '{:op ':leaf, :val Int}))
```

```

111 function nodes(t: {left: {op: string, val: number},
112                 op: string,
113                 right: {op: string, val: number}}
114                 | {op: string, val: number}) ...
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165

```

Figure 2. TypeWiz’s TypeScript annotations for Figure 1.

Porting Figure 1 to Typed Clojure involves writing a type definition `Tree`, and annotating `nodes` as

```
(ann nodes [Tree -> Int])
```

Existing annotation tools fail to generate these types since only classes are given recursive types, and so cannot infer recursive types for plain data, such as JSON or heterogeneous dictionaries. To demonstrate the current state-of-the-art in automatic annotation tools for optional type systems, we transliterate Figure 1 to JavaScript using plain objects. We use TypeWiz [21] to generate TypeScript annotations via dynamic analysis, as it is well maintained and generates comparable annotations to similar tools [4, 7, 10, 12, 15, 20].

Figure 2 shows the actual output of TypeWiz for a JavaScript translation of Figure 1. Unfortunately, the annotation is too specific: it only accepts trees of height 1 or 2. For fair comparison, even a class-based translation to JavaScript with a common `nodes` method yields a shallow annotation:

```

138 class Node {
139   public left: Leaf;
140   public right: Leaf;
141   ...
142 }
143
144 class Leaf {
145   public data: number;
146   ...
147 }

```

Since TypeWiz uses dynamic analysis, it is faithfully providing the exact types that are observed at runtime. Unfortunately, incrementally better test coverage does not quickly converge to useful annotations. For example, if we add an example of a nested `Node` on the just the left branch the annotation for the plain objects version is still not recursive (alarmingly, it instead grows linearly in the number of nodes), and the class-based version helpfully updates the type of `left` to `Leaf|Node`, but `right` still remains `Leaf`.

On the other hand, our approach recognizes Figure 1 as traversing recursively defined data with two cases, distinguished by the `:op` entry (Figure 3).

Our approach is sensitive enough to to compute optional keys for each constructor. Adding a unit test that includes a `:val` entry in a `:node` yields the annotation

```

160 (defalias Op
161   (U (HMap :mandatory
162       {:op ':node, :left Op, :right Op}
163       :optional {:val Int}))
164     '{:op ':leaf, :val Int}))
165

```

```

166 (defalias Op
167   (U '{:op ':node, :left Op, :right Op}
168     '{:op ':leaf, :val Int}))
169 (ann nodes [Op -> Int])
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220

```

Figure 3. Our tool’s Typed Clojure annotation of Figure 1. `defalias` introduces an equirecursive type alias, `U` is a set-theoretic union type constructor, and `' :node` is a singleton type containing just the keyword value `:node`.

```

177 (defalias Op
178   (U (HMap :mandatory
179       {:op ':node, :left Op, :right Op}
180       :optional {:val Int}))
181     '{:op ':node3, :left Op, :mid Op, :right Op}
182     '{:op ':leaf, :val Int}))
183 (ann nodes [Op -> Int])
184 (ann nodes' [Op -> Int])
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220

```

Figure 4. Optional entries and combined information across functions. `HMap` is the expanded heterogeneous map type constructor, specifying both mandatory and optional entries.

Furthermore, we aggressively combine recursive data used in the same file. Given a 3-way node example with the key `:node3`, used as input to a distinct function `nodes'`, our tool generates the combined type shown in Figure 4.

## Contributions

- We outline a generalized approach to automatically generating type annotations (Section 2).
- We describe a novel approach to reconstructing recursively defined structural records from fully unrolled examples in a formal model of our inference algorithm (Section 3).
- We show how to extend our approach with space-efficient and lazy runtime tracking (Section 4).
- We report our experience using this algorithm to generate types, tests, and contracts on several Clojure libraries and programs (Section 5).

## 2 Overview

Now that we have introduced the problem, we can flesh out our philosophy and overall approach. To concretize our discussion, Figure 5 demonstrates our tool’s output when generating types for the 1,776 line file `cljs.compiler`.

Its main function is `emit`, which effectfully converts a map-based AST to JavaScript. The AST is created by functions in `cljs.analyzer`, a significantly larger 4,366 line Clojure file. Without inspecting `cljs.analyzer`, our tool annotates `emit` on line 23 with a recursive AST type `Op` (lines 1-12).

An important question to address is “how accurate are these annotations?”. Unlike previous work in this area [4], we do not aim for soundness guarantees in our generated types. A significant contribution of our work is a tool that Clojure programmers can use to help learn about and specify their programs. In that spirit, we strive to generate annotations meeting more qualitative criteria. Each guideline by itself helps generate more useful annotations and, as we discuss in Section 5.1, they combine in interesting ways help to make up for shortcomings.

**Choose recognizable names** Assigning a good name for a type increases readability by succinctly conveying its purpose. Along those lines, a good name for the AST representation on lines 1-12 might be `AST` or `Expr`. However, these kinds of names can be very misleading when incorrect, so instead of guessing them, our tool takes a more consistent approach and generates *easily recognizable* names based on the type the name points to. Then, those with a passing familiarity with the data flowing through the program can quickly identify and rename them. For example,

- `Op` (lines 1-12) is chosen because `:op` is clearly the dispatch key (the `:op` entry is also helpfully placed as the first entry in each case to aid discoverability),
- `ColumnLineContextMap` (lines 13-15) enumerates the keys of the map type it points to,
- `NameShadowMap` and `FnScopeFnSelfNameNsMap` (referenced on lines 4 and 5) similarly, and
- `HMap49305` (lines 16-22) shows how our tool fails to give names to certain combinations of types (we discuss the severity of this particular situation in Section 5.1).

**Favor compact annotations** Literally translating runtime observations into annotations without compacting them leads to unmaintainable and impractical types resembling TypeWiz’s annotation for nodes (Figure 2). To avoid this, we use optional keys where possible, like line 15, infer recursive types like `Op`, and reuse type aliases in function annotations, like `emit` and `emit-dot` (lines 23, 24). These processes of compacting annotations often makes them more general, which leads into our next goal.

**Don’t overspecify types** Poor test coverage can easily skew the results of dynamic analysis tools, so we choose to err on the side of generalizing types where possible. Our opening example nodes (Figure 1) is a good example of this—our inferred type (Figure 3) is recursive, despite nodes only being tested with a tree of height 2. This has several benefits.

- We avoid exhausting the pool of easily recognizable names by generalizing types to communicate the general role of an argument or return position. For example, `emit-dot` (line 24) is annotated to take `Op`, but in reality accepts only a subset of `Op`. Programmers can combine the recognizability of `Op` with the suggestive

```

1 (defalias Op ; omitted some entries and 11 cases 276
2   (U (HMap :mandatory 277
3     {:op ':binding, 278
4      :info (U NameShadowMap 279
5            FnScopeFnSelfNameNsMap), ...}) 280
6     :optional 281
7     {:env ColumnLineContextMap, :init Op, 282
8      :shadow (U nil Op), ...}) 283
9     '(:op ':const, :env HMap49305, ...) 284
10    '(:op ':do, :env HMap49305, 285
11     :ret Op, :statements (Vec Nothing), ...) 286
12    ...)) 287
13 (defalias ColumnLineContextMap 288
14   (HMap :mandatory {:column Int, :line Int} 289
15     :optional {:context ':expr})) 290
16 (defalias HMap49305 ; omitted some entries 291
17   (U nil 292
18     '(:context ':statement, :column Int, ...) 293
19     '(:context ':return, :column Int, ...) 294
20     (t/HMap :mandatory 295
21       {:context ':expr, :column Int, ...) 296
22       :optional {...}})) 297
23 (ann emit [Op -> nil]) 298
24 (ann emit-dot [Op -> nil]) 299

```

**Figure 5.** Sample raw output from our tool inferring types for `cljs.compiler`, a 1,776 line file defining the code generation phase of a production-quality compiler. Its AST format is inferred as `Op` (lines 1-12) with 22 recursive references (like lines 7, 8, 11) and 14 cases distinguished by `:op` (like lines 3, 9, 10), 5 of which have optional entries (like lines 6-8). To improve inference time, only the code emission unit tests were exercised (299 lines containing 39 assertions) which normally take 40 seconds to run, from which we generated 448 lines of types and 517 lines of specs in 2.5 minutes on a 2011 MacBook Pro (16GB RAM, 2.4GHz i5).

name of `emit-dot` to decide whether, for instance, to split `Op` into smaller type aliases or add type casts in the definition of `emit-dot` to please the type checker (some libraries require more casts than others to type check, as discussed in Section 5.2).

- Generated Clojure spec annotations (an extension discussed in Section 4.3) are more likely to accept valid input with specs enabled, even with incomplete unit tests (we enable generated specs on several libraries in Section 5.3).
- Our approach becomes more amenable to extensions improving the running time of runtime observation without significantly deteriorating annotation quality, like lazy tracking (Section 4.2).

Our general approach to generating types is separated into two phases—**collection** and **inference**. The collection phase (Section 3.1), gathers observations about a running program. This is achieved by instrumenting the program and exercising it, usually by running its unit tests, with space-efficient tracking (Section 4.1) avoiding redundant traversals of values. The inference phase (Section 3.2) uses these runtime observations to generate the final type annotations, with recursive types, optional entries, and good names.

The first pass in the inference phase generates a naive type environment from runtime observations (described in Section 3.2.1). Types are fully unrolled, appearing similar to TypeWiz’s final annotation for nodes (Figure 2)—except it is the starting point of our algorithm.

A key hypothesis in our algorithm is that functions in the same file operate on related data. Our inference is built to be used per-file and aggressively merges all apparently-related HMap types—firstly types immediately nested within each other (Section 3.2.2), and then across type aliases (Section 3.2.3). This often yields useful types in our benchmarks. This approach has limited success for libraries providing polymorphic functions, since unit tests may use sample data that are unrelated to the details of a function. Our algorithm excels with programs that mainly operate on one or two (possibly recursive) map-based data representations, which, outside of polymorphic libraries, are common in the Clojure ecosystem in our experience.

### 3 Formalism

We present  $\lambda_{\text{track}}$ , an untyped  $\lambda$ -calculus describing the essence of our approach to automatic annotations. We split our model into two phases: the collection phase collect that runs an instrumented program and collects observations, and an inference phase infer that derives type annotations from these observations that can be used to automatically annotate the program.

We define the top-level driver function `annotate` that connects both pieces. It says, given a program  $e$  and top-level variables  $\bar{x}$  to infer annotations for, return an annotation environment  $\Delta$  with possible entries for  $\bar{x}$  based on observations from evaluating an instrumented  $e$ .

$$\begin{aligned} \text{annotate} &: e, \bar{x} \rightarrow \Delta \\ \text{annotate} &= \text{infer} \circ \text{collect} \end{aligned}$$

To contextualize the presentation of these phases, we begin a running example: inferring the type of a top-level function  $f$ , that takes a map and returns its `:a` entry, based on the following usage.

```
define f = λm.(get m :a)
(f {:a 42}) => 42
```

Plugging this example into our driver function we get a candidate annotation for  $f$ :

$$\text{annotate}((f \text{ {:a 42}}), [f]) = \{f : \{:\text{a N}\} \rightarrow \text{N}\}$$

$v$	$::= n \mid k \mid [\lambda x.e, \rho]_c \mid \{\overline{k v}\} \mid c$	Values	386
$e$	$::= x \mid v \mid (\mathbf{track} \ e \ \pi) \mid \lambda x.e$ $\mid \{\overline{e \bar{e}}\} \mid (e \ \bar{e})$	Expressions	388
$\rho$	$::= \{\overline{x \mapsto \bar{v}}\}$	Runtime environments	389
$l$	$::= x \mid \mathbf{dom} \mid \mathbf{rng} \mid \mathbf{key}_m(k)$	Path Elements	390
$\pi$	$::= \bar{l}$	Paths	391
$r$	$::= \{\overline{\tau \pi}\}$	Inference results	392
$\tau, \sigma$	$::= \text{N} \mid [\tau \rightarrow \tau] \mid (\mathbf{HMap}_o^m) \mid (\bigcup \bar{\tau})$ $\mid a \mid k \mid \mathbf{K} \mid \top \mid (\mathbf{Map} \ \tau \ \tau) \mid ?$	Types	393
$\Gamma$	$::= \{\overline{x : \tau}\}$	Type environments	394
$m, o$	$::= \{\overline{k \ \tau}\}$	HMap entries	395
$A$	$::= \{\overline{a \mapsto \tau}\}$	Type alias environments	396
$\Delta$	$::= (A, \Gamma)$	Annotation environments	397

Figure 6. Syntax of Terms, Types, Inference results, and Environments for  $\lambda_{\text{track}}$

#### 3.1 Collection phase

Now that we have a high-level picture of how these phases interact, we describe the syntax and semantics of  $\lambda_{\text{track}}$ , before presenting the details of collect. Figure 6 presents the syntax of  $\lambda_{\text{track}}$ . Values  $v$  consist of numbers  $n$ , Clojure-style keywords  $k$ , closures  $[\lambda x.e, \rho]_c$ , constants  $c$ , and keyword keyed hash maps  $\{\overline{k v}\}$ .

Expressions  $e$  consist of variables  $x$ , values, functions, maps, and function applications. The special form  $(\mathbf{track} \ e \ \pi)$  observes  $e$  as related to path  $\pi$ . Paths  $\pi$  record the source of a runtime value with respect to a sequence of path elements  $l$ , always starting with a variable  $x$ , and are read left-to-right. Other path elements are a function domain  $\mathbf{dom}$ , a function range  $\mathbf{rng}$ , and a map entry  $\mathbf{key}_{\overline{k_1}}(k_2)$  which represents the result of looking up  $k_2$  in a map with keyset  $\overline{k_1}$ .

Inference results  $\{\overline{\tau \pi}\}$  are pairs of paths  $\pi$  and types  $\tau$  that say the path  $\pi$  was observed to be type  $\tau$ . Types  $\tau$  are numbers  $\text{N}$ , function types  $[\tau \rightarrow \tau]$ , ad-hoc union types  $(\bigcup \tau \ \tau)$ , type aliases  $a$ , and unknown type  $?$  that represents a temporary lack of knowledge during the inference process. Heterogeneous keyword map types  $\{\overline{k \ \tau}\}$  for now represent a series of required keyword entries—we will extend them to have optional entries in later phases.

The big-step operational semantics  $\rho \vdash e \Downarrow v ; r$  (Figure 7) says under runtime environment  $\rho$  expression  $e$  evaluates to value  $v$  with inference results  $r$ . Most rules are standard, with extensions to correctly propagate inference results  $r$ . B-Track is the only interesting rule, which instruments its fully-evaluated argument with the track metafunction.

The metafunction  $\text{track}(v, \pi) = v' ; r$  (Figure 7) says if value  $v$  occurs at path  $\pi$ , then return a possibly-instrumented  $v'$  paired with inference results  $r$  that can be immediately derived from the knowledge that  $v$  occurs at path  $\pi$ . It has a case for every kind of value. The first three cases records the number input as type  $\text{N}$ . The fourth case, for closures, returns a wrapped value resembling higher-order function contracts [9], but we track the domain and range rather than

<p>441</p> <p>442</p> <p>443</p> <p>444</p> <p>445</p> <p>446</p> <p>447</p> <p>448</p> <p>449</p> <p>450</p> <p>451</p> <p>452</p> <p>453</p> <p>454</p> <p>455</p> <p>456</p> <p>457</p> <p>458</p> <p>459</p> <p>460</p> <p>461</p> <p>462</p> <p>463</p> <p>464</p> <p>465</p>	<div style="display: flex; justify-content: space-between;"> <div style="width: 45%;"> <p><b>B-TRACK</b></p> <math display="block">\frac{\rho \vdash e \Downarrow v; r}{\text{track}(v, \pi) = v'; r'}</math> <p><b>B-CLOS</b></p> <math display="block">\rho \vdash \lambda x. e \Downarrow [\lambda x. e, \rho]_c; \{ \}</math> <p><b>B-VAL</b></p> <math display="block">\rho \vdash v \Downarrow v; \{ \}</math> </div> <div style="width: 45%;"> <p><b>B-APP</b></p> <math display="block">\frac{\begin{array}{l} \rho \vdash e_1 \Downarrow [\lambda x. e, \rho']_c; r_1 \\ \rho \vdash e_2 \Downarrow v; r_2 \\ \rho'[x \mapsto v] \vdash e \Downarrow v'; r_3 \end{array}}{\rho \vdash (e_1 e_2) \Downarrow v'; \bigcup \bar{r}_i}</math> <p><b>B-DELTA</b></p> <math display="block">\frac{\begin{array}{l} \rho \vdash e \Downarrow c; r_1 \\ \rho \vdash e' \Downarrow v; r' \\ \delta(c, \bar{v}) = v'; r_2 \end{array}}{\rho \vdash (e \bar{e}') \Downarrow v'; r \cup r'}</math> <p><b>B-VAR</b></p> <math display="block">\rho \vdash x \Downarrow \rho(x); \{ \}</math> </div> </div> <p>466</p> <p>467</p> <p>468</p> <p>469</p> <p>470</p> <p>471</p> <p>472</p> <p>473</p> <p>474</p> <p>475</p> <p>476</p> <p>477</p> <p>478</p> <p>479</p> <p>480</p> <p>481</p> <p>482</p> <p>483</p> <p>484</p> <p>485</p> <p>486</p> <p>487</p> <p>488</p> <p>489</p> <p>490</p> <p>491</p> <p>492</p> <p>493</p> <p>494</p> <p>495</p>
--	---

**Figure 7.** Operational semantics,  $\text{track}(v, \pi) = v; r$  and constants

verify them. The remaining rules case, for maps, recursively tracks each map value, and returns a map with possibly wrapped values. Immediately accessible inference results are combined and returned. A specific rule for the empty map is needed because we otherwise only rely on recursive calls to **track** to gather inference results—in the empty case, we have no data to recur on.

Now we have sufficient pieces to describe the initial collection phase of our model. Given an expression  $e$  and variables  $\bar{x}$  to track,  $\text{instrument}(e, \bar{x}) = e'$  returns an instrumented expression  $e'$  that tracked usages of  $\bar{x}$ . It is defined via capture-avoiding substitution:

$$\text{instrument}(e, \bar{x}) = e[\overline{(\text{track } x [x])} / \bar{x}]$$

Then, the overall collection phase  $\text{collect}(e, \bar{x}) = r$  says, given an expression  $e$  and variables  $\bar{x}$  to track, returns inference results  $r$  that are the results of evaluating  $e$  with instrumented occurrences of  $\bar{x}$ . It is defined as:

$$\text{collect}(e, \bar{x}) = r, \text{ where } \vdash \text{instrument}(e, \bar{x}) \Downarrow v; r$$

For our running example of collecting for the program  $(f \{ :a 42 \})$ , we instrument the program by wrapping occurrences of  $f$  with **track** with path  $[f]$ .

$$\text{instrument}((f \{ :a 42 \}), [f]) = ((\text{track } f [f]) \{ :a 42 \})$$

Then we evaluate the instrumented program and derive two inference results (colored in red for readability):

$$\vdash ((\text{track } f [f]) \{ :a 42 \}) \Downarrow 42; \{ N_{[f, \text{dom}, \text{key}(:a)], N_{[f, \text{rng}]} \}$$

Here is the full derivation:

$$\begin{aligned} &=> ((\text{track } f [f]) \{ :a 42 \}) \\ &=> (\text{track } (\text{get } (\text{track } \{ :a 42 \} [f, \text{dom}]) :a) [f, \text{rng}]) \\ &=> (\text{track } (\text{get } \{ :a 42 \}; \{ N_{[f, \text{dom}, \text{key}(:a)]} \} :a) [f, \text{rng}]) \\ &=> (\text{track } 42; \{ N_{[f, \text{dom}, \text{key}(:a)]} \} [f, \text{rng}]) \\ &=> 42; \{ N_{[f, \text{dom}, \text{key}(:a)], N_{[f, \text{rng}]} \} \end{aligned}$$

Notice that intermediate values can have inference results (colored) attached to them with a semicolon, and the final value has inference results about both  $f$ 's domain and range.

### 3.2 Inference phase

After the collection phase, we have a collection of inference results  $r$  which can be passed to the metafunction  $\text{infer}(r) = \Delta$  to produce an annotation environment:

$$\begin{aligned} \text{infer} : r &\rightarrow \Delta \\ \text{infer} &= \text{inferRec} \circ \text{toEnv} \end{aligned}$$

The first pass  $\text{toEnv}(r) = \Gamma$  generates an initial type environment from inference results  $r$ . The second pass

$$\text{squashLocal}(\Gamma) = \Delta'$$

creates individual type aliases for each HMap type in  $\Gamma$  and then merges aliases that both occur inside the same nested type into possibly recursive types. The third pass  $\text{squashGlobal}(\Delta) = \Delta'$  merges type aliases in  $\Delta$  based on their similarity.

#### 3.2.1 Pass 1: Generating initial type environment

The first pass is given in Figure 8. The entry point  $\text{toEnv}$  folds over inference results to create an initial type environment via update. This style is inspired by occurrence typing [22], from which we also borrow the concepts of paths into types.

We process paths right-to-left in update, building up types from leaves to root, before joining the fully constructed type with the existing type environment via  $\sqcup$ . The first case handles the **key** path element. The extra map of type information preserves both keyset information and any entries that might represent tags (populated by the final case of **track**, Figure 7). This information helps us avoid prematurely collapsing tagged maps, by the side condition of the HMap  $\sqcup$  case. The  $\sqcup^H$  metafunction aggressively combines two HMaps—required keys in both maps are joined and stay required, otherwise keys become optional.

The second and third update cases update the domain and range of a function type, respectively. The  $\sqcup$  case for function types joins covariantly on the domain to yield more useful annotations. For example, if a function accepts  $N$  and  $K$ , it will have type  $[N \rightarrow ?] \sqcup [K \rightarrow ?] = [(\cup N K) \rightarrow ?]$ .

551	$\sqcup : \tau, \tau \rightarrow \tau$				
552	$(\bigcup \bar{\sigma}) \sqcup \tau = (\bigcup \overline{\sigma \sqcup \tau})$	$(\text{HMap}_{o_1}^{m_1}) \sqcup^H (\text{HMap}_{o_2}^{m_2}) = (\text{HMap}_o^m)$			
553	$\tau \sqcup (\bigcup \bar{\sigma}) = (\bigcup \overline{\sigma \sqcup \tau})$	where $\text{req} = \bigcup \overline{\text{dom}(m_i)}$			
554	$? \sqcup \tau = \tau$	$\text{opt} = \bigcup \overline{\text{dom}(o_i)}$			
555	$\tau \sqcup ? = \tau$	$\overline{k^r} = \bigcap \overline{\text{dom}(m_i)} \setminus \text{opt}$			
556	$[\tau_1 \rightarrow \sigma_1] \sqcup [\tau_2 \rightarrow \sigma_2] = [\tau_1 \sqcup \tau_2 \rightarrow \sigma_1 \sqcup \sigma_2]$	$\overline{k^o} = \text{opt} \cup (\text{req} \setminus \overline{k^r})$			
557	$(\text{HMap}_{o_1}^{m_1}) \sqcup (\text{HMap}_{o_2}^{m_2}) = (\text{HMap}_{o_1}^{m_1}) \sqcup^H (\text{HMap}_{o_2}^{m_2})$	$m = \overline{\{k^r \sqcup m_i[k^r]\}}$			
558		$(k, k_i) \in m_i \Rightarrow \overline{k_{i-1}} = k_i$			
559	$\tau \sqcup \sigma = (\bigcup \tau \sigma), \text{otherwise}$	$o = \overline{\{k^o \sqcup \overline{m_i[k^o]}, o_i[k^o]\}}$			
560					
561	$\text{fold} : \forall \alpha, \beta. (\alpha, \beta \rightarrow \alpha), \alpha, \overline{\beta} \rightarrow \alpha$	$\text{update} : \Gamma, \tau_\pi \rightarrow \Gamma$			
562	$\text{fold}(f, a_0, \overline{b^n}) = a_n$	$\text{update}(\Gamma, \tau_{\pi::[\text{key}_{\langle \overline{k^r} \sigma \rangle}(k)]}) = \text{update}(\Gamma, \overline{\{k^r \sigma k \tau\}_\pi})$			
563	where $\overline{a_i} = f(\overline{a_{i-1}}, \overline{b_i})^{1 \leq i \leq n}$	$\text{update}(\Gamma, \tau_{\pi::[\text{dom}]}) = \text{update}(\Gamma, [\tau \rightarrow ?]_\pi)$			
564		$\text{update}(\Gamma, \tau_{\pi::[\text{rng}]}) = \text{update}(\Gamma, [? \rightarrow \tau]_\pi)$			
565		$\text{update}(\Gamma[x \mapsto \sigma], \tau_{[x]}) = \Gamma[x \mapsto \tau \sqcup \sigma]$			
566	$\text{toEnv} : r \rightarrow \Gamma$	$\text{update}(\Gamma, \tau_{[x]}) = \Gamma[x \mapsto \tau]$			
567	$\text{toEnv}(r) = \text{fold}(\text{update}, \{\}, r)$				
568					
569					
570					
571					

Figure 8. Definition of  $\text{toEnv}(r) = \Gamma$ 

Returning to our running example, we now want to convert our inference results

$$r = \{\mathsf{N}_{[f, \text{dom}, \text{key}:(a)]}, \mathsf{N}_{[f, \text{rng}]}\}$$

into a type environment. Via  $\text{toEnv}(r)$ , we start to trace  $\text{update}(\{\}, \mathsf{N}_{[f, \text{dom}, \text{key}:(a)]})$

### 3.2.2 Pass 2: Squash locally

We now describe the algorithm for generating recursive type aliases. The first step `squashLocal` creates recursive types from directly nested types. It folds over each type in the type environment, first creating aliases with `aliasHMap`, and then attempting to merge these aliases by `squashAll`.

A type is aliased by `aliasHMap` either if it is a union containing a HMap, or a HMap that is not a member of a union. While we will use the structure of HMaps to determine when to create a recursive type, keeping surrounding type information close to HMaps helps create more compact and readable recursive types. The implementation uses a post-order traversal via `postwalk`, which also threads an annotation environment as it applies the provided function.

Then, `squashAll` follows each alias  $a_i$  reachable from the type environment and attempts to merge it with any alias reachable from  $a_i$ . The squash function maintains a set of already visited aliases to avoid infinite loops.

The logic for merging aliases is contained in `mergeAliases`. Merging  $a_2$  into  $a_1$  involves mapping  $a_2$  to  $a_1$  and  $a_1$  to the join of both definitions. Crucially, before joining, we rename occurrences of  $a_2$  to  $a_1$ . This avoids a linear increase in the width of union types, proportional to the number of merged aliases. The running time of our algorithm is proportional to the width of union types (due to the quadratic combination of unions in the join function) and this optimization greatly

helped the running time of several benchmarks. To avoid introducing infinite types, top-level references to other aliases we are merging with are erased with the helper `f`.

The `merge?` function determines whether two types are related enough to warrant being merged. We present our current implementation, which is simplistic, but is fast and effective in practice, but many variations are possible. Aliases are merged if they are all HMaps (not contained in unions), that contain a keyword key in common, with possibly disjoint mapped values. For example, our opening example has the `:op` key mapped to either `:leaf` or `:node`, and so aliases for each map would be merged. Notice again, however, the join operator does not collapse differently-tagged maps, so they will occur recursively in the resulting alias, but separated by union.

Even though this implementation of `merge?` does not directly utilize the aliased union types carefully created by `aliasHMap`, they still affect the final types. For example, squashing  $\mathsf{T}$  in

```
(defalias T
  (U nil '{:op :node :left '{:op :leaf ...} ...}))
```

results in

```
(defalias T
  (U nil '{:op :node :left T ...} '{:op :leaf ...}))
```

rather than

```
(defalias T2 (U '{:op :node :left T ...}
  '{:op :leaf ...}))
```

```
(defalias T (U nil T2))
```

An alternative implementation of `merge?` we experimented with included computing sets of keysets for each alias, and

<pre> 661 aliasHMap : Δ, τ → (Δ, τ) 662 aliasHMap(Δ, τ) = postwalk(Δ, τ, f) 663 where f(Δ, (HMap<sup>m<sub>1</sub></sup>)) = reg(Δ, (HMap<sup>m<sub>1</sub></sup>)) 664 f(Δ, (∪ τ̄)) = reg(Δ, (∪ resolve(τ̄))), 665 if a ∈ τ̄ 666 f(Δ, τ) = (Δ, τ), otherwise 667 668 aliases : τ → ā 669 aliases(a) = [a] 670 aliases(τ(σ̄)) = ∪ aliases(σ̄) 671 672 postwalk : Δ, τ, (Δ, τ → (Δ, τ)) → (Δ, τ) 673 postwalk(Δ<sub>0</sub>, τ(σ̄<sup>n</sup>), w) = w(Δ<sub>n</sub>, τ(σ̄<sup>n</sup>)) 674 where (Δ<sub>i</sub>, σ̄<sup>i</sup>) = postwalk(Δ<sub>i-1</sub>, σ<sub>i</sub>, w) 675 676 mergeAliases : Δ, ā → Δ 677 mergeAliases(Δ, []) = Δ 678 mergeAliases(Δ, [a<sub>1</sub>...a<sub>n</sub>]) = Δ[a<sub>1</sub> ↦ a<sub>1</sub>][a<sub>1</sub> ↦ σ] 679 where σ = ∪ f(resolve(Δ, a<sub>i</sub>))[a<sub>1</sub>/a<sub>i</sub>] 680 f(a') = (∪), if a' ∈ ā 681 f((∪ τ̄)) = (∪ f(τ̄)) 682 f(τ) = τ, otherwise 683 684 squashLocal : Γ → Δ 685 squashLocal(Γ) = <b>fold</b>(h, ({}, {}), Γ) 686 where h(Δ, x : τ) = Δ<sub>2</sub>[x ↦ τ<sub>2</sub>] 687 where (Δ<sub>1</sub>, τ<sub>1</sub>) = aliasHMap(Δ, τ) 688 (Δ<sub>2</sub>, τ<sub>2</sub>) = squashAll(Δ<sub>1</sub>, τ<sub>1</sub>) </pre>	<pre> reg : Δ, τ → (Δ, τ) reg(Δ, τ) = (Δ[a ↦ τ], a), where a is fresh resolve : Δ, τ → τ resolve(Δ, a) = resolve(Δ[a]) resolve(Δ, τ) = τ, otherwise squashAll : Δ, τ → Δ squashAll(Δ<sub>0</sub>, τ) = Δ<sub>n</sub> where ā<sup>n</sup> = aliases(τ) Δ<sub>i</sub> = squash(Δ<sub>i-1</sub>, [a<sub>i</sub>], []) squash : Δ, ā, ā → Δ squash(Δ, [], d) = Δ squash(Δ, a<sub>1</sub> :: w, d) = squash(Δ', w ∪ as, d ∪ {a<sub>1</sub>}) where as = aliases(Δ[a<sub>1</sub>]) \ d ap = d \ {a<sub>1</sub>} f(Δ, a<sub>2</sub>) = if ¬merge?(resolve(Δ, a)), then Δ else mergeAliases(Δ, ā<sub>i</sub>) Δ' = if a ∈ d, then Δ, else fold(f, Δ, ap ∪ as) merge? : τ̄ → <b>Bool</b> merge?((HMap<sup>m<sub>i</sub></sup>)) = ∃k.(k, k<sub>i</sub>) ∈ m<sub>i</sub> merge?(τ̄) = <b>F</b>, otherwise </pre>	<pre> 716 717 718 719 720 721 722 723 724 725 726 727 728 729 730 731 732 733 734 735 736 737 738 739 740 741 742 743 744 745 746 747 748 749 750 751 752 753 754 755 756 757 758 759 760 761 762 763 764 765 766 767 768 769 770 </pre>
--	--	--

Figure 9. Definition of squashLocal(Γ) = Δ

merging if the keysets overlapped. This, and many of our early experimentations, required expensive computations of keyset combinations and traversals over them that could be emulated with cruder heuristics like the current implementation.

### 3.2.3 Pass 3: Squash globally

The final step combines aliases without restriction on whether they occur “together”. This step combines type information between different positions (such as in different arguments or functions) so that any deficiencies in unit testing coverage are massaged away.

The squashGlobal function is the entry point in this pass, and is similar in structure to the previous pass. It first creates aliases for each HMap via aliasSingleHMap. Then, HMap aliases are grouped and merged in squashHorizontally.

The aliasSingleHMap function first traverses the type environment to create HMap aliases via singleHMap, and binds the resulting environment as Δ'. Then, alias environment entries are updated with f, whose first case prevents re-aliasing a top-level HMap, before we call singleHMap (singleHMap's

second argument accepts both x and a). The τ(σ̄) syntax represents a type τ whose constructor takes types σ̄.

After that, squashHorizontally creates groups of related aliases with groupSimilarReq. Each group contains HMap aliases whose required keysets are similar, but are never differently-tagged. The code creates a map r from keysets to groups of HMap aliases with that (required) keyset. Then, for every keyset  $\bar{k}$ , similarReq adds aliases to the group whose keysets are a subset of  $\bar{k}$ . The number of missing keys permitted is determined by thres, for which we do not provide a definition. Finally, remDiffTag removes differently-tagged HMaps from each group, and the groups are merged via mergeAliases as before.

### 3.2.4 Implementation

Further passes are used in the implementation. In particular, we trim unreachable aliases and remove aliases that simply point to another alias (like  $a_2$  in mergeAliases) between each pass.

771  $\text{req} : \Delta, a \rightarrow m$  826  
 772  $\text{req}(\Delta, a) = \text{req}(\Delta, a)$  827  
 773  $\text{req}(\Delta, (\text{HMap}_o^m)) = m$  828  
 774 829  
 775 830  
 776 831  
 777  $\text{squashHorizontally} : \Delta \rightarrow \Delta$  832  
 778  $\text{squashHorizontally}(\Delta) = \text{fold}(\text{mergeAliases}, \Delta, \text{groupSimilarReq}(\Delta))$  833  
 779 834  
 780 835  
 781  $\text{squashGlobal} : \Delta \rightarrow \Delta$  836  
 782  $\text{squashGlobal} = \text{squashHorizontally} \circ \text{aliasSingleHMap}$  837  
 783 838  
 784 839  
 785 840  
 786 841  
 787  $\text{groupSimilarReq} : \Delta \rightarrow \bar{a}$  842  
 788  $\text{groupSimilarReq}(\Delta) = [\bar{a} | \bar{k} \in \text{dom}(r), \bar{a} = \text{remDiffTag}(\text{similarReq}(\bar{k}))]$  843  
 789  $\text{where } r = \{(\bar{k}, \bar{a}) | (\text{HMap}_o^{\{\bar{k} \tau\}}) \in \text{rng}(\Delta[A]), \bar{a} = \text{matchingReq}(\bar{k})\}$  844  
 790  $\text{matchingReq}(\bar{k}) = [a | (a, (\text{HMap}_o^m)) \in \Delta]$  845  
 791  $\text{similarReq}(\bar{k}) = [a | \bar{k}'^n \subseteq \bar{k}^m, m - n \leq \text{thres}(m), a \in r[\bar{k}']]$  846  
 792  $\text{remDiffTag}(\bar{a}) = [a' | a' \in \bar{a}, \text{if } (k, k') \in \text{req}(\Delta, a') \text{ and } \bigvee (k, k'') \in \text{req}(\Delta, a) \text{ then } \overline{k' = k''}]$  847  
 793 848  
 794 849

Figure 10. Definition of  $\text{squashGlobal}(\Delta) = \Delta'$ 

## 4 Extensions

### 4.1 Space-efficient tracking

To reduce the overhead of runtime tracking, we can borrow the concept of “space-efficient” contract checking from the gradual typing literature [13]. Instead of tracking just one path at once, a space-efficient implementation of track threads through a set of paths. When a tracked value flows into another tracked position, we extract the unwrapped value, and then our new tracked value tracks the paths that is the set of the old paths with the new path.

To model this, we introduce a new kind of value  $[e, \rho]_c^{\bar{\pi}}$  that tracks old value  $v$  as new value  $[e, \rho]_c$  with the paths  $\bar{\pi}$ . Proxy expressions are introduced when tracking functions, where instead of just returning a new wrapped function, we return a proxy. We can think of function proxies as a normal function with some extra metadata, so we can reuse the existing semantics for function application—in fact we can support space-efficient function tracking just by extending **track**.

We present the extension in Figure 11. The first two **track** rules simply make inference results for each of the paths. The next rule says that a bare closure reduces to a proxy that tracks the domain and range of the closure with respect to the list of paths. Attached to the proxy is everything needed to extend it with more paths, which is the role of the final rule. It extracts the original closure from the proxy and creates a new proxy with updated paths via the previous rule.

795  $v ::= \dots \mid [\lambda x.e, \rho]_c^{\bar{\pi}}$  Values 850  
 796 851  
 797  $\text{track}(n, \bar{\pi}) = n ; \bigcup \{\mathbf{N}_{\bar{\pi}}\}$  852  
 798  $\text{track}(k, \bar{\pi}) = k ; \bigcup \{\mathbf{K}_{\bar{\pi}}\}$  853  
 799  $\text{track}([\lambda x.e, \rho]_c, \bar{\pi}) = [e', \rho]_c^{\bar{\pi}} ; \{\}$  854  
 800  $\text{where } y \text{ is fresh,}$  855  
 801  $e' = \lambda y. (\text{track } ((\lambda x.e) (\text{track } y \pi :: [\text{dom}]))$  856  
 802  $\pi :: [\text{rng}])$  857  
 803  $\text{track}([e', \rho']_c^{\bar{\pi}'}, [\lambda x.e, \rho]_c, \bar{\pi}) = \text{track}([\lambda x.e, \rho]_c, \bar{\pi} \cup \bar{\pi}')$  858  
 804 859

Figure 11. Space-efficient tracking extensions (changes)

### 4.2 Lazy tracking

Building further on the extension of space-efficient functions, we apply a similar idea for tracking maps. In practice, eagerly walking data structures to gather inference results is expensive. Instead, waiting until a data structure is used and tracking its contents lazily can help ease this tradeoff, with the side-effect that fewer inference results are discovered.

Figure 12 extends our system with lazy maps. We add a new kind of value  $\{\bar{k} v\}^{\{k' \overline{m \bar{\pi}}\}}$  that wraps a map  $\{\bar{k} v\}$  with tracking information. Keyword entries  $k'$  are associated with pairs of type information  $m$  with paths  $\bar{\pi}$ . The first **track** rule demonstrates how to create a lazily tracked map. We calculate the possibly tagged entries in our type information in advance, much like the equivalent rule in Figure 7, and store them for later use. Notice that non-keyword entries are



881  $v ::= \dots \mid \overline{\{k \ v\}}^{\overline{\{k \ m \ \bar{\pi}\}}}$  Values  
 882  
 883  $\text{track}(\overline{\{k \ k' \ k'' \ v\}}, \bar{\pi}) = \overline{\{k \ k' \ k'' \ v\}}^{\overline{\{k \ t \ k'' \ t\}}}; \{\}$   
 884 where  $t = \{\{k \ k' \ k'' \ ?\} \bar{\pi}\}$   
 885  $\text{track}(\overline{\{k \ v\}}^{\overline{\{k' \ m \ \bar{\pi}'\}}}, \bar{\pi}) = \overline{\{k \ v\}}^{\overline{\{k' \ m \ (\bar{\pi} \cup \bar{\pi}')\}}}; \{\}$   
 886  
 887  $\delta(\text{assoc}, \overline{\{k \ v\}}^{\overline{\{k' \ t', \ k'' \ t\}}}, k', v') = \overline{\{k \ v\}}^{\overline{\{k' \ t' \ v'\}}}; \{\}$   
 888  $\delta(\text{assoc}, \overline{\{k \ v\}}^{\overline{\{k'' \ t\}}}, k', v') = \overline{\{k \ v\}}^{\overline{\{k' \ t' \ v'\}}}; \{\}$   
 889  $\delta(\text{get}, \overline{\{k \ v, \ k' \ v'\}}^{\overline{\{k \ t, \ k'' \ t'\}}}, k) = \text{track}(v, \bar{\pi})$   
 890 where  $\bar{\pi} = [\pi :: [\text{key}_m(k) \mid (m, \bar{\pi}) \in t, \pi \in \bar{\pi}]]$   
 891  $\delta(\text{get}, \overline{\{k \ v, \ k' \ v'\}}^{\overline{\{k'' \ t'\}}}, k) = v$   
 892  $\delta(\text{dissoc}, \overline{\{k \ v, \ k' \ v'\}}^{\overline{\{k \ t, \ k'' \ t'\}}}, k) = \overline{\{k' \ v'\}}^{\overline{\{k'' \ t'\}}}; \{\}$   
 893  $\delta(\text{dissoc}, \overline{\{k \ v, \ k' \ v'\}}^{\overline{\{k'' \ t'\}}}, k) = \overline{\{k' \ v'\}}^{\overline{\{k'' \ t'\}}}; \{\}$   
 894

895 **Figure 12.** Lazy tracking extensions (changes)  
 896  
 897  
 898

899 not yet traversed, and thus no inference results are derived  
 900 from them. The second **track** rule adds new paths to track.

901 The subtleties of lazily tracking maps lie in the  $\delta$  rules. The  
 902 assoc and dissoc rules ensure we no longer track overwritten  
 903 entries. Then, the get rules perform the tracking that was  
 904 deferred from the **track** rule for maps in Figure 7 (if the  
 905 entry is still tracked).

906 In our experience, some combination of lazy and eager  
 907 tracking of maps strikes a good balance between perfor-  
 908 mance overhead and quantity of inference results. Intuitively,  
 909 if a function does not access parts of its argument, they  
 910 should not contribute to that function’s type signature. How-  
 911 ever, our inference algorithm combines information *across*  
 912 function signatures to deduce useful, recursive type aliases.  
 913 Some eager tracking helps normalize the quality of function  
 914 annotations with respect to unit test coverage.

915 For example, say functions  $f$  and  $g$  operate on the same  
 916 types of (deeply nested) arguments, and  $f$  has complete test  
 917 coverage (but does not traverse all of its arguments), and  
 918  $g$  has incomplete test coverage (but fully traverses its ar-  
 919 guments). Eagerly tracking  $f$  would give better inference  
 920 results, but lazily tracking  $g$  is more efficient. Forcing sev-  
 921 eral layers of tracking helps strike this balance, which our  
 922 implementation exposes as a parameter.

923 This can be achieved in our formal system by adding fuel  
 924 arguments to **track** that contain depth and breadth tracking  
 925 limits, and defer to lazy tracking when out of fuel.  
 926

### 927 4.3 Automatic contracts with clojure.spec

928 While we originally designed our tool to generate Typed  
 929 Clojure annotations, it also supports generating “specs” for  
 930 clojure.spec, Clojure’s runtime verification system. There  
 931 are key similarities between Typed Clojure and clojure.spec,  
 932 such as extensive support for potentially-tagged keyword  
 933 maps, however spec features a global registry of names via  
 934 **s/def** and an explicit way declare unions of maps with a  
 935

936 common dispatch key in s/multi-spec. These require dif-  
 937 ferences in both type and name generation.

938 The following generated specs correspond to the first Op  
 939 case of Figure 5 (lines 2-8).  
 940

```

941 1 (defmulti op-multi-spec :op) ;dispatch on :op key
942 2 (defmethod op-multi-spec :binding ;match :binding
943 3   [_] ;s/keys matches keyword maps
944 4   (s/keys :req-un [::op ...] ;required keys
945 5     :opt-un [::column ...])) ;optional keys
946 6 (s/def ::op #{:js :let ...}) ;:op key maps to keywords
947 7 (s/def ::column int?) ;:column key maps to ints
948 8 ; register ::Op as union dispatching on :op entry
949 9 (s/def ::Op (s/multi-spec op-multi-spec :op))
950 10 ; emit's first argument :ast has spec ::Op
951 11 (s/def emit :args (s/cat :ast ::Op) :ret nil?)
  
```

## 952 5 Evaluation

953 We performed a quantitative evaluation of our algorithm on  
 954 several open source programs.  
 955

956 **startrek-clojure** A reimplementation of a Star Trek text  
 957 adventure game, created as a way to learn Clojure.  
 958

959 **math.combinatorics** The core library for common com-  
 960 binatorial functions on collections, with implementations  
 961 based on Knuth’s Art of Computer Programming, Volume 4.  
 962

963 **fs** A Clojure wrapper library over common file-system op-  
 964 erations.  
 965

966 **data.json** A library for working with JSON.  
 967

968 **mini.occ** A model of occurrence typing by an author of  
 969 the current paper. It utilizes three mutually recursive ad-hoc  
 970 structures to represent expressions, types, and propositions.  
 971

972 **cljs.compiler** ClojureScript (CLJS) is a Clojure variant that  
 973 runs on JavaScript virtual machines. We infer types for its  
 974 compiler (written in Clojure) which emits JavaScript from a  
 975 recursively defined map-based abstract syntax tree format.  
 976

977 Our approach is to carry out three experiments.  
 978

### 979 5.1 Experiment 1: Manual inspection

980 In the first experiment, we generate types for a program  
 981 and manually inspect the resulting annotations. We follow  
 982 the criteria outlined in Section 2 to judge the quality of our  
 983 output, namely:  
 984

- 985 • using recognizable names,
- 986 • favoring compact annotations, and
- 987 • not overspecifying types.

988 Since it is our largest benchmark with 448 lines of gen-  
 989 erated type annotations, we concentrated on manually in-  
 990 specting cljs.compiler’s generated types. We commented  
 991 (Section 2) on how successfully Figure 5 follows our goals.  
 992

Here, we further elaborate on Figure 5 (with line references), and contextualize them with the elided annotations.

One remarkable success in the generated types was the automatic inference `Op` (lines 1-12) with 14 distinct cases, and other features described in Figure 5. Further investigation reveals that the compiler actually features 36 distinct AST nodes—unsurprisingly, 39 assertions was not sufficient test coverage to discover them all. However, because of the recognizable name and organization of `Op`, it's clear where to add the missing nodes if no further tests are available.

A failure of `cljs.compiler`'s generated types was `HMap49305`. It clearly fails to be a recognizable name. However, all is not lost: the compactness and recognizable names of other adjacent annotations makes it plausible for a programmer with some knowledge of the AST representation to recover. In particular 13/14 cases in `Op` have entries from `:env` to `HMap49305`, (like lines 9 and 10), and the only exception (line 7) maps to `ColumnLineContextMap`. From this information the user can decide to combine these aliases.

Several instances of overspecification are evident, such as the `:statements` entry of a `:do` AST node being inferred as an always-empty vector (line 11). In some ways, this is useful information, showing that test coverage for `:do` nodes could be improved. To fix the annotation, we could rerun the tool with better tests. If no such test exists, we would have to fall back to reverse-engineering code to identify the correct type of `:statements`, which is `(Vec Op)`.

Finally, 19 functions in `cljs.compiler` are annotated to take or return `Op` (like lines 23, 24). This kind of alias reuse enables annotations to be relatively compact (only 16 type aliases are used by the 49 functions that were exercised).

## 5.2 Experiment 2: Changes needed to type check

In this experiment, we first generated types with our algorithm by running the tests, then amended the program so that it type checks. We observed some frequent reasons for why changes were needed (summarized in Figure 13).

**Uncalled functions** Some functions not called at all in the unit tests. This results in very general type annotations that need manual changes to be useful. For example, the `startrek-clojure` game has several exit conditions, one of which is running out of time. Since the tests do not specifically call this function, nor play the game long enough to invoke this condition, no useful type is inferred.

```
(ann game-over-out-of-time AnyFunction)
```

In this case, minimal effort is needed to amend this type signature: the appropriate type alias already exists:

```
(defalias CurrentKlingonsCurrentSectorEnterpriseMap
  (HMap :mandatory
    {:current-klingons (Vec EnergySectorMap),
     :current-sector (Vec Int), ...}
    :optional {:lrs-history (Vec Str)}))
```

So we amend the signature as

```
(ann game-over-out-of-time
  [(Atom1 CurrentKlingonsCurrentSectorEnterpriseMap)
   -> Boolean])
```

**Over-precision** Returns of function types are often too restrictive, as the unit tests may have not been complete.

There are several instances of this in `math.combinatorics`. The `all-different?` function takes a collection and returns true only if the collection contains distinct elements. As evidenced in the generated type, the tests exercise this functions with collections of integers, atoms, keywords, and characters.

```
(ann all-different?
  [(Coll (U Int (Atom1 Int) ':a ':b Character))
   -> Boolean])
```

In our experience, the union is very rarely a good candidate for a Typed Clojure type signature, so a useful heuristic to improve the generated types would be to upcast such unions to a more permissive type, like `Any`. When we performed that case study, we did not yet add that heuristic to our tool, so in this case, we manually amend the signature as

```
(ann all-different? [(Coll Any) -> Boolean])
```

Another example of overprecision is the generated type of `initial-perm-numbers` a helper function taking a *frequency map*—a hash map from values to the number of times they occur—which is the shape of the return value of the `core` `frequencies` function.

The generated type shows only a frequency map where the values are integers are exercised.

```
(ann initial-perm-numbers
  [(Map Int Int) -> (Coll Int)])
```

A more appropriate type instead takes `(Map Any Int)`. In many examples of overprecision, while the generated type might not be immediately useful to check programs, they serve as valuable starting points and also provide an interesting summary of test coverage.

**Missing polymorphism** We do not attempt to infer polymorphic function types, so these amendments are expected. However, it is useful to compare the optimal types with our generated ones.

For example, the `remove-nth` function in `math.combinatorics` returns a functional delete operation on its argument. Here we can see the tests only exercise this function with collections of integers.

```
(ann remove-nth [(Coll Int) Int -> (Vec Int)])
```

However, the overall shape of the function is intact, and the manually amended type only requires a few keystrokes.

```
(ann remove-nth
  (All [a] [(Coll a) Int -> (Vec a)]))
```

Lib	LOC	GT	LA	MD	C	I	P	L	S	O	U	N	V	R	K	F	H	LS	RS	IT	MS	UT
sc	166	133	3	70/41	5	0	0	2	13	1	5	1	1	2	0	0	0	25	0	10	0	Y
mc	923	395	147	124/120	23	1	11	19	2	5	0	9	3	2	4	1	3	601	0	320	0	Y
fs	588	157	1	119/86	50	0	0	2	3	4	4	11	2	9	0	0	0	543	0	215	0	Y
dj	528	168	9	94/125																		
mo	530	49	1	46/26																		
cc	1776	448	4	N/A																		

**Figure 13.** The number of type annotations generated for each program: Lib = Abbreviated library names in the order we introduce them on page 9, LOC = Number of lines of code we generate types for, GT = Total number of lines of generated types after running our tool, LA = The number of local annotations generated by our tools. *Number of manual changes needed to type check, and why they were needed:* MD = Lines added/removed diff from git comparing initial generated types to the manual amendments needed to type check with Typed Clojure (unless it was too difficult to port), C = Casts, I = Instantiation, P = Polymorphic annotation, L = Local annotation, S = Work around type system Shortcoming, O = Overprecise argument type, U = Uncalled function due to bad test coverage, N = Add No-check annotation to skip checking function, V = Add Variable arity argument type, R = Overprecise return type, K = Add Keyword argument types, F = Added filter annotation, H = Erase/upcast HVec annotation. *Generated specs:* LS = Number of lines of spec generated, RS = No. recursive specs, IT = No. instance testing specs, MS = Useful map types, UT = Passed unit tests with specs enabled.

Similarly, `iter-perm` could be polymorphic, but its type is generated as

```
(ann iter-perm [(Vec Int) -> (U nil (Vec Int))])
```

We decided this function actually works over any number, and bounded polymorphism was more appropriate, encoding the fact that the elements of the output collection are from the input collection.

```
(ann iter-perm
  (All [a]
    [(Vec (I a Num)) -> (U nil (Vec (I a Num)))]))
```

**Missing argument counts** Often, variable argument functions are given very precise types. Our algorithm does not apply any heuristics to approximate variable arguments – instead we emit types that reflect only the arities that were called during the unit tests.

A good example of this phenomenon is the type inferred for the `plus` helper function from `math.combinatorics`. From the generated type, we can see the tests exercise this function with 2, 6, and 7 arguments.

```
(ann plus (IFn [Int Int Int Int Int Int Int -> Int]
               [Int Int Int Int Int Int -> Int]
               [Int Int -> Int]))
```

Instead, `plus` is actually variadic and works over any number of arguments. It is better annotated as the following, which is easy to guess based on both the annotated type and manually viewing the function implementation.

```
(ann plus [Int * -> Int])
```

A similar issue occurs with `mult`.

```
(ann mult [Int Int -> Int]) ;; generated
(ann mult [Int * -> Int])   ;; amended
```

A similar issue is inferring keyword arguments. Clojure implements keyword arguments with normal variadic arguments. Notice the generated type for `lex-partitions-H`, which takes a fixed argument, followed by some optional integer keyword arguments.

```
(ann lex-partitions-H
  (IFn [Int -> (Coll (Coll (Vec Int)))]
    [Int ':min Int ':max Int
     -> (Coll (Coll (Coll Int)))]))
```

While the arity of the generated type is too specific, we can conceivably use the type to help us write a better one.

```
(ann lex-partitions-H
  [Int & :optional {:min Int :max Int}
  -> (Coll (Coll (Coll Int)))])
```

**Weaknesses in Typed Clojure** We encountered several known weaknesses in Typed Clojure's type system that we worked around. The most invasive change needed was in `startrek-clojure`, which strongly updated the global mutable configuration map on initial `play`. We instead initialized the map with a dummy value when it is first created.

`cljs.compiler` uses many polymorphic idioms that Typed Clojure is poor at checking, so we deemed it too difficult to attempt to type check. In particular, there are many usages of the core functions `get-in` and `update-in` (functions that deeply lookup and manipulate maps) which are not even assigned types in Typed Clojure. Many function definitions would need to be ignored by the type checker to work around this. Furthermore, many manual instantiations would be needed to check transducers and polymorphic functions passed to other polymorphic functions.

### 5.3 Experiment 3: Specs pass unit tests

Our final experiment uses our tool to generate specs (4.3) instead of types. Specs are checked at runtime, so to verify the utility of generated specs, we enable spec checking while rerunning the unit tests that were used in the process of creating them.

At first this might seem like a trivial property, but it serves as a valuable test of our inference algorithm. The aggressive merging strategies to minimize aliases and maximize recognizability, while unsound transformations, are based on hypotheses about Clojure idioms and how Clojure programs are constructed. If, hypothetically, we generated singleton specs for numbers like we do for keywords and did not eventually upcast them to `number?`, the specs might be too strict to pass its unit tests. Some function specs also perform generative testing based on the argument and return types provided. If we collapse a spec too much and include it in such a spec, it might feed a function invalid input.

Thankfully, we avoid such pitfalls, and so our generated specs pass their tests for the benchmarks we tried. The right of Figure 13 shows our preliminary results. All inferred specs pass the unit tests when enforced, which tells us they are at least well formed. Since hundreds of invariants are checked, we can also be more confident that the specs are useful.

## 6 Related work

**Automatic annotations** There are two common implementation strategies for automatic annotation tools. The first strategy, “ruling-out” (for invariant detection), assumes all invariants are true and then use runtime analysis results to rule out impossible invariants. The second “building-up” strategy (for dynamic type inference) assumes nothing and uses runtime analysis results to build up invariant/type knowledge.

Examples of invariant detection tools include Daikon [8], DIDUCE [11], and Carrot [19], and typically enhance statically typed languages with more expressive types or contracts. Examples of dynamic type inference include our tool, Rubydust [4], JSTrace [20], and TypeDevil [18], and typically target untyped languages.

Both strategies have different space behavior with respect to representing the set of known invariants. The ruling-out strategy typically uses a lot of memory at the beginning, but then can free memory as it rules out invariants. For example, if `odd(x)` and `even(x)` are assumed, observing `x = 1` means we can delete and free the memory recording `even(x)`. Alternatively, the building-up strategy uses the least memory storing known invariants/types at the beginning, but increases memory usage as more the more samples are collected. For example, if we know `x : Bottom`, and we observe `x = "a"` and `x = 1` at different points in the program, we must use more memory to store the union `x : String ∪ Integer` in our set of known invariants.

**Daikon** Daikon can reason about very expressive relationships between variables using properties like ordering ( $x < y$ ), linear relationships ( $y = ax + b$ ), and containment ( $x \in y$ ). It also supports reasoning with “derived variables” like fields ( $x.f$ ), and array accesses ( $a[i]$ ). Typed Clojure’s dynamic inference can record heterogeneous data structures like vectors and hash-maps, but otherwise cannot express relationships between variables.

There are several reasons for this. The most prominent is that Daikon primarily targets Java-like languages, so inferring simple type information would be redundant with the explicit typing disciplines of these languages. On the other hand, the process of moving from Clojure to Typed Clojure mostly involves writing simple type signatures without dependencies between variables. Typed Clojure recovers relevant dependent information via occurrence typing [22], and gives the option to manually annotate necessary dependencies in function signatures when needed.

**Other Annotation Tools** Static analyzers TSInfer [15], Typete [12] and Pytype [10] automatically annotate JavaScript and Python code, respectively. We were unable to install TSInfer, and unable to run Typete without a runtime error. Pytype inferred nodes in Figure 1 as `(? -> int)`—our tool generates a compact recursive type (Figure 3). It inferred a class-based translation of Figure 1 similarly—fields `left`, `right`, and `val` were inferred as `Any`, and method nodes as `(self -> int)` on `Leaf`, and `(self -> Any)` on `Node`.

NoRegrets [17] uses dynamic analysis to learn how a program is used, and automatically runs the tests of downstream projects to improve test coverage. Their *dynamic access paths* represented as a series of *actions* are analogous to our paths of path elements.

## 7 Conclusion

This paper shows how to generate recursive heterogeneous type annotations for untyped programs that use plain data. We use a novel algorithm to “squash” the observed structure of program values into named recursive types suitable for optional type systems, all without the assistance of record, structure, or class definitions. We test this approach on thousands of lines of Clojure code, optimizing generated annotations for programmer comprehensibility over soundness.

In our experience, our guidelines to automatically name, group, and reuse types yield insightful annotations for those with some familiarity with the original programs, even if the initial annotations are imprecise, incomplete, and always require some changes to type check. Most importantly, many of these changes will involve simply rearranging or changing parts of existing annotations, so programmers are no longer left alone with the daunting task of reverse-engineering such programs completely from scratch.

## References

- [1] 2018. *Flow*. <https://flow.org>
- [2] 2018. *Hack Language*. <https://hacklang.org>
- [3] 2018. *TypeScript*. <https://www.typescriptlang.org>
- [4] Jong-hoon (David) An, Avik Chaudhuri, Jeffrey S. Foster, and Michael Hicks. 2011. Dynamic Inference of Static Types for Ruby. *SIGPLAN Not.* 46, 1 (Jan. 2011), 459–472. <https://doi.org/10.1145/1925844.1926437>
- [5] Ambrose Bonnaire-Sergeant, Rowan Davies, and Sam Tobin-Hochstadt. 2016. Practical Optional Types for Clojure. In *European Symposium on Programming Languages and Systems*. Springer, 68–94.
- [6] Gilad Bracha. 2004. Pluggable type systems. In *OOPSLA workshop on revival of dynamic languages*, Vol. 4.
- [7] Dropbox. [n. d.]. *PyAnnotate*. <https://github.com/dropbox/pyannotate>
- [8] Michael D Ernst, Jake Cockrell, William G Griswold, and David Notkin. 2001. Dynamically discovering likely program invariants to support program evolution. *IEEE Transactions on Software Engineering* 27, 2 (2001), 99–123.
- [9] Robert Bruce Findler and Matthias Felleisen. 2002. Contracts for higher-order functions. In *ACM SIGPLAN Notices*, Vol. 37. ACM, 48–59.
- [10] Google. [n. d.]. *PyType*. <https://github.com/google/pytype>
- [11] Sudheendra Hangal and Monica S Lam. 2002. Tracking down software bugs using automatic anomaly detection. In *Software Engineering, 2002. ICSE 2002. Proceedings of the 24rd International Conference on*. IEEE, 291–301.
- [12] Mostafa Hassan, Caterina Urban, Marco Eilers, and Peter Müller. 2018. MaxSMT-Based Type Inference for Python 3. In *Computer Aided Verification*, Hana Chockler and Georg Weissenbacher (Eds.). Springer International Publishing, Cham, 12–19.
- [13] David Herman, Aaron Tomb, and Cormac Flanagan. 2010. Space-efficient Gradual Typing. *Higher Order Symbol. Comput.* 23, 2 (June 2010), 167–189. <https://doi.org/10.1007/s10990-011-9066-z>
- [14] Rich Hickey. 2008. The Clojure programming language. In *Proc. DLS*.
- [15] Erik Krogh Kristensen and Anders Møller. 2017. Inference and Evolution of TypeScript Declaration Files. In *International Conference on Fundamental Approaches to Software Engineering*. Springer, 99–115.
- [16] Jukka Lehtosalo. [n. d.]. *mypy*. <http://mypy-lang.org/>
- [17] Gianluca Mezzetti, Anders Møller, and Martin Toldam Torp. 2018. Type Regression Testing to Detect Breaking Changes in Node.js Libraries. In *Proc. 32nd European Conference on Object-Oriented Programming (ECOOP)*.
- [18] Michael Pradel, Parker Schuh, and Koushik Sen. 2015. TypeDevil: Dynamic type inconsistency analysis for JavaScript. In *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, Vol. 1. IEEE, 314–324.
- [19] Brock Pytlik, Manos Renieris, Shriram Krishnamurthi, and Steven P Reiss. 2003. Automated fault localization using potential invariants. *arXiv preprint cs/0310040* (2003).
- [20] Claudiu Saftoiu. 2010. *JSTrace: Run-time type discovery for JavaScript*. Technical Report. Technical Report CS-10-05, Brown University.
- [21] Uri Shaked. 2018. *TypeWiz*. <https://github.com/urish/typewiz>
- [22] Sam Tobin-Hochstadt and Matthias Felleisen. 2010. Logical Types for Untyped Languages. In *Proc. ICFP (ICFP '10)*.

1376  
1377  
1378  
1379  
1380  
1381  
1382  
1383  
1384  
1385  
1386  
1387  
1388  
1389  
1390  
1391  
1392  
1393  
1394  
1395  
1396  
1397  
1398  
1399  
1400  
1401  
1402  
1403  
1404  
1405  
1406  
1407  
1408  
1409  
1410  
1411  
1412  
1413  
1414  
1415  
1416  
1417  
1418  
1419  
1420  
1421  
1422  
1423  
1424  
1425  
1426  
1427  
1428  
1429  
1430