

# 1 Transient Typechecks are (Almost) Free

2 **Richard Roberts** 

3 School of Design, Victoria University of Wellington  
4 rykardo.r@gmail.com

5 **Stefan Marr** 

6 School of Computing, University of Kent  
7 s.marr@kent.ac.uk

8 **Michael Homer** 

9 School of Engineering and Computer Science, Victoria University of Wellington  
10 mwh@ecs.vuw.ac.nz

11 **James Noble** 

12 School of Engineering and Computer Science, Victoria University of Wellington  
13 kjx@ecs.vuw.ac.nz

---

## 14 Abstract

15 Transient gradual typing imposes run-time type tests that typically cause a linear slowdown in  
16 programs' performance. This performance impact discourages the use of type annotations because  
17 adding types to a program makes the program slower. A virtual machine can employ standard just-  
18 in-time optimizations to reduce the overhead of transient checks to near zero. These optimizations  
19 can give gradually-typed languages performance comparable to state-of-the-art dynamic languages,  
20 so programmers can add types to their code without affecting their programs' performance.

21 **2012 ACM Subject Classification** Software and its engineering → Just-in-time compilers; Software  
22 and its engineering → Object oriented languages; Software and its engineering → Interpreters

23 **Keywords and phrases** dynamic type checking, gradual types, optional types, Grace, Moth, object-  
24 oriented programming

25 **Digital Object Identifier** 10.4230/LIPIcs.ECOOP.2019.15

26 **Funding** This work is supported by the Royal Society of New Zealand Marsden Fund

## 27 **1** Introduction

28 *“It is a truth universally acknowledged, that a dynamic language in possession of a*  
29 *good user base, must be in want of a type system.”*  
30 with apologies to Jane Austen.

31 Dynamic languages are increasingly prominent in the software industry. Building on  
32 the pioneering work of Self [20], much work in academia and industry has gone into making  
33 them more efficient [13, 14, 66, 24, 23, 25]. Just-in-time compilers have, for example, turned  
34 JavaScript from a naïvely interpreted language barely suitable for browser scripting, into  
35 a highly efficient ecosystem, eagerly adopted by professional programmers for a very wide  
36 range of tasks [44].

37 A key advantage of these dynamic languages is the flexibility offered by the lack of a  
38 static type system. From the perspective of many computer scientists, software engineers,  
39 and computational theologians, this flexibility has the disadvantage that programs without  
40 types are more difficult to read, to understand, and to analyze than programs with types.  
41 Gradual Typing aims to remedy this disadvantage, adding types to dynamic languages while  
42 maintaining their flexibility [16, 48, 50].

43 There is a spectrum of different approaches to gradual typing [22, 28]. At one end — “pluggable  
44 types” as in Strongtalk [17] or “erasure semantics” as in TypeScript [8] — all types are erased



© Richard Roberts, Stefan Marr, Michael Homer, James Noble;  
licensed under Creative Commons License CC-BY

33rd European Conference on Object-Oriented Programming (ECOOP 2019).

Editor: Alastair F. Donaldson; Article No. 15; pp. 15:1–15:29

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

## 15:2 Transient Typechecks are (Almost) Free

45 before the execution, limiting the benefit of types to the statically typed parts of programs,  
46 and preventing programs from depending on type checks at run time. In the middle, “tran-  
47 sient” or “type-tag” checks as in Reticulated Python offer first-order semantics, checking  
48 whether an object’s type constructor or supported methods match explicit type declarations  
49 [49, 11, 46, 60, 29]. Reticulated Python also supports an alternative “monotonic” semantics  
50 which mutates an object to narrow its concrete type when it is passed into a more spe-  
51 cific type context. At the other end of the spectrum, behavioral typechecks as in Typed  
52 Racket [59, 57], Gradualtalk [3], and Reticulated Python’s proxies, support higher-order  
53 semantics, retaining types until run time, performing the checks eagerly, and giving detailed  
54 information about type violations as soon as possible via blame tracking [63, 2]. Finally,  
55 Ductile typing dynamically interprets a static type system at runtime [7]. Unfortunately,  
56 any gradual system with run-time semantics (i.e. everything more complex than erasure)  
57 currently imposes a significant run-time performance overhead to provide those semantics  
58 [56, 62, 42, 6, 45, 55, 29, 30].

59 The performance cost of run-time checks is problematic in itself, but also creates perverse  
60 incentives. Rather than the ideal of gradually adding types in the process of hardening a  
61 developing program, the programmer is incentivized to leave the program untyped or even  
62 to *remove* existing types in search of speed. While the Gradual Guarantee [50] requires that  
63 removing a type annotation does not affect the result of the program, the performance profile  
64 can be drastically shifted by the overhead of ill-placed checks. For programs with crucial  
65 performance constraints, for new programmers, and for gradual language designers, juggling  
66 this overhead can lead to increased complexity, suboptimal software-engineering choices, and  
67 code that is harder to maintain, debug, and analyze.

68 In this paper, we focus on the centre of the gradual typing spectrum: the transient,  
69 first-order, type-tag checks as used in Reticulated Python and similar systems. Several  
70 studies have found that these type checks have a negative impact on programs’ performance.  
71 Chung, Li, Nardelli and Vitek, for example, found that “*The transient approach checks types*  
72 *at uses, so the act of adding types to a program introduces more casts and may slow the*  
73 *program down (even in fully typed code).*” and say “*transient semantics... is a worst case*  
74 *scenario... , there is a cast at almost every call*” [22]. Greenman and Felleisen find that  
75 the slowdown is predictable, as transient checking “*imposes a run-time checking overhead*  
76 *that is directly proportional to the number of [type annotations] in the program*” [28], and  
77 Greenman and Migeed found a “*clear trend that adding type annotations adds performance*  
78 *overhead. The increase is typically linear.*” [29].

79 In contrast, we demonstrate that transient type checks can be “almost free” via a just-  
80 in-time compiler to an optimizing virtual machine. We insert gradual checks naïvely, for  
81 each gradual type annotation. Whenever an annotated method is called or returns, or an  
82 annotated variable is accessed, we check types dynamically, and terminate the program with  
83 a type error if the check fails. Despite this simplistic approach, a just-in-time compiler can  
84 eliminate redundant checks—removing almost all of the checking overhead, resulting in a  
85 performance profile aligned with untyped code.

86 We evaluate our approach by adding transient type checks to Moth, an implementation  
87 of the Grace programming language built on top of Truffle and the Graal just-in-time  
88 compiler [67, 66]. Inspired by Richards *et al.* [45] and Bauman *et al.* [6], our approach  
89 conflates types with information about the dynamic object structure (maps [20] or object  
90 shapes [65]), which allows the just-in-time compiler to reduce redundancy between checking  
91 structure and checking types; consequently, most of the overhead that results from type  
92 checking is eliminated.

93 The contributions of this paper are:

- 94 ■ demonstrating that VM optimizations enable transient gradual type checks with low  
95 performance cost
- 96 ■ an implementation approach that requires only small changes to existing abstract-syntax-  
97 tree interpreters
- 98 ■ an evaluation based on classic benchmarks and benchmarks from the literature on gradual  
99 typing

## 100 **2** Gradual Types in Grace

101 This section introduces Grace, and motivates supporting transient gradual typing in the  
102 language.

### 103 **2.1** The Grace Programming Language

104 Grace is an object-oriented, imperative, educational programming language, with a focus  
105 on introductory programming courses, but also intended for more advanced study and  
106 research [9, 19]. While Grace’s syntax draws from the so-called “curly bracket” tradition of  
107 C, Java, and JavaScript, the structure of the language is in many ways closer to Smalltalk:  
108 all computation is via dynamically dispatched “method requests” where the object receiving  
109 the request decides which code to run, and returns within lambdas that are “non-local”,  
110 returning to the method activation in which the block is instantiated [27]. In other ways,  
111 Grace is closer to JavaScript than Smalltalk: Grace objects can be created from object  
112 literals, rather than by instantiating classes [10, 35] and objects and classes can be deeply  
113 nested within each other [37].

114 Critically, Grace’s declarations and methods’ arguments and results can be annotated  
115 with types, and those types can be checked either statically or dynamically. This means the  
116 type system is intrinsically gradual: type annotations should not affect the semantics of a  
117 correct program [50], and the type system includes a distinguished “Unknown” type which  
118 matches any other type and is the implicit type for untyped program parts.

119 The static core of Grace’s type system is well described elsewhere [34]; here we explain  
120 how these types can be understood dynamically, from the Grace programmer’s point of view.  
121 Grace’s types are structural [9], that is, an object implements a type whenever it implements  
122 all the methods required by that type, rather than requiring classes or objects to declare  
123 types explicitly. Methods match when they have the same name and arity: argument and  
124 return types are ignored. A type thus expresses the requests an object can respond to, for  
125 example whether a particular accessor is available, rather than a nominal location in an  
126 explicit inheritance hierarchy.

127 Grace then checks the types of values at run time:

- 128 ■ the values of arguments are checked after a method is requested, but before the body of  
129 the method is executed;
- 130 ■ the value returned by a method is checked after its body is executed; and
- 131 ■ the values of variables are checked whenever written or read by user code.<sup>1</sup>

132 In the spectrum of gradual typing, these semantics are closest to the transient typechecks of  
133 Reticulated Python [60, 29]. Reticulated Python inserts transient checks only when a value

---

<sup>1</sup> Checking on read in addition to writes may seem unnecessary. For the rational, see Section 6.2.

## 15:4 Transient Typechecks are (Almost) Free

134 flows from untyped to typed code, while Grace inserts transient checks only at explicit type  
135 annotations (but in principle checks every annotation every time).

### 136 2.2 Why Gradual Typing?

137 Our primary motivation for this work is to provide a flexible system to check consistency  
138 between an execution of a program and its type annotations. A key part of the design  
139 philosophy of Grace is that the language should not force students to annotate programs  
140 with types until they are ready, so that teachers can choose whether to introduce types, early,  
141 late, or even not at all.

142 A secondary goal is to have a design that can be implemented with only a small set of  
143 changes to facilitate integration in existing systems.

144 Both of these goals are shared with much of the other work on gradual type systems, but  
145 our context leads to some different choices. First, while checking Grace’s type annotations  
146 statically may be optional, checking them dynamically should not be: any value that flows  
147 into a variable, argument, or result annotated with a type must conform to that type  
148 annotation. Second, adding type annotations should not degrade a program’s performance,  
149 or rather, programmers should not be encouraged to improve performance by removing  
150 type annotations. And third, we allow the programmer to execute a program even when  
151 not statically type-correct. Allowing such execution is useful to students, where they can  
152 see concrete examples of dynamic type errors. This is possible because Grace’s static type  
153 checking is optional, which means that an implementation cannot depend on the correctness  
154 or mutual compatibility of a program’s type annotations.

155 Unfortunately, existing gradual type implementations do not meet these goals, particularly  
156 regarding performance; hence the ongoing debate about whether gradual typing is alive,  
157 dead, or some state in between [56, 62, 42, 6, 45, 29, 30].

### 158 2.3 Using Grace’s Gradual Types

159 We now illustrate how the gradual type checks work in practice in the context of a simple  
160 program to record information about vehicles. Suppose the programmer starts developing  
161 this vehicle application by defining an object intended to represent a car (Listing 1, Line 1)  
162 and writes a method that, given the car object, prints out its registration number (Line 5).

```
1 def car = object {  
2   var registration is public := "J03553"  
3 }  
4  
5 method printRegistration(v) {  
6   print "Registration: {v.registration}"  
7 }
```

■ **Listing 1** The start of a simple Grace program for tracking vehicle information.

163 Next, the programmer adds a check to ensure any object passed to the `printRegistra-`  
164 `tion` method will respond to the `registration` request; they define the structural type  
165 `Vehicle` [58] naming just that method (Listing 2, Line 1), and annotate the `printRegis-`  
166 `tration` method’s argument with that type (Listing 2, Line 5). The annotation ensures  
167 that a type error will be thrown if an object, passed to the `printRegistration` method,

```

1 type Vehicle = interface {
2   registration
3 }
4
5 method printRegistration(v: Vehicle) {
6   print "Registration: {v.registration}"
7 }

```

■ **Listing 2** Adding a type annotation to a method parameter.

168 cannot respond to the `registration` message. Without the type check, the `print` method  
 169 would cause a run-time error when interpolating the string. However, since type errors cause  
 170 termination, the run-time error in the middle of the `print` implementation will now be  
 171 avoided.

172 In Listing 3, the programmer continues development and creates two car objects (Lines 9  
 173 and 17), that conform to an expanded `Vehicle` type (Line 1).

```

1 type Vehicle = interface {
2   registration
3   registerTo(_)
4 }
5
6 type Person = interface { name }
7 type Department = interface { code }
8
9 var personalCar : Vehicle :=
10  object {
11    var registration is public := "DLS018"
12    method registerTo(p: Person) {
13      print "{p.name} registers {self}"
14    }
15  }
16
17 var governmentCar : Vehicle :=
18  object {
19    var registration is public := "FKD218"
20    method registerTo(d: Department) {
21      print "some department {self}"
22    }
23  }
24
25 governmentCar.registerTo(
26  object {
27    var name is public := "Richard"
28  }
29 )

```

■ **Listing 3** A program in development with inconsistently typed `registerTo` methods.

174 Note that each version of the `registerTo` method declares a different type for its parameter  
 175 (Lines 12 and 20). When the programmer executes this program, both `personalCar` and  
 176 `governmentCar` can be assigned to a variable declared as `Vehicle` because checking that  
 177 assignment considers only that the vehicle has a `registerTo` method, but not the required  
 178 argument type of that method. At Line 25 the developer attempts to register a government

## 15:6 Transient Typechecks are (Almost) Free

179 car to a person: only when the method (Line 20) is *invoked* will the gradual type test on  
180 the argument fail (the object that is passed in is not a `Department` because it lacks a `code`  
181 `method`).

### 182 **3 Graal, Truffle, Self-Optimization and Dynamic Adaptive** 183 **Compilation**

184 This section gives a brief introduction into just-in-time compilation, and the main techniques  
185 we rely on for our optimizations.

#### 186 **3.1 Self-Optimizing Interpreters**

187 Self-optimizing abstract-syntax-tree (AST) interpreters [68] are the foundation for the work  
188 presented here. Together with partial evaluation [66], self-optimization enables efficient  
189 dynamic language implementations that reach the performance of custom state-of-the-art  
190 virtual machines (cf. Section 5.2 and [41]). The framework for building such interpreters is  
191 called Truffle.

192 The key idea is that an AST rewrites itself based on a program’s run-time values to  
193 reflect the minimal set of operations needed to execute the program correctly.

194 As an example, consider the addition of two numbers in a dynamic language, possibly  
195 written simply as: `a + b`. Because there are no static types known, the run-time values  
196 for `a` and `b` could potentially be anything from an integer or a double, to a string or a  
197 collection, or any arbitrary objects that have a “+” method. In an self-optimizing interpreter,  
198 the expression may be represented by an `add` node, with two child nodes that each read a  
199 variable. The first time the `add` node executes, it may find that both values to be added  
200 are integers. It will then speculate that all future executions also see integers, and thus,  
201 rewrite itself to an `add-integer` node. This `add-integer` node will simply confirm that  
202 both values are integers, and then directly perform the integer addition. Compared to a  
203 general `add` node, we do not have to cover the cases for doubles, strings, and other kinds of  
204 objects, which results in much simpler code that can be more easily optimized. All other  
205 cases are supported by rewriting the `add` node to more general versions. This happens for  
206 instance, when the values are not integers, however, programs are often very monomorphic  
207 in practice, and so the speculation is highly beneficial.

208 As a consequence of the rewriting, what often starts out as something close to a traditional  
209 AST, in the end incorporates additional knowledge about the execution. Thus, such trees  
210 should be referred to more correctly as *execution trees* rather than ASTs.

#### 211 **3.2 Polymorphic Inline Caches for Optimizing Dynamic Behavior**

212 Polymorphic inline caches (PICs) [32] are a variation on the theme of caching run-time values  
213 to improve performance. Originally, they focused on method invocation in dynamic languages  
214 to avoid costly method lookups by caching the looked-up method for a specific type. For  
215 dynamic languages, PICs can be generalized to not only consider the receiver type, but  
216 instead for instance the object shape (cf. Section 3.3), which enables the optimizations we  
217 are aiming for.

218 In a language such as JavaScript, a PIC would be used for instance for the following  
219 expression: `obj.toString()`. The dot can be thought of as the lexical representation of the  
220 method lookup. An implementation would keep a small cache for each such dot in the code.  
221 This means, for each lexical lookup location, we have a separate cache. PICs benefit from

222 the relatively monomorphic behavior of programs. A specific lexical lookup is likely to see  
223 only one kind of object in the `obj` variable, so the cache will usually have the correct method  
224 for the object ready and can avoid a costly lookup.

### 225 3.3 Object Shapes: Metadata for Dynamic Objects

226 Object shapes [65], which are also known as maps [20] or hidden classes, are in the most general  
227 case a type and usage profile for groups of objects. In languages such as Self, JavaScript, and  
228 Grace, we do not have traditional classes that define the set of fields for an object. The set of  
229 fields might even change over time. Furthermore, fields can theoretically store any possible  
230 value. However, in practice, the behavior of programs is again relatively monomorphic and  
231 objects created in a specific part of a program are likely to have always the same set of fields,  
232 which each are used to store only a small number of different kinds of values. For example,  
233 an object representing a counter would have a field `count`, which always stores integers, while  
234 an object representing a person may have always a field `name` that stores a string, but never  
235 an integer.

236 Object shapes represent this run-time information in a way that allows a just-in-time  
237 compiler to represent objects in memory similarly to C structs, and then to generate highly  
238 efficient code. Object shapes can be conflated with additional information, for instance to  
239 represent knowledge about types [6, 45]. For the use of PICs, object shapes are important,  
240 because they give objects a form of identification that groups them, and which in practice,  
241 has similar properties with respect to monomorphic behavior as classes have.

### 242 3.4 Just-in-Time Compilation with Graal and Truffle

243 The Graal compiler is a just-in-time compiler for Java. For languages built on the Truffle  
244 framework, Graal comes with additional support for partial evaluation, which enables efficient  
245 native code generation for Truffle interpreters [66].

246 As such, Graal is a metacompiler. This means that instead of compiling a specific  
247 program, in our case a Grace program, Graal compiles our Grace interpreter Moth for the  
248 execution of a specific Grace method. For simplicity, partial evaluation can be thought  
249 of a highly aggressive inlining strategy. It starts with the root node of a specific Grace  
250 method and inlines all code reachable from it, while considering the execution tree to be  
251 constant. To enable further optimizations, Graal does further inlining on the level of the  
252 Grace program, which is important to expose the same optimization opportunities classic  
253 just-in-time compilers have. The applied optimizations include for instance constant folding,  
254 common subexpression elimination, and loop-invariant code motion.

255 Especially loop-invariant code motion and common subexpression elimination are im-  
256 portant to generate efficient native code for dynamic languages. Since we rely on techniques  
257 such as self-optimizing nodes, PICs and object shapes, which all introduce various checks, a  
258 compiler needs to move these out of loops, and remove redundant checks.

259 By combining all the techniques sketched in this section, Graal and Truffle are able to  
260 execute dynamic languages as efficiently as virtual machines built for a specific language –  
261 but with much less implementation effort.

## 262 4 Moth: Grace on Graal and Truffle

263 Implementing dynamic languages as state-of-the-art virtual machines can require enorm-  
264 ous engineering efforts. Meta-compilation approaches [41] such as RPython [12, 14] and

## 15:8 Transient Typechecks are (Almost) Free

265 GraalVM [67, 66] reduce the necessary work dramatically, because they allow language  
266 implementers to leverage existing VMs and their support for just-in-time compilation and  
267 garbage collection.

268 Moth [47] adapts SOMNS [39] to leverage this infrastructure for Grace. SOMNS is a  
269 Newspeak implementation [18] on top of the Truffle framework and the Graal just-in-time  
270 compiler, which are part of the GraalVM project. One key optimization of SOMNS for this  
271 work is the use of object shapes [65], also known as maps [20] or hidden classes. They represent  
272 the structure of an object and the types of its fields. In SOMNS, shapes correspond to the class  
273 of an object and augment it with run-time type information. With Moth's implementation,  
274 SOMNS was changed to parse Grace code, adapting a few of the self-optimizing abstract-  
275 syntax-tree nodes to conform to Grace's semantics. Despite these changes Moth preserves the  
276 peak performance of SOMNS, which reaches that of V8, Google's JavaScript implementation  
277 (cf. Section 5.2 and Marr *et al.* [40]).

### 278 4.1 Adding Gradual Type Checking

279 One of the goals for our approach to gradual typing was to keep the necessary changes to  
280 an existing implementation small, while enabling optimization in highly efficient language  
281 runtimes. In an AST interpreter, we can implement this approach by attaching the checks  
282 to the relevant AST nodes: the expected types for the argument and return values can be  
283 included with the node for requesting a method, and the expected type for a variable can  
284 be attached to the nodes for reading from and writing to that variable. In practice, we  
285 encapsulate the logic of the check within a new class of AST nodes, specially to support  
286 gradual type checking. Moth's front end was adapted to parse and record type annotations  
287 and attach instances of this checking node as children of the existing method, variable read,  
288 and variable write nodes.

289 The check node uses the internal representation of a Grace type (cf. Listing 4, Line 13)  
290 to test whether an observed object conforms to that type. An object satisfies a type if all  
291 members required by the type are provided by that object (Line 5).

292 Note, we use a pseudo code syntax similar to Python for all code examples that represent  
293 the implementation of Moth. We chose this syntax to avoid any confusion with our Grace  
294 examples (even though Moth is implemented in Java).

```
1 class Type:
2     def init(members):
3         self._members = members
4
5     def is_satisfied_by(other: Type):
6         for m in self._members:
7             if m not in other._members:
8                 return False
9         return True
10
11    def check(obj: Object):
12        t = obj.get_type()
13        return self.is_satisfied_by(t)
```

■ **Listing 4** Sketch of a `Type` in our system and its `check()` semantics.



```

1 global record: Matrix
2
3 class TypeCheckNode(Node):
4
5     expected: Type
6
7     @Spec(static_guard=`expected.check(obj)`)
8     def check(obj: Number):
9         pass
10
11    @Spec(static_guard=`expected.check(obj)`)
12    def check(obj: String):
13        pass
14
15    ...
16
17    @Spec(guard=`obj.shape==cached_shape`, static_guard=`expected.check(obj)`)
18    def check(obj: Object, @Cached(obj.shape) cached_shape: Shape):
19        pass
20
21    @Fallback
22    def check(obj: Any):
23        T = obj.get_type()
24
25        if record[T, expected] is unknown:
26            record[T, expected] = T.is_subtype_of(expected)
27
28        if not record[T, expected]:
29            raise TypeError("{obj} doesn't implement {expected}")

```

■ **Listing 5** A sketch of the specializations in `TypeCheckNode` to minimize the run-time overhead of type checking. A specialization is a minimal set of operations for one specific situation, e.g., that the value to be checked is some type of number.

## 295 4.2 Optimization

296 There are two aspects to our implementation that are critical for a minimal-overhead solution:

- 297 ■ specialized executions of the type checking node, along with guards to protect these
- 298 specialized versions, and
- 299 ■ a matrix to cache sub-typing relationships to eliminate redundant exhaustive subtype
- 300 tests.

301 **Optimized Type Check Node** The first performance-critical aspect to our implementation  
 302 is the optimization of the type checking node. We rely on Truffle and its TruffleDSL [31]. This  
 303 means we provide a number of special cases, which are selected during execution based on the  
 304 observed concrete kinds of objects. A sketch of our type checking node using a pseudo-code  
 305 version of the DSL is given in Listing 5. A simple optimization is for well known types such  
 306 as numbers (Line 8) or strings (Line 12). The methods annotated with `@Spec` (shorthand  
 307 for `@Specialization`) correspond to possible states in a state machine that is generated by  
 308 the TruffleDSL. Thus, if a check node observes a number or a string, it will check on the  
 309 first execution only that the expected type, i.e., the one defined by some type annotation, is  
 310 satisfied by the object by using a `static_guard`. If this is the case, the DSL will activate  
 311 this state. For just-in-time compilation, only the activated states and their normal guards

## 15:10 Transient Typechecks are (Almost) Free

```
1 class VariableReadNode(Node):
2     slot: FrameSlot
3     type_check: TypeCheckNode
4
5     @Spec
6     def do_read(frame: VirtualFrame):
7         value = frame.read(slot)
8         if type_check:
9             type_check.check(value)
10        return value
```

■ **Listing 6** Sketch of a `VariableReadNode` using the `TypeCheckNode` to ensure Grace’s transient semantics.

312 are considered. A `static_guard` is not included in the optimized code. If a check fails, or  
313 no specialization matches, the fallback, i.e., `check_generic` is selected (Line 22), which may  
314 raise a type error.

315 For generic objects, we rely on the specialization on Line 18, which checks that the object  
316 satisfies the expected type. If that is the case, it reads the shape of the object (cf. Section 4)  
317 at specialization time, and caches it for later comparisons. Thus, during normal execution,  
318 we only need to read the shape of the object and then compare it to the cached shape with  
319 a simple reference comparison. If the shapes are the same, we can assume the type check  
320 passed successfully. Note that shapes are not equivalent to types, however, shapes imply  
321 the set of members of an object, and thus, do imply whether an object fulfills one of our  
322 structural types.

323 The `TypeCheckNode` is used in Moth in all places that need to check types, which includes  
324 reading and writing variables as well as method requests and returns. Listing 6 shows a  
325 sketch of an AST node that implements reading from a local variable, which is stored in a  
326 frame object. A frame corresponds to a stack frame, sometimes also called an environment.

327 Line 8 first checks whether a type check needs to be performed. Since type annotations  
328 are optional, it may not be necessary to check for a type. Note that `type_check` is a constant  
329 for just-in-time compilation (cf. Section 3.4), which enables subsequent optimizations. Line 9  
330 then calls the `check()` method on the `TypeCheckNode`, which may result in a type error. For  
331 a variable that only contains numbers, the `type_check` object would activate the number  
332 specialization in its state machine. For just-in-time compilation, this means only the code  
333 for checking numbers needs to be compiled, but none of the other specializations.

334 In many cases, the specialization for objects would be activated in a `TypeCheckNode`,  
335 which checks the shape of an object against a cache. This check is identical to the check  
336 performed by a polymorphic inline cache (PIC, cf. Section 3.2). Since PICs are used for all  
337 method calls, they are very common in most programs, and these additional checks can often  
338 be removed easily via common subexpression elimination.

339 **Subtype Cache Matrix** The other performance-critical aspect to our implementation is  
340 the use of a matrix to cache sub-typing relationships. The matrix compares types against  
341 types, featuring all known types along the columns and the same types again along the rows.  
342 A cell in the table corresponds to a sub-typing relationship: does the type corresponding  
343 to the row implement the type corresponding to the column? All cells in the matrix begin  
344 as unknown and, as encountered in checks during execution, we populate the table. If a  
345 particular relationship has been computed before we can skip the check and instead recall the

346 previously-computed value (Line 26 in Listing 5). Using this table we are able to eliminate  
347 the redundancy of evaluating the same type to type relationships across different checks in  
348 the program. To reduce redundancy further we also unify types in a similar way to Java’s  
349 string interning; during the construction of a type we first check to see if the same set of  
350 members is expressed by a previously-created type and, if so, we avoid creating the new  
351 instance and provide the existing one instead.

352 Together the self-specializing type check node and the cache matrix ensure that our  
353 implementation eliminates redundancy, and consequently, we are able to minimize the  
354 run-time overhead of our system.

## 355 **5 Evaluation**

356 To evaluate our approach to gradual type checking, we first establish the baseline performance  
357 of Moth compared to Java and JavaScript, and then assess the impact of the type checks  
358 themselves.

### 359 **5.1 Method and Setup**

360 To account for the complex warmup behavior of modern systems [4] as well as the non-  
361 determinism caused by e.g. garbage collection and cache effects, we run each benchmark for  
362 1000 iterations in the same VM invocation.<sup>2</sup> Afterwards, we inspected the run-time plots  
363 over the iterations and manually determined a cutoff of 350 iterations for warmup, i.e., we  
364 discard iterations with signs of compilation. As a result, we use a large number of data  
365 points to compute the average, but outliers, caused by e.g. garbage collection, remain visible  
366 in the plots. All reported averages use the geometric mean since they aggregate ratios.

367 All experiments were executed on a machine running Ubuntu Linux 16.04.4, with Kernel  
368 3.13. The machine has two Intel Xeon E5-2620 v3 2.40GHz, with 6 cores each, for a total  
369 of 24 hyperthreads. We used ReBench 0.10.1 [38], Java 1.8.0\_171, Graal 0.33 (a13b888),  
370 Node.js 10.4, and Higgs from 9 May 2018 (aa95240). Benchmarks were executed one by  
371 one to avoid interference between them. The analysis of the results was done with R 3.4.1,  
372 and plots are generated with ggplot 2.2.1 and tikzDevice 0.11. Our experimental setup is  
373 available online to enable reproductions.<sup>3</sup>

### 374 **5.2 Are We Fast Yet?**

375 To establish the performance of Moth, we compare it to Java and JavaScript. Moth is used in  
376 its untyped version, i.e., without type checks. For JavaScript we chose two implementations,  
377 Node.js with V8 as well as the Higgs VM. The Higgs VM is an interesting point of comparison,  
378 because Richards *et al.* [45] used it in their study. The goal of this comparison is to determine  
379 whether our approach could be applicable to industrial strength language implementations  
380 without adverse effects on their performance.


381 We compare across languages based on the Are We Fast Yet benchmarks [40], which are  
382 designed to enable a comparison of the effectiveness of compilers across different languages.  
383 To this end, they use only a common set of core language elements. While this reduces the  
384 performance-relevant differences between languages, the set of core language elements covers

---

<sup>2</sup> For the Higgs VM, we only use 100 iterations, because of its lower performance. This is sufficient since Higgs’s compilation approach induces less variation and leads to more stable measurements.

<sup>3</sup> **SM TODO** *merge changes, and tag final version* <https://github.com/gracelang/moth-benchmarks>

## 15:12 Transient Typechecks are (Almost) Free



■ **Figure 1** Comparison of Java 1.8, Node.js 10.4, Higgs VM, and Moth. The boxplot depicts the peak-performance results for the Are We Fast Yet benchmarks, each benchmark normalized individually based on the result for Java, which means all results for Java are 1.0, and its box appears as a line. The dots on the plot represent the geometric mean reported as averages. For these benchmarks, Moth is within the performance range of JavaScript, as implemented by Node.js, which makes Moth an acceptable platform for our experiments.

385 only common object-oriented language features with first-class functions. Consequently, these  
386 benchmarks are not necessarily a predictor for application performance, but can give a good  
387 indication for basic mechanisms such as type checking.

388 Figure 1 shows the results. We use Java as baseline since it is the fastest language  
389 implementation in this experiment. Note that we perform a unit conversion on the results  
390 separately for each benchmark, using the average of Java as 1 unit. While this conversion  
391 does not change the distribution of the data, it allows us to show it neatly on one plot.

392 We see that Node.js (V8) is about 1.8x (min. 0.8x, max. 2.7x) slower than Java. Moth is  
393 about 2.3x (min. 0.9x, max. 4.3x) slower than Java. As such, it is on average 31% (min.  
394  $-16\%$ , max. 2.3x) slower than Node.js. Compared to the Higgs VM, which is on these  
395 benchmarks 10.4x (min. 1.5x, max. 163x) slower than Java, Moth reaches the performance of  
396 Node.js more closely. With these results, we argue that Moth is a suitable platform to assess  
397 the impact of our approach to gradual type checking, because its performance is close enough  
398 to state-of-the-art VMs, and run-time overhead is not hidden by slow baseline performance.

### 399 5.3 Performance of Transient Gradual Type Checks

400 The performance overhead of our transient gradual type checking system is assessed based  
401 on the Are We Fast Yet benchmarks as well as benchmarks from the gradual-typing literature.  
402 The goal was to complement our benchmarks with additional ones that are used for similar  
403 experiments and can be ported to Grace. To this end, we surveyed a number of papers [56,  
404 62, 42, 6, 45, 55, 29] and selected benchmarks that have been used by multiple papers. Some  
405 of these benchmarks overlapped with the Are We Fast Yet suite, or were available in different  
406 versions. While not always behaviorally equivalent, we chose the Are We Fast Yet versions  
407 since we already used them to establish the performance baseline. The selected benchmarks  
408 as well as the papers in which they were used are shown in Table 1.

409 The benchmarks were modified to have complete type information. To ensure correctness  
410 and completeness of these experiments, we added an additional check to Moth that reports  
411 absent type information to ensure each benchmark is completely typed. To assess the  
412 performance overhead of type checking, we compare the execution of Moth with all checks

■ **Table 1** Benchmarks selected from literature.

Fannkuch	[62, 29]	
Float	[62, 42, 29]	
Go	[62, 42, 29]	
NBody	[36, 62, 29]	used [40]
Queens	[62, 42, 29]	used [40]
PyStone	[62, 42, 29]	
Sieve	[56, 42, 6, 45, 30]	used [40]
Snake	[56, 42, 6, 45, 30]	
SpectralNorm	[62, 42, 29]	

413 disabled, i.e., the baseline version from Section 5.2, against an execution that has all checks  
 414 enabled. We did not measure programs that mix typed and untyped code because with our  
 415 implementation technique a fully typed program is expected to have the largest overhead.

#### 416 **Peak Performance**

417 Figure 2 depicts the overall results comparing Moth, with all optimizations, against the  
 418 untyped version. The run-time overhead, after discarding the warmup iterations, is on  
 419 average 5% (min. -13%, max. 79%).

420 The benchmark with the highest overhead of 79% is List. The benchmark traverses a  
 421 linked list and has to check the list elements individually. Unfortunately, the structure of  
 422 this list introduces checks that do not coincide with shape checks on the relevant objects.  
 423 We consider this benchmark a pathological case and discuss it in detail in Section 6.1.

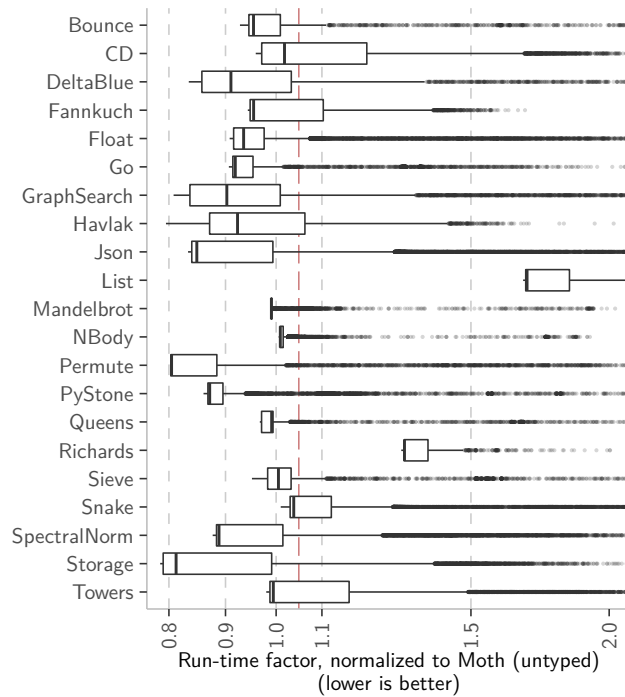
424 Beside List, the highest overheads are on Richards (33%), CD (12%), Snake (14%), and  
 425 Towers (12%). Richards has one major component, also a linked list traversal, similar to  
 426 List. Snake and Towers primarily access arrays in a way that introduces checks that do not  
 427 coincide with behavior in the unchecked version.

428 In some benchmarks, however, the run time decreased; notably Permute (-13%), Graph-  
 429 Search (-3%), and Storage (-8%). Permute simply creates the permutations of an array.  
 430 GraphSearch implements a page rank algorithm and thus is primarily graph traversal. Storage  
 431 stresses the garbage collector by constructing a tree of arrays. For these benchmarks the  
 432 introduced checks seem to coincide with shape-check operations already performed in the  
 433 untyped version. The performance improvement is possibly caused by having checks earlier,  
 434 which enables the compiler to more aggressively move them out of loops. Another reason  
 435 could simply be that the extra checks shift the boundaries of compilation units. In such cases,  
 436 checks might not be eliminated completely, but the shifted boundary between compilation  
 437 units might mean that the generated native code interacts better with the instruction cache  
 438 of the processor.

#### 439 **Warmup Performance**

440 To more precisely measure warmup, all relevant experiments were executed 30 times, with  
 441 each running for 100 iterations. The resulting Figure 3 shows the first 100 iterations for each  
 442 benchmark. For each iteration  $n$ , we normalized the measurements to the mean of iteration  
 443  $n$  of the untyped Moth implementation. Thus, any increase indicates a slow down because of

## 15:14 Transient Typechecks are (Almost) Free



■ **Figure 2** A boxplot comparing the performance of Moth with and without type checking. The plot depicts the run-time overhead on peak performance over the untyped performance. On average, transient type checking introduces an overhead of 5% (min. -13%, max. 79%). The average is indicated as a line with long dashes. The visible outliers correspond to various complex aspects of the overall system, e.g., including garbage collection and cache effects. Note that the axis is logarithmic to avoid distorting the proportions of relative speedups and slowdowns.

444 typing. The darker lines indicate the means, while the lighter area indicates a 95% confidence  
445 interval.

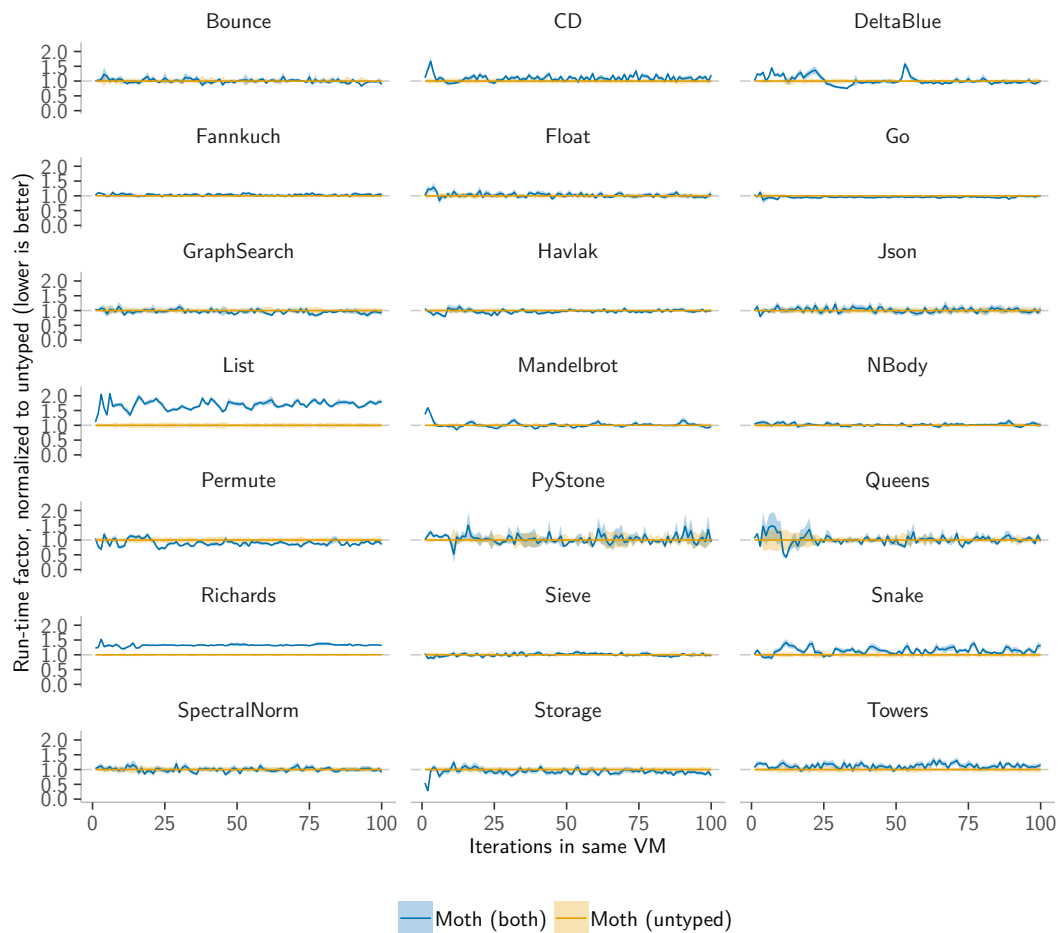
446 Looking only at the first few iterations, where we assume that most code is executed in  
447 the interpreter and might be affected by compilation activity, the overhead appears minimal.  
448 Only the Mandelbrot and CD benchmarks shows a noticeable slowdown.

449 Mandelbrot with its distinctly slow first iteration can be explained by its code structure.  
450 Since it is a computational kernel with many primitive operations, but no method calls,  
451 optimized code is only reached after the first full benchmark iteration. The problem could  
452 be alleviated with on-stack-replacement for loops, which is currently not done. Since other  
453 benchmarks use methods, they reach compiled code earlier and do not exhibit the same  
454 first-iteration slowdown.

455 PyStone however show various spikes. Since spikes appear in both directions (speedups  
456 and slowdowns), we assume that they indicate a shift, for instance, of garbage collection  
457 pauses, which may happen because of different heap configurations triggered by the additional  
458 data structures for type information.

## 459 5.4 Effectiveness of Optimizations

460 To characterize the concrete impact of our two optimizations, i.e., the optimized type checking  
461 node, which replaces complex type tests with checks for object shapes, and our matrix to  
462 cache sub-typing information, we look at the number of type checks performed by the



■ **Figure 3** Plot of the run time for the first 100 iterations. The lines indicate the mean at iteration  $n$  normalized to the untyped result, the lighter area indicates a 95% confidence interval. The first iteration, i.e., mostly interpreted, seems to be affected significantly only for Mandelbrot, though CD shows slower behavior in early warmup, too.

463 benchmarks, as well as the impact on peak performance.

#### 464 Impact on Performed Type Tests

465 Table 2 gives an overview of the number of type tests done by the benchmarks during execution.  
 466 We distinguish two operations `check_generic` and `is_subtype_of`, which correspond to  
 467 the operations in Line 22 and Line 5 of Listing 4. Thus, `check_generic` is the test called  
 468 whenever a full type check has to be performed, and `is_subtype_of` is the part of the check  
 469 that determines the relationship between two types. The second column of Table 2 indicates  
 470 which optimization is applied, and the following columns show the mean, minimum, and  
 471 maximum number of invocations of the tests over all benchmarks.

472 The baselines without optimizations are the rows with the results for neither of the  
 473 optimizations being enabled. Depending on the benchmark, we see that the type tests are  
 474 done tens of millions to hundreds of millions times for a single iteration of a benchmark.

## 15:16 Transient Typechecks are (Almost) Free

■ **Table 2** Type Test Statistics over all Benchmarks. This table shows how many of the type tests can be avoided based on our two optimizations. As indicated by the numbers, the number of type tests can vary significantly between benchmarks. Thus, the table shows the mean, minimum, and maximum number of type tests across all benchmarks for a given configuration. With the use of an optimized node that replaces type checks with simple object shape checks, `check_generic` is invoked only for the first time that a lexical location sees a specific object shape, which eliminates run-time type checks almost completely. Using our subtype matrix that caches type-check results, invocations of `is_subtype_of` are further reduced by an order of magnitude.

Type Test	Enabled Optimization	mean #invocations	min	max
check_generic	Neither	137,525,845	11,628,068	896,604,537
	Subtype Cache	137,525,845	11,628,068	896,604,537
	Optimized Node	292	68	1,012
	Both	292	68	1,012
is_subtype_of	Neither	134,125,215	11,628,067	896,604,534
	Subtype Cache	16	10	29
	Optimized Node	292	68	1,012
	Both	16	10	29


475 Our optimizations reduce the number of type test invocations dramatically. As a result,  
476 the full check for the subtyping relationship is done only once for any specific type and a  
477 possible super type. Similarly, the generic type check is replaced by a shape check and thus  
478 minimizes the number of expensive type checks to the number of lexical locations that verify  
479 types combined with the number of shapes a specific lexical location sees at run time.

### 480 Impact on Performance

481 Figure 4 shows how our optimizations contribute to the peak performance. The figure depicts  
482 Moth's peak performance over all benchmarks, depending on the activated optimizations. As  
483 for Figure 1, we do a per-benchmark unit conversion using Moth (untyped), preserving the  
484 distribution properties of the results, but enabling us to show the results on a single plot.

485 As seen before in Figure 2, the untyped version is faster by 5%. Moth with both  
486 optimizations enabled as well as Moth with the optimized type-check node (cf. Listing 4)  
487 reach the same performance. This indicates that the subtype cache matrix is not strictly  
488 necessary for the peak performance. However, we can see that the subtype cache matrix  
489 improves performance by an order of magnitude over the Moth version without any type  
490 check optimizations. This shows that it is a relevant and useful optimization. Based on the  
491 numbers of Table 2, we see that this optimization is relevant for the very first execution  
492 of code. For code that has not executed before, having the global cache for the subtype  
493 relations gives the most benefit. After the first execution, the lexical caches in form of the  
494 type check nodes are primed with the same information, and the subtype cache matrix is  
495 only rarely needed. An example for code that benefits from the subtype cache matrix is unit  
496 test code, because most of the code is executed only once. While the performance of unit  
497 tests is not always critical, it can have a major impact on developer productivity.





■ **Figure 4** Performance Impact of the Optimizations on the Peak Performance over all Benchmarks. The boxplot shows the performance of Moth normalized to the untyped version, i.e., without any type checks. This means all results for Moth (untyped) are 1.0 and its box appears as a line. The dots on the plot represent the geometric mean reported as averages. The performance of Moth with both optimizations and Moth with only the node for optimized type checks are identical. The subtype check cache improves performance over the unoptimized version, but does not contribute to the peak performance.

#### 498 Impact on Memory Usage

499 In our implementation, the subtype cache matrix is the largest additional data structure. We  
 500 initialize it for up to 1000 types and use 1 byte per type combination. Java utilizes ca. 1MB  
 501 of memory for the matrix. Additional memory is used to represent the type-check nodes  
 502 at every lexical location. Since they behave like polymorphic inline caches (PIC) [32], their  
 503 memory usage depends on the specific program execution. For the benchmarks used in this  
 504 paper, the extra memory use can be up to 200KB.

505 In the context of Graal and Truffle, this additional memory usage is small, since the  
 506 metacompilation approach uses a lot of memory [41]. In a dedicated virtual machine, memory  
 507 use can be further optimized and be as efficient as for other kinds of PICs.

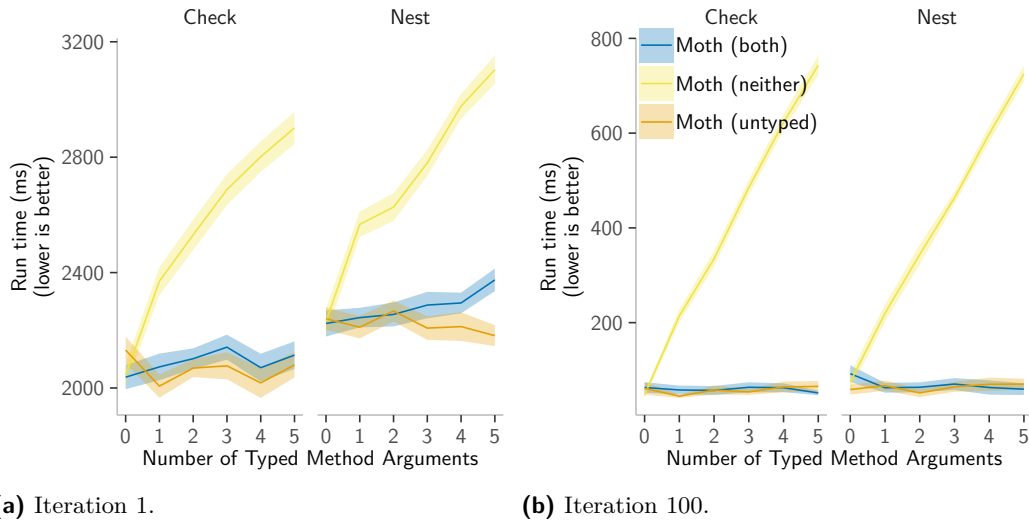
### 508 5.5 Transient Typechecks are (Almost) Free

509 As discussed in the introduction, in many existing gradually typed systems, one would expect  
 510 a linear increase of the performance overhead with the increasing number of type annotations.

511 In this section, we show that this is not necessarily the case on our system. For this  
 512 purpose, we use two microbenchmarks Check and Nest, which have at their core method  
 513 calls with 5 parameters. The Check benchmark calls the same method 10 times in a row, i.e.,  
 514 it has 10 call sites. The Nest benchmark has 10 methods with identical signatures, which  
 515 recurse from the first one to the last one. Thus, there are still 10 method calls, but they  
 516 are nested in each other. In both benchmarks, each method increments a counter, which  
 517 is checked at the end of the execution to verify that both do the same number of method  
 518 activations, and only the shape of the activation stack differs.

519 Each benchmark exists in six variants, each variant in a separate file, going from having  
 520 no type annotations over annotating only the first method parameter to annotating all 5  
 521 parameters. To demonstrate the impact of compilation, we present the results for the first

## 15:18 Transient Typechecks are (Almost) Free



■ **Figure 5** Transient Typechecks are (Almost) Free. Two microbenchmarks, each with six variants, demonstrate the common scenario of adding type annotations over time, which in our system does not have an impact on peak performance. The benchmark variants differ only in the increasing number of method arguments that have type annotations. We show the result for the first benchmark iteration (a) and the one hundredth (b). Moth (neither), i.e., Moth without our two optimizations sees a linear increase in run time. For the first iteration, we see some difference between Moth (both) and Moth (untyped). By the hundredth iteration, however, the compiler has eliminated the overhead of the type checks and both Moth variants essentially have the same performance (independent of the number of method arguments with type annotations).

522 iteration as well as the hundredth iteration. The first iteration is executed at least partially  
523 in the interpreter, while the hundredth iteration executes fully compiled.

524 Figure 5 shows that such a common scenario of methods being gradually annotated with  
525 types does not incur an overhead on peak performance in our system. The plot shows the  
526 mean of the run time for each benchmark configuration. Furthermore, it indicates a band  
527 with the 95% confidence interval. The yellow line, Moth (neither), corresponds to our Moth  
528 with type checking but without any optimizations. For this case, we see that the performance  
529 overhead grows linearly with the number of type annotations.

530 For Moth (both) and Moth (untyped), we see for the first iteration that the band of  
531 confidence intervals diverges, indicating that the additional type checks have an impact on  
532 startup performance. However, for the hundredth iteration, the confidence intervals overlap  
533 for the optimized Moth as well as the one that does not perform typechecks. This means that  
534 Moth does not suffer from a general linear overhead for adding type checks. Instead, most  
535 type checks do not have an impact on peak performance. However, as previously argued for  
536 the List benchmark, this is only the case for checks that can be subsumed by shape checks  
537 (shape checks are performed whether or not type checks are present).

## 538 5.6 Changes to Moth

539 Outlined earlier in Section 4, a secondary goal of our design was to enable the implementation  
540 of our approach to be realized with few changes to the underlying interpreter. This helps to  
541 ensure that each Grace implementation can provide type checking in a uniform way.

542 By examining the history of changes maintained by our version control, we estimate that

```

1 type ListElement = interface {
2   next
3 }
4
5 var elem: ListElement := headOfList
6 while (...) do {
7   elem := elem.next
8 }

```

■ **Listing 7** Example for dynamic type checks not corresponding to existing checks.

543 our implementation of Moth required 549 new lines and 59 changes to existing lines. The  
544 changes correspond to the implementation of new modules for the type class (179 lines) and  
545 the self-specializing type checking node (139 lines), modifications to the front end to extract  
546 typing information (115 new lines, 14 lines changes) and finally the new fields and amended  
547 constructors for AST nodes (116 new lines, 45 lines changes).

## 548 6 Discussion

### 549 6.1 The VM Could Not Already Know That

550 One of the key optimizations for our work and the work of others [6, 45] is the use of object  
551 shapes to encode information about types (in our case), or type casts and assumptions (in  
552 the case of gradually typed systems). The general idea is that a VM will already use object  
553 shapes for method dispatches, field accesses, and other operations on objects. Thus any  
554 further use to also imply type information can often be optimized away when the compiler  
555 sees that the same checks are done, and therefore can be combined by optimizations such as  
556 common subexpression elimination.

557 This assumption breaks, however, when checks are introduced that do not correspond  
558 to those that exist already. As described in Section 4, our approach introduces checks for  
559 reading from and writing to variables. Listing 7 gives an example of a pathological case. It  
560 is a loop traversing a linked list. For this example our approach introduces a check, for the  
561 `ListElement` type, when (1) assigning to and reading from `elem` and (2) when activating  
562 the `next` method. The checks for reading from `elem` and activating the method can be  
563 combined with the dispatch’s check on object shape. Unfortunately, the compiler cannot  
564 remove the check when writing to `elem`, because it has no information about what value will  
565 be returned from `next`, and so it needs to preserve the check to be able to trigger an error  
566 on the assignment. For our List benchmark, this check induces an overhead of 79%.

567 Compiler optimizations such as inlining are also insufficient for this particular case,  
568 because there are no guarantees about what `elem` does to implement `next`. The `next` method  
569 of a specific kind of `ListElement` may even have a type annotation for a return value. The  
570 best Graal can do in this example is to combine the check for the return value with the one  
571 writing to `elem`.

572 Since the example shows a somewhat generic data structure, there is the question of  
573 whether the issue applies to other data structures as well. Our benchmarks use a range of  
574 data structures including hash maps, sets, and vectors, none of which show the issue, because  
575 in more complex programs the chance of already having a check there is high, and cases  
576 where there has not been one before seem to be rare — although one can always consider  
577 additional optimizations to eliminate further checks. For generic data structures, storage

## 15:20 Transient Typechecks are (Almost) Free

578 strategies [13] could be used to encode type information about elements. This would allow  
579 the VM to check only once before a loop, and the loop could then rely on that check for  
580 guarantees about the elements of the data structure.

### 581 6.2 Optimizations

582 **Read and Write Checks.** As a simplification, we currently check variable access on both  
583 reads and writes. This approach simplifies the implementation, because we do not need to  
584 adapt all built-ins, i.e., all primitive operations provided by the interpreter. One optimization  
585 could be to avoid read checks. A type violation can normally only occur when writing to  
586 a variable, but not when reading. However, to maintain the semantics, this would require  
587 us to adapt many primitives. Examples for primitives are operations that activate blocks,  
588 which need to check their arguments or return values as well as any primitives that write to  
589 variables or fields. Given the number of primitives, this is error prone and incompleteness  
590 would result in missing type checks.

591 By checking reads and writes in a few well defined locations, we get errors as soon as user  
592 code accesses fields and variables. Moreover, only a small set of locations required changes  
593 to the code, which simplified the implementation. Given the good results (cf. Sections 5.4  
594 and 5.6), we decided to keep read checks, because it is a more uniform and maintainable  
595 approach for an academic project.

596 **Dynamic Type Propagation.** Another optimization could be to use Truffle’s approach to  
597 self-specialization [68] and propagate type information to avoid redundant checks. At the  
598 moment, Truffle interpreters typically use self-specialization to specialize the AST to avoid  
599 boxing of primitive types. This is done by speculating that some subtree always returns  
600 the expected type. If this is not the case, the return value of the subtree is going to be  
601 propagated via an exception, which is caught and triggers respecialization. This idea could  
602 possibly be used to encode higher-level type information for return values, too. This could  
603 be used to remove redundant checks in the interpreter by simply discovering at run time  
604 that whole subexpressions conform to the type annotations.

605 **Performance Impact of Types** As seen in Section 6.1, there are cases where adding types  
606 may reduce performance, even so, in the best case this does not happen (cf. Section 5.5).

607 While the expectation is that adding more types may result in higher potential for  
608 performance issues, in the context of dynamic and adaptive compilation as used for Moth,  
609 this is not necessarily the case. Since compilers rely on various heuristics, for instance for  
610 inlining, there may be situations where a fully typed program is faster than a program  
611 with fewer types. Since the checks need to be compiled themselves, they also influence  
612 such heuristics. This means it is possible that partially typed programs may show worse  
613 performance than fully typed ones.

### 614 6.3 Threats to Validity

615 This work is subject to many of the threats to validity common to evaluations of experimental  
616 language implementations. Our underlying implementation may contain undetected bugs that  
617 affect the semantics or performance of the gradual typing checks, affecting construct validity  
618 — we may not have implemented what we think we have. Given that, our benchmarking  
619 harness run on the same implementation is subject to the same risks, thus also affecting  
620 internal validity — we may not be measuring the implementation correctly. Moth is built on

621 the Truffle and Graal toolchain, so we expect external validity there at least — we expect the  
622 results would transfer to other Graal VMs doing similar AST-based optimizations. We have  
623 less external validity regarding other kinds of VMs (such as VMs specialized to particular  
624 languages, or VMs built via meta-tracing rather than partial evaluation). Nevertheless, we  
625 expect our results should be transferable as we rely on common techniques.

626 **Generalizability** Finally, because we are working in Grace, it is less obvious that our results  
627 generalize to other gradually typed-languages. We have worked to ensure e.g. our benchmarks  
628 do not depend on any features of Grace that are not common in other gradually-typed  
629 object-oriented languages, but as Grace lacks a large corpus of programs the benchmarks  
630 are necessarily artificial, and it is not clear how the results would transfer to the kinds of  
631 programs actually written in practice. The advantage of Grace (and Moth) for this research  
632 is that their relative simplicity means we have been able to build an implementation that  
633 features competitive performance with significantly less effort than would be required for  
634 larger and more complex languages. On the other hand, more effort on optimisations could  
635 well lead to even better performance.

636 Another aspect which limits generalizability is the specific semantics of Grace. Reticulated  
637 Python, Typed Racket, and Gradualtalk have semantics that need additional runtime support,  
638 and thus, we cannot draw any conclusions without further research.

639 For languages such as Newspeak, Strongtalk, or TypeScript, where types do not have  
640 run-time semantics, one could add termination based on type errors to these languages, or  
641 simply avoid termination and report the errors after program completion as a debugging aid.  
642 For either option, our approach should apply and we would expect similar results.

## 643 **7 Related Work**

644 Although syntaxes for type annotations in dynamic languages go back at least as far as  
645 Lisp [54], the first attempts at adding a comprehensive static type system to a dynamic-  
646 ally typed language involved Smalltalk [33], with the first practical system being Bracha’s  
647 Strongtalk [17]. Strongtalk (independently replicated for Ruby [26]) provided a powerful and  
648 flexible static type system, where crucially, the system was *optional* (also known as pluggable  
649 [16]). Programmers could run the static checker over their Smalltalk code (or not); either way  
650 the type annotations had no impact whatsoever of the semantics of the underlying Smalltalk  
651 program.

652 Siek and Taha [48] introduced the term “gradual typing” to describe the logical extension  
653 of this scheme: a dynamic language with type annotations that could, if necessary, be checked  
654 at runtime. Siek and Taha build on earlier complementary work extending fully statically  
655 typed languages with a “DYNAMIC” type—Abadi *et al.* ’s 1991 TOPLAS paper [1] is an  
656 important early attempt and also surveys previous work.

657 Revived practical adoption of dynamic languages generated revived research interest,  
658 leading to the formulation of the *gradual guarantee* to characterize sound gradual type  
659 systems: informally “removing type annotations always produces a program that is still well  
660 typed” and also “evaluates to an equivalent value” [50], drawing on Boyland’s critical insight  
661 that such a guarantee must by its nature exclude code that reflects on the presence or absence  
662 of type declarations [15]. Moth ensures that the values passing through type annotations  
663 cannot be incompatible with those annotations and that type annotations cannot change  
664 program values, and Moth’s type tests consider only method names (not any further type  
665 annotations). This means that removing type annotations cannot cause a program to fail

## 15:22 Transient Typechecks are (Almost) Free

666 or change its behaviour, satisfying the informal statement of the gradual guarantee. Moth  
667 does not meet the refined formal statement of the guarantee in Siek *et al.*'s [50]'s Theorem 5,  
668 however, because Theorem 5 requires all intermediate values conform to their inferred static  
669 types. Moth only checks at explicit type declarations, not inferred intermediate types.

670 Type errors in gradual, or other dynamically checked, type systems will be detected  
671 at the type declarations, but often those declarations will not be at fault — indeed in a  
672 correctly typed program in a sound gradually typed system, the declarations cannot be at  
673 fault because they will have passed the static type checker. Rather, the underlying fault  
674 must be somewhere within the barbarian dynamically typed code *trans vallum*. Blame  
675 tracking [63, 52, 2] localizes these faults by identifying the point in the program where the  
676 system makes an assumption about dynamically typed objects, so can identify the root  
677 cause should the assumption fail. Different semantics for blame detect these faults slightly  
678 differently, and impose more or less implementation overhead [60, 51, 62].

679 The diversity of semantics and language designs incorporating gradual typing has been  
680 captured recently via surveys incorporating formal models of different design options.  
681 Chung *et al.* [22] present an object-oriented model covering optional semantics (erasure),  
682 transient semantics, concrete semantics (from Thorn [11]), and behavioural semantics (from  
683 Typed Racket), and give a series of programs to clarify the semantics of a particular language.  
684 Greenman *et al.* take a more functional approach, again modelling erasure, transient (“first  
685 order”), and behavioural (“higher order”) semantics [28], and also present performance in-  
686 formation based on Typed Racket. Wilson *et al.* take a rather different approach, employing  
687 questionnaires to investigate the semantics programmers expect of a gradual typing system  
688 [64].

689 As with languages more generally, there seem to be two main implementation strategies for  
690 languages mixing dynamic and static type checks: either adding static checks into a dynamic  
691 language implementation, or adding support for dynamic types to an implementation that  
692 depends on static types for efficiency. Typed Racket, for example, optimizes code with a  
693 combination of type inference and type declarations—the Racket IDE “optimizer coach” goes  
694 as far as to suggest to programmers type annotations that may improve their program’s  
695 performance [53]. In these implementations, values flowing from dynamically to statically  
696 typed code must be checked at the boundary. Fully statically typed code needs no dynamic  
697 type checks, and so generally performs better than dynamically typed code. Adopting a  
698 gradual type system such as Typed Racket [59] allows programmers to explicitly declare types  
699 that can be checked statically, removing unnecessary overhead. Ortin *et al.*'s [43] approach  
700 takes this to a logical extreme using a rule base to guide program specialisation at compile  
701 time based on abstract interpretation.

702 On the other hand, systems such as Reticulated Python [60], SafeTypeScript [45], and  
703 our work here, take the opposite approach. These systems do not use information from  
704 type declarations to optimize execution speed, rather the necessity to perform (potentially  
705 repeated) dynamic type checks tends to slow programs down, so here code with no type  
706 annotations generally performs better than statically typed code, or rather, code with many  
707 type annotations. In the limit, these kinds of systems may only ever check types dynamically  
708 and may not involve a static type checker at all.

709 As gradual typing systems have come to wider attention, the question of their imple-  
710 mentation overheads has become more prominent. Takikawa *et al.* [56] asked “is sound  
711 gradual typing dead?” based on a systematic performance measurement on Typed Racket.  
712 The key here is their evaluation method, where they constructed a number of different  
713 permutations of typed and untyped code, and evaluated performance along the spectrum [30].

714 Bauman *et al.* [6] replied to Takikawa *et al.*'s study, in which they used Pycket [5] (a tracing  
715 JIT for Racket) rather than the standard Racket VM, but maintained full gradually-typed  
716 Racket semantics. Bauman *et al.* are able to demonstrate most benchmarks with a slowdown  
717 of 2x on average over all configurations. Note that this is not directly comparable to our  
718 system, since typed modules do not need to do any checks at run time. Typed Racket only  
719 needs to perform checks at boundaries between typed and untyped modules, however, they  
720 use the same essential optimization technique that we apply, using object shapes to encode  
721 information about gradual types. Muehlboeck and Tate [42] also replied to Takikawa *et al.*,  
722 using a similar benchmarking method applied to Nom, a language with features designed to  
723 make gradual types easier to optimize, demonstrating speedups as more type information is  
724 added to programs. Their approach enables such type-driven optimizations, but relies on a  
725 static analysis which can utilize the type information, and the underlying types are nominal,  
726 rather than structural.

727 Most recently, Kuhlenschmidt *et al.* [36] employ an ahead of time (i.e. traditional, static)  
728 compiler for a custom language called Grift and demonstrate good performance for code  
729 where more than half of the program is annotated with types, and reasonable performance  
730 for code without type annotations.

731 Perhaps the closest to our approach are Vitousek *et al.* [60] (incl. [62, 29]) and  
732 Richards *et al.* [45]. Vitousek *et al.* describe dynamically checking transient types for  
733 Reticulated Python (termed “tag-type” soundness by Greenman and Migeed [29]). As with  
734 our work, Vitousek *et al.*'s transient checks inspect only the “top-level” type of an object.  
735 Reticulated Python undertakes these transient type checks at different places to Moth. Moth  
736 only checks explicit type annotations, while Reticulated Python implicitly checks whenever  
737 values flow from dynamic to static types. We refrain from a direct performance comparison  
738 since Reticulated Python is an interpreter without just-in-time compilation and thus per-  
739 formance tradeoffs are different. In recent experimental work, however, Vitousek *et al.* [61]  
740 have evaluated Reticulated Python's transient semantics running on top of an unmodified  
741 PyPy JIT metacompiler. These results are broadly consistent with those presented here,  
742 finding similarly small slowdowns using just the tracing JIT, and reducing those slowdowns  
743 even further when some tests are eliminated via static type inference.

744 Richards *et al.* [45] take a similar implementation approach to our work, demonstrating  
745 that key mechanisms such as object shapes used by a VM to optimize dynamic languages can  
746 be used to eliminate most of the overhead of dynamic type checks. Unlike our work, Richards  
747 implement “monotonic” gradual typing with blame, rather than the simpler transient checks,  
748 and do so on top of an adapted Higgs VM. The Higgs VM implements a baseline just-in-time  
749 compiler based on basic-block versioning [21]. In contrast, our implementation of dynamic  
750 checks is built on top of the Truffle framework for the Graal VM, and reaches performance  
751 approaching that of V8 (cf. Section 5.2). The performance difference is of relevance here  
752 since any small constant factors introduced into a VM with a lower baseline performance  
753 can remain hidden, while they stand out more prominently on a faster baseline.

754 Overall, it is unclear whether our results confirm the ones reported by Richards *et al.* [45],  
755 because our system is simpler. It does not introduce the polymorphism issues caused by  
756 accumulating cast information on object shapes, which could be important for performance.  
757 Considering that Richards *et al.* report ca. 4% overhead on the classic Richards benchmark,  
758 while we see 33%, further work seems necessary to understand the performance implications  
759 of their approach for a highly optimizing just-in-time compiler.

760 **8 Conclusion**

761 As gradually typed languages become more common, and both static and dynamically  
762 typed languages are extended with gradual features, efficient techniques for gradual type  
763 checking become more important. In this paper, we have demonstrated that optimizing  
764 virtual machines enable transient gradual type checks with relatively little overhead, and  
765 with only small modifications to an AST interpreter. We evaluated this approach with Moth,  
766 an implementation of the Grace language on top of Truffle and Graal.

767 In our implementation, types are structural and shallow: a type specifies only the names  
768 of members provided by objects, and not the types of their arguments and results. These  
769 types are checked on access to variables, when assigning to method parameters, and also on  
770 return values. The information on types is encoded as part of an object's shape, which means  
771 that shape checks already performed in an optimizing dynamic language implementation can  
772 be used to check types, too. Being able to tie checks to the shapes in this way is critical for  
773 reducing the overhead of dynamic checking.

774 Using the Are We Fast Yet benchmarks as well as a collection of benchmarks from the  
775 gradual typing literature, we find that our approach to dynamic type checking introduces an  
776 overhead of 5% (min. -13%, max. 79%) on peak performance. In addition to the results  
777 from further microbenchmarks, we take this as a strong indication that transient gradual  
778 types do not need to imply a linear overhead compared to untyped programs. However,  
779 we also see that interpreter and startup performance is indeed reduced by additional type  
780 annotations.

781 Since Moth reaches the performance of a highly optimized JavaScript VM such as V8, we  
782 believe that these results are a good indication for the low peak-performance overhead of our  
783 approach.

784 In specific cases, the overhead is still significant and requires further research to be  
785 practical. Thus, future research should investigate how the number of gradual type checks  
786 can be reduced without causing the type feedback to become too imprecise to be useful.  
787 One approach might increase the necessary changes to a language implementation, but  
788 avoid checking every variable read. Another approach might further leverage Truffle's  
789 self-specialization to propagate type requirements and avoid unnecessary checks.

790 Finally, we hope to apply our approach to other parts of the spectrum of gradual typing,  
791 eventually reaching full structural types with blame that support the gradual guarantee.  
792 This should let us verify that Richards *et al.* [45]'s results generalize to highly optimizing  
793 virtual machines, or alternatively, show that other optimizations for precise blame need to  
794 be investigated.

795 **References**

- 796 **1** Martín Abadi, Luca Cardelli, Benjamin C. Pierce, and Gordon D. Plotkin. Dynamic  
797 typing in a statically typed language. *ACM Trans. Program. Lang. Syst.*, 13(2):237–268,  
798 1991.
- 799 **2** Amal Ahmed, Robert Bruce Findler, Jeremy G. Siek, and Philip Wadler. Blame for  
800 all. In *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of*  
801 *Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011*, pages  
802 201–214, 2011.
- 803 **3** Esteban Allende, Oscar Callaú, Johan Fabry, Éric Tanter, and Marcus Denker. Gradual  
804 typing for Smalltalk. *Sci. Comput. Program.*, 96:52–69, 2014. doi:10.1016/j.scico.  
805 2013.06.006.



- 806 4 Edd Barrett, Carl Friedrich Bolz-Tereick, Rebecca Killick, Sarah Mount, and Laurence  
807 Tratt. Virtual Machine Warmup Blows Hot and Cold. *Proc. ACM Program. Lang.*,  
808 1(OOPSLA):52:1–52:27, October 2017.
- 809 5 Spenser Bauman, Carl Friedrich Bolz, Robert Hirschfeld, Vasily Kirilichev, Tobias Pape,  
810 Jeremy G. Siek, and Sam Tobin-Hochstadt. Pycket: a tracing JIT for a functional  
811 language. In *Proceedings of the 20th ACM SIGPLAN International Conference on*  
812 *Functional Programming, ICFP 2015, Vancouver, BC, Canada, September 1-3, 2015*,  
813 pages 22–34, 2015.
- 814 6 Spenser Bauman, Carl Friedrich Bolz-Tereick, Jeremy Siek, and Sam Tobin-  
815 Hochstadt. Sound gradual typing: Only mostly dead. *Proc. ACM Program. Lang.*,  
816 1(OOPSLA):54:1–54:24, October 2017.
- 817 7 Michael Bayne, Richard Cook, and Michael D. Ernst. Always-available static and dynamic  
818 feedback. In *Proceedings of the 33rd International Conference on Software Engineering*  
819 *(ICSE)*, pages 521–530, 2011.
- 820 8 Gavin M. Bierman, Martín Abadi, and Mads Torgersen. Understanding TypeScript.  
821 In *ECOOP 2014 - Object-Oriented Programming - 28th European Conference, Uppsala,*  
822 *Sweden, July 28 - August 1, 2014. Proceedings*, pages 257–281, 2014.
- 823 9 Andrew P. Black, Kim B. Bruce, Michael Homer, and James Noble. Grace: the absence  
824 of (inessential) difficulty. In *Onward! '12: Proceedings 12th SIGPLAN Symp. on New*  
825 *Ideas in Programming and Reflections on Software*, pages 85–98, New York, NY, 2012.  
826 ACM.
- 827 10 Andrew P. Black, Norman C. Hutchinson, Eric Jul, and Henry M. Levy. The development  
828 of the Emerald programming language. In *Proceedings of the Third ACM SIGPLAN*  
829 *History of Programming Languages Conference (HOPL-III), San Diego, California, USA,*  
830 *9-10 June 2007*, pages 1–51, 2007.
- 831 11 Bard Bloom, John Field, Nathaniel Nystrom, Johan Östlund, Gregor Richards, Rok  
832 Strniša, Jan Vitek, and Tobias Wrigstad. Thorn: Robust, concurrent, extensible scripting  
833 on the JVM. In *Proceedings of the ACM SIGPLAN-SIGACT Symposium on Principles*  
834 *of Programming Languages (POPL)*, pages 117–136, 2009.
- 835 12 Carl Friedrich Bolz, Antonio Cuni, Maciej Fijalkowski, and Armin Rigo. Tracing the  
836 Meta-level: PyPy’s Tracing JIT Compiler. In *Proceedings of the 4th Workshop on the Im-*  
837 *plementation, Compilation, Optimization of Object-Oriented Languages and Programming*  
838 *Systems, ICPOOLPS '09*, pages 18–25. ACM, 2009.
- 839 13 Carl Friedrich Bolz, Lukas Diekmann, and Laurence Tratt. Storage Strategies for Col-  
840 lections in Dynamically Typed Languages. In *Proceedings of the 2013 ACM SIGPLAN*  
841 *International Conference on Object Oriented Programming Systems Languages & Applic-*  
842 *ations, OOPSLA'13*, pages 167–182. ACM, 2013.
- 843 14 Carl Friedrich Bolz and Laurence Tratt. The Impact of Meta-Tracing on VM Design and  
844 Implementation. *Science of Computer Programming*, 98:408–424, February 2013.
- 845 15 John Tang Boyland. The problem of structural type tests in a gradual-typed language.  
846 In *FOOL*, 2014.
- 847 16 Gilad Bracha. Pluggable Type Systems. OOPSLA Workshop on Revival of Dynamic  
848 Languages, October 2004.
- 849 17 Gilad Bracha and David Griswold. Stongtalk: Typechecking Smalltalk in a production  
850 environment. In *OOPSLA*, 1993.
- 851 18 Gilad Bracha, Peter von der Ahé, Vassili Bykov, Yaron Kishai, William Maddox, and  
852 Eliot Miranda. Modules as Objects in Newspeak. In *European Conference on Object-*

- 853 *Oriented Programming (ECOOP)*, volume 6183 of *Lecture Notes in Computer Science*,  
854 pages 405–428. Springer, 2010.
- 855 **19** Kim Bruce, Andrew Black, Michael Homer, James Noble, Amy Ruskin, and Richard  
856 Yannow. Seeking Grace: a new object-oriented language for novices. In *Proceedings 44th*  
857 *SIGCSE Technical Symposium on Computer Science Education*, pages 129–134. ACM,  
858 2013.
- 859 **20** Craig Chambers, David Ungar, and Elgin Lee. An Efficient Implementation of SELF a  
860 Dynamically-Typed Object-Oriented Language Based on Prototypes. In *Proceedings on*  
861 *Object-Oriented Programming Systems, Languages and Applications*, OOPSLA'89, pages  
862 49–70. ACM, October 1989.
- 863 **21** Maxime Chevalier-Boisvert and Marc Feeley. Interprocedural Type Specialization of  
864 JavaScript Programs Without Type Analysis. In *30th European Conference on Object-*  
865 *Oriented Programming (ECOOP 2016)*, volume 56 of *LIPICs*, pages 7:1–7:24. Schloss  
866 Dagstuhl–Leibniz-Zentrum fuer Informatik, 2016. doi:10.4230/LIPICs.ECOOP.2016.7.
- 867 **22** Benjamin Chung, Paley Li, Francesco Zappa Nardelli, and Jan Vitek. KafKa: gradual  
868 typing for objects. In *32nd European Conference on Object-Oriented Programming,*  
869 *ECOOP 2018, July 16-21, 2018, Amsterdam, The Netherlands*, pages 12:1–12:24, 2018.  
870 doi:10.4230/LIPICs.ECOOP.2018.12.
- 871 **23** Daniel Clifford, Hannes Payer, Michael Stanton, and Ben L. Titzer. Memento Mori:  
872 Dynamic Allocation-site-based Optimizations. In *Proceedings of the 2015 International*  
873 *Symposium on Memory Management*, ISMM '15, pages 105–117. ACM, 2015.
- 874 **24** Benoit Daloz, Stefan Marr, Daniele Bonetta, and Hanspeter Mössenböck. Efficient  
875 and Thread-Safe Objects for Dynamically-Typed Languages. In *Proceedings of the 2016*  
876 *ACM International Conference on Object Oriented Programming Systems Languages &*  
877 *Applications*, OOPSLA'16, pages 642–659. ACM, 2016.
- 878 **25** Ulan Degenbaev, Jochen Eisinger, Manfred Ernst, Ross McIlroy, and Hannes Payer.  
879 Idle Time Garbage Collection Scheduling. In *Proceedings of the 37th ACM SIGPLAN*  
880 *Conference on Programming Language Design and Implementation*, PLDI'16, pages  
881 570–583. ACM, 2016.
- 882 **26** M. Furr, J.-H. An, J. Foster, and M.J. Hicks. Static type inference for Ruby. In *Symposium*  
883 *on Applied Computation*, pages 1859–1866, 2009.
- 884 **27** Adele Goldberg and David Robson. *Smalltalk-80: The Language and its Implementation*.  
885 Addison-Wesley, 1983.
- 886 **28** Ben Greenman and Matthias Felleisen. A spectrum of type soundness and performance.  
887 *PACMPL*, 2(ICFP):71:1–71:32, 2018. doi:10.1145/3236766.
- 888 **29** Ben Greenman and Zeina Migeed. On the cost of type-tag soundness. In *Proceedings*  
889 *of the ACM SIGPLAN Workshop on Partial Evaluation and Program Manipulation*,  
890 PEPM'18, pages 30–39. ACM, 2018.
- 891 **30** Ben Greenman, Asumu Takikawa, Max S. New, Daniel Feltey, Robert Bruce Findler, Jan  
892 Vitek, and Matthias Felleisen. How to evaluate the performance of gradual type systems.  
893 *Journal of Functional Programming*, 29:45, 2019. doi:10.1017/S0956796818000217.
- 894 **31** Christian Humer, Christian Wimmer, Christian Wirth, Andreas Wöß, and Thomas  
895 Würthinger. A Domain-Specific Language for Building Self-Optimizing AST Inter-  
896 preters. In *Proceedings of the 13th International Conference on Generative Pro-*  
897 *gramming: Concepts and Experiences*, GPCE '14, pages 123–132. ACM, 2014. doi:  
898 10.1145/2658761.2658776.
- 899 **32** Urs Hölzle, Craig Chambers, and David Ungar. Optimizing Dynamically-Typed Object-  
900 Oriented Languages With Polymorphic Inline Caches. In *ECOOP '91: European Con-*

- ference on Object-Oriented Programming, volume 512 of *LNCIS*, pages 21–38. Springer, 1991. doi:10.1007/BFb0057013.
- 901  
902
- 903 **33** Ralph E. Johnson. Type-checking Smalltalk. In *Conference on Object-Oriented Pro-*  
904 *gramming Systems, Languages, and Applications (OOPSLA’86), Portland, Oregon, USA,*  
905 *Proceedings.*, pages 315–321, 1986.
- 906 **34** Timothy Jones. *Classless Object Semantics*. PhD thesis, Victoria University of Wellington,  
907 2017.
- 908 **35** Timothy Jones, Michael Homer, James Noble, and Kim Bruce. Object inheritance without  
909 classes. In *30th European Conference on Object-Oriented Programming (ECOOP 2016)*,  
910 volume 56, pages 13:1–13:26, 2016.
- 911 **36** Andre Kuhlenschmidt, Deyaaeldeen Almahallawi, and Jeremy G. Siek. Efficient Gradual  
912 Typing. *CoRR*, abs/1802.06375, 2018. arXiv:1802.06375.
- 913 **37** Ole Lehrmann Madsen, Birger Møller-Pedersen, and Kristen Nygaard. *Object-Oriented*  
914 *Programming in the BETA Programming Language*. Addison-Wesley, 1993.
- 915 **38** Stefan Marr. Rebench: Execute and document benchmarks reproducibly, July 2018.  
916 Version 0.10.1. doi:10.5281/zenodo.1311762.
- 917 **39** Stefan Marr. SOMns: a newspeak for concurrency research. [https://github.com/-](https://github.com/smarr/SOMns)  
918 [smarr/SOMns](https://github.com/smarr/SOMns), 2018.
- 919 **40** Stefan Marr, Benoit Dalozé, and Hanspeter Mössenböck. Cross-Language Compiler  
920 Benchmarking—Are We Fast Yet? In *Proceedings of the 12th Symposium on Dynamic*  
921 *Languages*, DLS’16, pages 120–131. ACM, November 2016.
- 922 **41** Stefan Marr and Stéphane Ducasse. Tracing vs. Partial Evaluation: Comparing Meta-  
923 Compilation Approaches for Self-Optimizing Interpreters. In *Proceedings of the 2015*  
924 *ACM International Conference on Object Oriented Programming Systems Languages &*  
925 *Applications*, OOPSLA ’15, pages 821–839. ACM, October 2015.
- 926 **42** Fabian Muehlboeck and Ross Tate. Sound gradual typing is nominally alive and well.  
927 *Proc. ACM Program. Lang.*, 1(OOPSLA):56:1–56:30, October 2017.
- 928 **43** Francisco Ortín, Miguel Garcia, and Seán McSweeney. Rule-based program specialization  
929 to optimize gradually typed code. *Knowledge-Based Systems*, 2019. doi:[https://doi.](https://doi.org/10.1016/j.knosys.2019.05.013)  
930 [org/10.1016/j.knosys.2019.05.013](https://doi.org/10.1016/j.knosys.2019.05.013).
- 931 **44** Aaron Pang, Craig Anslow, and James Noble. What programming languages do developers  
932 use? A theory of static vs dynamic language choice. In *2018 IEEE Symposium on Visual*  
933 *Languages and Human-Centric Computing, VL/HCC 2018, Lisbon, Portugal, October*  
934 *1-4, 2018*, pages 239–247, 2018. doi:10.1109/VLHCC.2018.8506534.
- 935 **45** Gregor Richards, Ellen Arteca, and Alexi Turcotte. The VM already knew that: Lever-  
936 aging compile-time knowledge to optimize gradual typing. *Proc. ACM Program. Lang.*,  
937 1(OOPSLA):55:1–55:27, October 2017.
- 938 **46** Gregor Richards, Francesco Zappa Nardelli, and Jan Vitek. Concrete types for TypeScript.  
939 In *29th European Conference on Object-Oriented Programming, ECOOP 2015, July 5-10,*  
940 *2015, Prague, Czech Republic*, pages 76–100, 2015. doi:10.4230/LIPIcs.ECOOP.2015.  
941 76.
- 942 **47** Richard Roberts, Stefan Marr, Michael Homer, and James Noble. Toward virtual machine  
943 adaption rather than reimplementaion. In *MoreVMs’17: 1st International Workshop on*  
944 *Workshop on Modern Language Runtimes, Ecosystems, and VMs at <Programming>*  
945 *2017*, 2017. Presentation.
- 946 **48** Jeremy G. Siek and Walid Taha. Gradual typing for functional languages. In *Seventh*  
947 *Workshop on Scheme and Functional Programming*, volume Technical Report TR-2006-06,  
948 pages 81–92. University of Chicago, September 2006.

- 949 **49** Jeremy G. Siek and Walid Taha. Gradual typing for objects. In *ECOOP 2007 - Object-*  
950 *Oriented Programming, 21st European Conference, Berlin, Germany, July 30 - August 3,*  
951 *2007, Proceedings*, pages 2–27, 2007.
- 952 **50** Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, and John Tang Boyland. Refined  
953 Criteria for Gradual Typing. In Thomas Ball, Rastislav Bodik, Shriram Krishnamurthi,  
954 Benjamin S. Lerner, and Greg Morrisett, editors, *1st Summit on Advances in Programming*  
955 *Languages (SNAPL 2015)*, volume 32 of *Leibniz International Proceedings in Informatics*  
956 *(LIPIcs)*, pages 274–293. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2015.
- 957 **51** Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, Sam Tobin-Hochstadt, and Ronald  
958 Garcia. Monotonic references for efficient gradual typing. In *European Symposium on*  
959 *Programming (ESOP)*, pages 432–456, 2015.
- 960 **52** Jeremy G. Siek and Philip Wadler. Threesomes, with and without blame. In *Proceedings of*  
961 *the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages,*  
962 *POPL 2010, Madrid, Spain, January 17-23, 2010*, pages 365–376, 2010.
- 963 **53** Vincent St-Amour, Sam Tobin-Hochstadt, and Matthias Felleisen. Optimization coaching:  
964 optimizers learn to communicate with programmers. In *Proceedings of the 27th Annual*  
965 *ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and*  
966 *Applications, OOPSLA 2012, part of SPLASH 2012, Tucson, AZ, USA, October 21-25,*  
967 *2012*, pages 163–178, 2012.
- 968 **54** G.L. Steele. *Common Lisp the Language*. Digital Press, 1990.
- 969 **55** Nataliia Stulova, José F. Morales, and Manuel V. Hermenegildo. Reducing the overhead  
970 of assertion run-time checks via static analysis. In *Proceedings of the 18th International*  
971 *Symposium on Principles and Practice of Declarative Programming*, PPDP’16, pages  
972 90–103. ACM, 2016.
- 973 **56** Asumu Takikawa, Daniel Feltey, Ben Greenman, Max S. New, Jan Vitek, and Matthias  
974 Felleisen. Is Sound Gradual Typing Dead? In *Proceedings of the 43rd Annual ACM*  
975 *SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL’16,  
976 pages 456–468. ACM, 2016.
- 977 **57** Asumu Takikawa, T. Stephen Strickland, Christos Dimoulas, Sam Tobin-Hochstadt,  
978 and Matthias Felleisen. Gradual typing for first-class classes. In *Proceedings of the*  
979 *27th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems,*  
980 *Languages, and Applications, OOPSLA 2012, part of SPLASH 2012, Tucson, AZ, USA,*  
981 *October 21-25, 2012*, pages 793–810, 2012. doi:10.1145/2384616.2384674.
- 982 **58** The Clean. *Vehicle*. Flying Nun Records, FN147, 1990.
- 983 **59** Sam Tobin-Hochstadt and Matthias Felleisen. The design and implementation of Typed  
984 Scheme. In *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles*  
985 *of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12,*  
986 *2008*, pages 395–406, 2008.
- 987 **60** Michael M. Vitousek, Andrew M. Kent, Jeremy G. Siek, and Jim Baker. Design and  
988 evaluation of gradual typing for Python. In *DLS’14, Proceedings of the 10th ACM*  
989 *Symposium on Dynamic Languages, part of SPLASH 2014, Portland, OR, USA, October*  
990 *20-24, 2014*, pages 45–56, 2014.
- 991 **61** Michael M. Vitousek, Jeremy G. Siek, and Avik Chaudhuri. Optimizing and evaluating  
992 transient gradual typing. *CoRR*, abs/1902.07808, 2019. URL: [http://arxiv.org/abs/](http://arxiv.org/abs/1902.07808)  
993 [1902.07808](http://arxiv.org/abs/1902.07808), arXiv:1902.07808.
- 994 **62** Michael M. Vitousek, Cameron Swords, and Jeremy G. Siek. Big Types in Little  
995 Runtime: Open-world Soundness and Collaborative Blame for Gradual Type Systems.

- 996 In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming*  
997 *Languages*, POPL'17, pages 762–774. ACM, 2017.
- 998 **63** Philip Wadler and Robert Bruce Findler. Well-typed programs can't be blamed. In  
999 *European Symposium on Programming Languages and Systems (ESOP)*, pages 1–16,  
1000 2009.
- 1001 **64** Preston Tunnell Wilson, Ben Greenman, Justin Pombrio, and Shriram Krishnamurthi.  
1002 The behavior of gradual types: A user study. In *Dynamic Language Symposium (DLS)*,  
1003 2018.
- 1004 **65** Andreas Wöß, Christian Wirth, Daniele Bonetta, Chris Seaton, Christian Humer, and  
1005 Hanspeter Mössenböck. An Object Storage Model for the Truffle Language Implement-  
1006 ation Framework. In *Proceedings of the 2014 International Conference on Principles*  
1007 *and Practices of Programming on the Java Platform: Virtual Machines, Languages, and*  
1008 *Tools*, PPPJ'14, pages 133–144. ACM, 2014.
- 1009 **66** Thomas Würthinger, Christian Wimmer, Christian Humer, Andreas Wöß, Lukas Stadler,  
1010 Chris Seaton, Gilles Duboscq, Doug Simon, and Matthias Grimmer. Practical Partial  
1011 Evaluation for High-performance Dynamic Language Runtimes. In *Proceedings of the*  
1012 *38th ACM SIGPLAN Conference on Programming Language Design and Implementation*,  
1013 PLDI'17, pages 662–676. ACM, 2017.
- 1014 **67** Thomas Würthinger, Christian Wimmer, Andreas Wöß, Lukas Stadler, Gilles Duboscq,  
1015 Christian Humer, Gregor Richards, Doug Simon, and Mario Wolczko. One VM to Rule  
1016 Them All. In *Proceedings of the 2013 ACM International Symposium on New Ideas, New*  
1017 *Paradigms, and Reflections on Programming & Software*, Onward! 2013, pages 187–204.  
1018 ACM, 2013.
- 1019 **68** Thomas Würthinger, Andreas Wöß, Lukas Stadler, Gilles Duboscq, Doug Simon, and  
1020 Christian Wimmer. Self-Optimizing AST Interpreters. In *Proceedings of the 8th Dynamic*  
1021 *Languages Symposium*, DLS'12, pages 73–82, October 2012. doi:10.1145/2384577.  
1022 2384587.