Transitioning from Structural to Nominal Code with Efficient Gradual Typing

FABIAN MUEHLBOECK, IST Austria, Austria

ROSS TATE, Cornell University, United States of America

Gradual typing is a principled means for mixing typed and untyped code. But typed and untyped code often exhibit different programming patterns. There is already substantial research investigating gradually giving types to code exhibiting typical untyped patterns, and some research investigating gradually removing types from code exhibiting typical typed patterns. This paper investigates how to extend these established gradual-typing concepts to give formal guarantees not only about how to change types as code evolves but also about how to change such programming patterns as well.

In particular, we explore mixing untyped “structural” code with typed “nominal” code in an object-oriented language. But whereas previous work only allowed “nominal” objects to be treated as “structural” objects, we also allow “structural” objects to dynamically acquire certain nominal types, namely interfaces. We present a calculus that supports such “cross-paradigm” code migration and interoperation in a manner satisfying both the static and dynamic gradual guarantees, and demonstrate that the calculus can be implemented efficiently.

CCS Concepts: • **General and reference** → **Performance**; • **Software and its engineering** → Formal language definitions; *Object oriented languages*; **Multiparadigm languages**; **Interoperability**.

Additional Key Words and Phrases: Gradual Typing, Gradual Guarantee, Nominal, Structural, Call Tags

ACM Reference Format:

Fabian Muehlboeck and Ross Tate. 2021. Transitioning from Structural to Nominal Code with Efficient Gradual Typing. *Proc. ACM Program. Lang.* 5, OOPSLA, Article 127 (October 2021), 29 pages. <https://doi.org/10.1145/3485504>

1 INTRODUCTION

Many typed object-oriented languages are implemented markedly differently than untyped object-oriented languages. For example, many compilers for typed object-oriented languages determine the memory layouts of class instances at compile time, and similarly translate field accesses to offset memory loads during compilation; whereas objects in untyped object-oriented languages are often represented as hashtables, with the compiler translating field accesses to key-value lookups in object hashtables. While an object has a number of *structural* properties, like what fields and methods it has, many typed object-oriented languages use *nominal* type systems, permitting the compiler to use the name of a type to establish and rely upon memory-layout invariants while also abstracting such low-level details from the programmer. Beyond types, these implementation considerations can prompt typed languages to use entirely different or restricted *expressions* for specifying and constructing objects. For example, whereas untyped languages often allow one to allocate a “structural” object by manually specifying its structure through explicit fields and

Authors’ addresses: Fabian Muehlboeck, IST Austria, Klosterneuburg, Austria, fabian.muehlboeck@ist.ac.at; Ross Tate, Computer Science, Cornell University, Ithaca, New York, United States of America, ross@cs.cornell.edu.

This work is licensed under a Creative Commons Attribution-NoDerivatives 4.0 International License.

© 2021 Copyright held by the owner/author(s).

2475-1421/2021/10-ART127

<https://doi.org/10.1145/3485504>

methods, typed languages often require one to allocate only “nominal” objects of some nominal class, deriving the object’s structure from the fields and methods of the class.

These considerations have prompted us to explore a calculus and compiler for an object-oriented language with which we examine the following two research questions:

- (1) How can one extend gradual-typing concepts to make strong guarantees about bridging untyped and typed languages when there are more substantial differences between the two than just the absence or presence of type annotations?
- (2) What design considerations and implementation techniques can one employ to maintain the performance of typed object-oriented languages using compile-time memory layout of typed nominal objects and method tables while also providing principled and efficient interoperation with untyped structural objects?

More concretely, we explore how to design gradually typed object-oriented languages whose programs can transition between untyped structural and typed nominal “paradigms” while incurring low overheads and still ensuring strong desirable properties, such as the gradual guarantees [Siek et al. 2015a] and soundness, that major mixed-typed industry languages like TypeScript, Flow, Hack, and C# [Bierman et al. 2010] do not offer. Wrigstad et al. [2010] did preliminary investigation in this space with like types in Thorn, but like types are limited in that they only let code treat nominal objects as their structural counterpart, providing no way for structural objects to be treated as their nominal counterpart. While our system also limits how objects can cross the paradigm boundary, we exploit the heavy use of interfaces in typed nominal code to provide a way for structural objects to masquerade as nominal objects. In particular, unlike Nom [Muehlboeck and Tate 2017], we do not require objects to be allocated with a nominal type in order to cross the boundary—untyped code can still use flexible object manipulation to create objects that can still cross the boundary into typed code. In doing so, we provide substantially more interoperation between structural and nominal code.

In this paper, we make the following contributions. In Section 2, we illustrate what transitioning between structural and nominal paradigms within a code base could look like and how it relates to existing concepts in gradual typing. In Section 3, we provide a calculus—MonNom—wherein untyped structural code can be interwoven with typed nominal code. In Section 4, we specify a **generalized precision relation** that can bridge more substantial differences than type annotations, along with a **static gradual guarantee** that formally describes the kinds of transitions our calculus supports. In Section 5, we define a semantics for our calculus along with a **dynamic gradual guarantee** ensuring that such supported transitions do not change the run-time behavior of programs (assuming inserted type annotations correctly classify the existing behavior of the program). In Section 6, we examine the semantics in more detail and discuss their connection to **implementation techniques** we employed to preserve common optimizations of typed code (such as determining memory layouts at compile time) while also providing low-overhead interoperation with untyped code. In Section 7, we demonstrate that these techniques indeed perform well: performance typically *improves* proportionate to implementation effort when using our recommended migration strategy, and even if one strays from this strategy the worst-case overheads stay under 25%.

2 OVERVIEW

Gradual typing is roughly the ability to mix typed and untyped program components together within the same program [Gronski et al. 2006; Matthews and Findler 2007; Siek and Taha 2006; Tobin-Hochstadt and Felleisen 2006]. In this work, we broaden this to support mixing not just differences in typing discipline but also differences in paradigms that, at least in object-oriented languages, often coincide with differences in typing discipline. As an example, consider Figure 1,

Fig. 1. An example program setup with different variants for the files **twice.mn**, **fooeey.mn**, and **fooe.mn**

which provides multiple variations of two different program components, functions **twice** and **fooeey**, connected together in **main.mn**. One application of gradual typing is code migration, and so the variations are connected by arrows indicating how we would expect a program component to be migrated through the variations. The transition **fooeey-B** $\Rightarrow$ **fooeey-C** is a traditional transition in gradual typing, simply replacing occurrences of the “dynamic” type **dyn** with “static” types—in this case **int**. But the subsequent transition **fooeey-C** $\Rightarrow$ **fooeey-D** is beyond traditional gradual typing as it replaces an allocation of a structural object with an allocation of a newly defined nominal class—moving from a structural paradigm to a nominal paradigm. This work focuses on how to support and benefit from such mixing of paradigms as well as typing disciplines.

2.1 Code Migration and Gradual Guarantees

Gradual typing is often motivated in terms of code migration [Campora et al. 2017; Greenman and Felleisen 2018; Tobin-Hochstadt and Felleisen 2006; Tobin-Hochstadt et al. 2017]. That is, one might want to add types to an untyped codebase in order to increase confidence in its correctness, improve its maintainability and/or extensibility, and/or accelerate its run-time performance. Adding types throughout the entire program all at once can be an overwhelming undertaking, and so gradual typing is designed to enable one to do so bit by bit, i.e. gradually.

Given that the codebase presumably has an existing user base and possibly other codebases that depend on it, one would like to be assured that the migration to (correct/intended) types

does not break backwards compatibility, either with respect to compilability with other programs (i.e. *static* backwards compatibility) or with respect to its run-time behavior (i.e. *dynamic* backwards compatibility). To this end, the static and dynamic gradual guarantees were defined [Siek et al. 2015a], which ensure that adding “correct” type annotations to a program will preserve, respectively, its typeability and its run-time behavior.

These gradual guarantees are defined in terms of a *precision* relation ($\sqsubseteq$), where the left-hand program component is considered to be more (rather than less) “precise” than the right-hand component. Traditionally the only difference between more precise and less precise programs is that the more precise program has “static” types (such as `int`) at some syntactic positions where the less precise program has “dynamic” types (such as `dyn`). So, as an example, `foeey-C` would be considered more precise than `foeey-B`. The static and dynamic gradual guarantees then state that all program components related by the given precision relation type-check and execute identically when “all goes well”, and that “problems” occur for the less precise program only when “problems” occur for the more precise program (reflecting the fact that untyped programs are more dynamic/flexible than typed programs).

We emphasize that these gradual guarantees are parameterized by a precision relation. Although this precision relation traditionally just relates syntactic constructs componentwise, plus a $\tau \sqsubseteq \text{dyn}$ rule for type refinement, one *can* relate more kinds of programs. For example, New et al. [2019] explored incorporating β and η equivalences into the precision relation in order deduce semantic theorems in addition to syntactic theorems. In this paper, we explore incorporating the correspondences between structural object-oriented constructs and nominal object-oriented constructs into the precision relation. For example, the structural counterpart to allocating a new instance of a nominal class is allocating a record with the same fields and methods that were defined in the class. As such, we consider `foeey-D` to be more precise than `foeey-C` (provided interface `Foo` has been defined in `foo.mn`). And, more generally, every transition in Figure 1 corresponds to our precision relation (modulo syntactic conveniences such as semicolons). Since our calculus—MonNom—satisfies the static and dynamic gradual guarantees (Theorems 4.1 and 5.3), this in turn guarantees that code can be migrated according to the transitions in Figure 1.

2.2 Mixing Paradigms

Because the precision relation is defined componentwise on most program constructs, the gradual guarantees require gradually typed languages to support mixing typed and untyped program components together in the same program. For example, these guarantees ensure that if the combination of `twice-C` and `foeey-E` works correctly (which it does), then so does the combination of `twice-A` and `foeey-E` despite the fact that the two use completely different forms of method invocation: the former looks up the method implementation according to a fixed offset within an interface method table whereas the latter performs a dictionary lookup using the method’s name. But more challengingly, the combination of `twice-C` and `foeey-A` must also work correctly despite the fact that it means that a structural record dynamically extended with a field holding a functional value must be treated as if it implemented a nominal interface requiring a method of the same name whose implementation is expected to be found at a fixed offset within an interface method table. The former was possible in both Thorn and Nom, but not the latter, and that is due to the above implementation challenges that we have determined how to address efficiently in this work.

These mixed programs are often viewed as simply intermediate states on the way to the end goal of fully typed programs. But in this work we view many of these mixed programs as the end goal themselves. This is because we are bridging paradigms, and each paradigm is better suited to different tasks, so a mixed program is able to match each program component to the paradigm that best fits its purpose. Furthermore, the `dyn` type can be more than simply lack of type annotations: it

is also a means to explicitly circumvent the limitations of the static type system. That is, rather than `dyn` being just a type that has yet to be determined, we also view `dyn` as potentially conveying the programmer's intent to reason about the dynamics of the program beyond what the type system is capable of expressing. For example, a closure can be cast to an interface representing a function from $\mathbb{Z}$ to $\mathbb{Z}$ and then back to `dyn` and, unlike most (sound) gradually typed systems, our calculus permits the resulting cast value to still be supplied, say, $\mathbb{R}$ s and to return $\mathbb{R}$ s when invoked from untyped code with the understanding that the programmer may be well aware that the dynamics of the program guarantees that this particular function also operates on other numerical values even if the static type system is unaware of or unable to express that fact.

2.3 Changing Hierarchies

Consider the combination `twice-C` with `foo-B`. Because this is well typed, the static gradual guarantee requires that any less precise combination must also be well typed. Because `twice` and `foo` are separate components, a traditional componentwise precision relation would consider the `twice-C` with `foo-A` combination to be less precise. But clearly this combination should be ill-typed as `twice-C` references an undefined type `name`, an issue that does not arise in structural type systems and did not arise in Nom [Muehlboeck and Tate 2017] wherein the nominal hierarchy was fixed. This introduces a new theoretical challenge to formulating our guarantees.

We address this challenge by defining our precision relation on program components within the context of their respective nominal hierarchies. For example, our variant of the standard rule $\tau \sqsubseteq \tau'$ requires τ' to at least be valid in the nominal hierarchy of the less precise (i.e. right-hand) program.

This necessary change turns out to offer some useful flexibility. In particular, the precision relation on *expressions* must now be parameterized by the respective nominal hierarchies, and as such we can use differences in these hierarchies to guide how we bridge differences in expressions and thereby manage how we mix the structural and nominal paradigms. In particular, we allow class-instance allocations in the more precise program to refine structural-object allocations in the less precise program only when the class is declared in the more precise nominal hierarchy (and so prescribes a structure) *but not* in the less precise nominal hierarchy. This allows the compiler to assume the fields of class instances lie at fixed offsets within the instance, while also permitting developers to change a structural record into a new class provided that they simultaneously change all other structural records (if any) corresponding to the class at the same time. As such, we can support the `fooe-C` $\Rightarrow$ `fooe-D` transition in Figure 1.

2.4 Casts, Coercions, and Run-Time Overhead

In this work we focus on (sound) gradual typing—as opposed to (unsound) optional typing—meaning that type annotations, on say variables, actually provide dynamic guarantees, say about what values those variables can hold. Sound gradual typing is often associated with performance issues, often causing multiple factors of overhead, but in this work we *rely* on sound gradual typing to provide good performance, such as to guarantee that a class field can always be found at a fixed offset.

The performance problems with gradual typing are due to the casts that are inserted at the boundary points between typed and untyped code in order to dynamically enforce the contracts expected by typed code. These casts are particularly problematic in the case of types enforcing higher-order contracts, such as functions and objects. In many gradually typed implementations, such casts require one to create a proxy for the cast value that performs the relevant casts on all inputs to and outputs from the value. The sieve (of Eratosthenes) benchmark was designed to highlight this problem and was found to incur over 100x overhead in Typed Racket [Takikawa et al. 2016]. And while works such as Grift [Kuhlenschmidt et al. 2019] have proposed using

“space-efficient” coercions [Herman et al. 2010] to mitigate this issue, Feltey et al. [2018] found that—while their analogous “collapsible contracts” did address *some* pathological benchmarks—many benchmarks were unaffected by the technique. Indeed, Grift found no difference between the two techniques for sieve with coarse-grained mixing, and while Grift’s overhead for sieve is much improved compared to Typed Racket (down to 3–4x slowdown), that seems to be due to low-level implementation improvements rather than high-level strategy improvements, and it still leaves unacceptable overhead for the benchmark.

On the other hand, there are some benchmarks where Grift performs extremely well. These benchmarks tend to be numerical, specifically on floating-point values. This seems to be because Grift uses type information not just to eliminate dynamic checks but furthermore to change how values are represented. In particular, untyped Grift represents 64-bit floating-point numbers as *boxed* values allocated in the heap, whereas typed Grift represents them as *unboxed* values. Thus untyped numerical programs are constantly boxing and unboxing the values they compute with whereas typed numerical programs have few heap allocations. This connects to our observation that untyped object-oriented languages tend to be “structural” whereas typed object-oriented languages tend to be “nominal”—performance can improve significantly when type information affects *representation*. But whereas Grift gets this benefit for first-order values like numbers, we want this benefit even for higher-order values. Our calculus achieves this through two means.

For interfaces, we utilize the fact that interface-method dispatch is often already implemented using a form of indirection required to support multiple interface inheritance. So we hijack this existing indirection to make it appear as if structural objects implement arbitrary interfaces.

Classes, on the other hand, often implement field access and method dispatch through fixed offsets. That is, they use type information to optimize object representation. We cannot mimic this like we can interface methods, and adding proxies would require typed code to check whether an object is a proxy before every field access or method dispatch. Given that one benefit of types in our setting is specifically performance, we do not want to add such overhead to the high-performance path. So, while we allow structural objects to be coerced to interfaces, we do not allow them to be coerced to classes, consciously trading off some interoperation flexibility to maintain performance (while still providing a migration path through the means discussed above).

The result of this consideration of tradeoffs is that we can still support higher-order benchmarks such as sieve, since the function type corresponds to an interface, but with much improved run-time performance. In particular, most of our configurations of mixing typing disciplines and paradigms exhibit *better* run-time performance than fully untyped code—as one would like to see from gradually typed languages—and even in our worst-case configurations the overheads are only percentages rather than factors. This is achieved without any program analysis (beyond simple type-checking), heuristics, speculation, or dynamic (re-)compilation, and as such there can be high confidence our techniques would scale to large, complex programs without hard-to-predict performance cliffs. Of course, incorporating advanced techniques such as the speculative optimization used by Richards et al. [2017] or the interprocedural program analysis used by Moy et al. [2021] could improve the performance of our language even further.

3 MONNOM

Our calculus, MonNom, is built around a nominally typed, object-oriented core, which on its own already supports a version of gradual typing similar to Nom [Muehlboeck and Tate 2017]. However, MonNom also supports structural records and lambdas. But the structural types one would normally expect for these records and lambdas are not reflected in MonNom’s type system; instead, they are simply typed using the special **dyn** type from the gradual portion of MonNom’s core. On its own, the structural part of the calculus behaves like major untyped object-oriented languages such

	Class Name C	Interface Name I	Field Name f	Variable Name x
Program	$\mathcal{P} ::= \langle \mathcal{H} \mid \mathcal{S} \mid \mathcal{I} \mid e \rangle$		(Notation:	$\mathcal{P} = \langle \mathcal{H}_{\mathcal{P}} \mid \mathcal{S}_{\mathcal{P}} \mid \mathcal{I}_{\mathcal{P}} \mid e_{\mathcal{P}} \rangle$)
Hierarchy	$\mathcal{H} ::= \emptyset \mid \mathcal{H}; N \leq I, \dots$		Nominal Type	$N ::= C \mid I$
Signature	$\mathcal{S} ::= \emptyset \mid \mathcal{S}; N\{ms; \dots\} \mid \mathcal{S}; C(\vec{\tau})\{f : \tau; \dots\}$			
Implementation	$\mathcal{I} ::= \emptyset \mid \mathcal{I}; x : C(\Gamma)\{f := e; \dots \mid mb; \dots\}$			
Type	$\tau ::= N \mid \mathbb{B} \mid \mathbf{dyn}$		Method Name	$m ::= f \mid \lambda$
Type List	$\vec{\tau} ::= \emptyset \mid \vec{\tau}, \tau$		Method Signature	$s ::= (\vec{\tau}) : \tau$
Type Context	$\Gamma ::= \emptyset \mid \Gamma, x : \tau$		Method Body	$b ::= (\Gamma) \mapsto e : \tau$
Expression	$e ::= x \mid \text{let } \langle \Gamma \rangle := \langle e, \dots \rangle \text{ in } e \mid \text{false} \mid \text{true} \mid e == e$ $\mid e.f \mid e.f := e \mid e(e, \dots)$ $\mid \text{new } C(e, \dots) \mid \text{new } \lambda\langle b \rangle \mid \text{new } x := \{f := e; \dots \mid mb; \dots\}$			

Fig. 2. Grammar

as Python and JavaScript. For example, every record has a dictionary for fields that are added at run time, and fields that contain lambdas can also be treated just like methods and called directly. Unlike Nom, no type annotations are necessary even on lambdas.

More importantly, MonNom provides a rich model of interaction between its nominal and structural parts: values originating in nominal code can interact with and flow into structural code, and vice versa. Moreover, we extend the static and dynamic gradual guarantees [Siek et al. 2015a] to cover the transition from structural to nominal code. The gradual guarantee is the key property of well-behaved program migration—it allows us to refine the parts of our example program in Figure 1 along the lineages depicted there and know that the behavior of those programs will stay the same. While MonNom takes care to align concepts between the structural and nominal paradigms and typing disciplines, it does effectively contain multiple languages plus their interactions, and consequently its formalism is too large to fit into this paper. The full formalism can be found in the technical report [Muehlboeck and Tate 2021b]. Here we will elaborate upon only the key concepts.

3.1 Grammar

Figure 2 shows the syntax of MonNom. The first part of the figure specifies the structure of a MonNom program $\mathcal{P}$. A program has four components. The last component is essentially the `main` expression of the program. But that expression exists in the context of the first three components, which altogether define a traditional nominal class and interface hierarchy. In a real language, and in our actual implementation, these three components would be defined via one syntactic unit, but for the purposes of characterizing and proving the formal guarantees of MonNom we found it useful to separate them. The first component $\mathcal{H}$ defines the inheritance hierarchy between classes and interfaces. (Note that the calculus distinguishes class names C and interface names I syntactically, using N when either is applicable. Thus the syntax of $\mathcal{H}$ indicates that only interfaces can be inherited in the calculus, although this is only for simplicity as our implementation also supports single inheritance of classes.) The second component $\mathcal{S}$ defines the signatures of class and interface methods and of class constructors and fields. The third component $\mathcal{I}$ specifies how class fields are initialized from their constructor parameters, and how class methods are implemented (where the variable x represents **this** or **self** in method bodies). So, in short, given a program $\mathcal{P}$, the hierarchy $\mathcal{H}_{\mathcal{P}}$ establishes the program's space of types; the signature $\mathcal{S}_{\mathcal{P}}$ provides the information necessary for type-checking; and the implementation $\mathcal{I}_{\mathcal{P}}$ defines the overall execution of the program, with the expression $e_{\mathcal{P}}$ specifying how to kick off the execution.

$$\boxed{\mathcal{H} \vdash \tau \leq \tau} \quad \frac{}{\mathcal{H} \vdash \tau \leq \tau} \quad \frac{N \leq I_1, \dots \in \mathcal{H} \quad \mathcal{H} \vdash I_i \leq \tau}{\mathcal{H} \vdash N \leq \tau}$$

Fig. 3. Nominal ($\leq$) Subtyping

$$\boxed{\vdash \tau \sim \tau} \quad \frac{}{\vdash \tau \sim \tau} \quad \frac{}{\vdash \tau \sim \mathbf{dyn}} \quad \frac{}{\vdash \mathbf{dyn} \sim \tau} \quad \boxed{\vdash \tau \sqsubseteq \tau} \quad \frac{}{\vdash \tau \sqsubseteq \tau} \quad \frac{}{\vdash \tau \sqsubseteq \mathbf{dyn}}$$

Fig. 4. Consistency ($\sim$) and Precision ($\sqsubseteq$)

The second part of Figure 2 specifies types and type contexts, as well as method names and signatures. The types are either nominal, primitive (which in the calculus is just the booleans $\mathbb{B}$, but in our implementation includes 64-bit signed integers and 64-bit IEEE-754 floating-point numbers), and the dynamic type **dyn**. Note that MonNom has no structural types for records or lambdas—they are simply assigned the dynamic type **dyn**. Method names are either field names or the special method name λ that allows objects to be invocable directly, just like a lambda closure.

The last part specifies expressions. The first line contains a number of different standard expressions whose meaning should be unsurprising. Note, though, that the equality operator in MonNom checks for equality of references. This means that, in order to satisfy the gradual guarantee, we must ensure that transitions preserve the behavior of object identities. Furthermore, although not formally discussed in this paper, we ensure this equality has the property that returning true guarantees identical behavior, i.e. if $x == x'$ then e_t else e_f is semantically equivalent to if $x == x'$ then $e_t[x' \mapsto x]$ else e_f , which we found to conflict with devices such as the “chaperone” references used by Typed Racket [Tobin-Hochstadt and Felleisen 2006].

The second line contains field access and mutation, along with application (which invokes λ -methods). Note that (named) method invocation is expressed in the grammar as application of arguments to the result of a field access. In typed code, the type of the receiver distinguishes the two, but untyped code has no such disambiguator. Indeed, MonNom’s type system has a special rule for this particular combination, and one challenge was figuring how to design a semantics and implementation that could handle the combined case efficiently in typed settings while still satisfying the gradual guarantee with respect to untyped settings.

The third line of expressions contains the constructors for class instances, lambdas, and records, in that order. Lambdas only specify their signature and implementation and (not so importantly) have no self-reference. Records feature a variable x that, like classes, is used to represent **this** or **self** in method bodies. Records also have fields and method implementations—we can implement these more efficiently when specified as part of the initial record allocation rather than added after the allocation, and our precision relation connects the two ways of constructing records.

3.2 Type System

MonNom’s type system is built around subtyping. Without taking gradual typing into account, the *nominal* subtyping relation for MonNom is defined in Figure 3. Since we omit generics in the calculus, this relation is extremely simple.

Besides being object-oriented, MonNom is also gradually typed. At the core of gradual typing is typically what is called the *consistency* relation, shown in Figure 4. Intuitively, two types are consistent with each other if there is a way to replace the occurrences of **dyn** on both sides such that the resulting types are identical (this again becomes more interesting with generics, but even then is straightforward). Consistency is then used to relax typing rules, expressing the fundamental

$$\begin{array}{c}
\boxed{\mathcal{H} \vdash \tau \triangleleft \tau} \quad \frac{\mathcal{H} \vdash \tau \leq \tau' \quad \vdash \tau' \sim \tau''}{\mathcal{H} \vdash \tau \triangleleft \tau''} \qquad \boxed{\mathcal{H} \vdash \tau \blacktriangleleft \tau} \quad \frac{\mathcal{H} \vdash \tau \leq \tau' \quad \vdash \tau' \sqsubseteq \tau''}{\mathcal{H} \vdash \tau \blacktriangleleft \tau''}
\end{array}$$

Fig. 5. Optimistic ($\triangleleft$) and Pessimistic ($\blacktriangleleft$) Subtyping

$$\begin{array}{c}
\boxed{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \downarrow \tau} \qquad \frac{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \uparrow \tau' \quad \mathcal{H} \vdash \tau' \triangleleft \tau}{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \downarrow \tau} \\
\\
\boxed{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \uparrow \tau} \qquad \frac{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \uparrow \tau \quad \mathcal{S} \vdash \tau.\lambda(\tau_1, \dots) : \tau_\lambda \quad \forall i. \mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e_i \downarrow \tau_i}{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e.\lambda(\tau_1, \dots) \uparrow \tau_\lambda} \qquad \frac{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \uparrow \tau \quad \mathcal{S} \vdash \tau.f(\tau_1, \dots) : \tau_f \quad \forall i. \mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e_i \downarrow \tau_i}{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e.f(\tau_1, \dots) \uparrow \tau_f} \\
\\
\frac{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \uparrow \tau \quad \mathcal{S} \vdash \tau.f : \tau_f}{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e.f \uparrow \tau_f} \qquad \frac{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash b}{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash \text{new } \lambda(b) \uparrow \text{dyn}} \qquad \frac{\forall i, i'. f_i = f_{i'} \Rightarrow i = i' \quad \forall i. \mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e_i \uparrow \tau_i \quad \nexists i, i'. f_i = m_{i'} \quad \forall i, i'. m_i = m_{i'} \Rightarrow i = i' \quad \forall i. \mathcal{H} \mid \mathcal{S} \mid \Gamma, x : \text{dyn} \vdash b_i}{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash \text{new } x := \{f_1 := e_1; \dots \mid m_1 b_1; \dots\} \uparrow \text{dyn}}
\end{array}$$

Fig. 6. Expression Typing (selected rules of $\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \uparrow \tau$ from the [technical report](#))

optimism underlying gradual typing that a program is worth executing if it is plausible that the types of the eventual run-time values could match.

The gradual guarantee depends on an asymmetric variant of consistency, called *precision*, shown in Figure 4. Intuitively, one type is more precise than another if the two can be made equal by replacing occurrences of **dyn** in the less precise type. Despite what the standard notation $\vdash \tau \sqsubseteq \tau'$ might suggest, the left-hand type τ is more (not less) precise than the right-hand type τ' . As we will see in Section 4, this concept can be extended to whole programs and provides a formal means to talk about program migration from less to more precise programs.

In type-checking, subtyping commonly plays a role similar to consistency—both relax what would otherwise be a type-unification requirement. A language featuring both subtyping and consistency needs to combine them, and [Siek and Taha \[2007\]](#) showed early in the history of gradual typing that the right way of doing that is what they called *consistent* subtyping, but which we call *optimistic*¹ subtyping, defined in Figure 5. The optimism in optimistic subtyping is useful for static type-checking. At run time, however, we care about safety and thus not about whether two types *could* plausibly match, but whether the values of one are guaranteed to belong to the other. Pessimistic subtyping, also defined in Figure 5, corresponds to this latter relationship.

Figure 6 shows the most interesting expression-typing rules of MonNom. MonNom uses bidirectional type-checking. For example, application and method invocation *synthesize* ($\uparrow$) the type of the receiver, use that synthesized type to lookup the corresponding signature, and then *check* ($\downarrow$) that the arguments have the expected parameter types. Note that the lookup judgements $\mathcal{S} \vdash \tau.f : \tau_f$ and $\mathcal{S} \vdash \tau.ms$ (defined in the [technical report](#)) account for **dyn** receivers—so that they return **dyn** and, in the latter case, accept all **dyn** parameters—which means that some expressions can be type-checked in multiple ways: as a dynamically typed field access followed by a dynamically typed application, or as a dynamically typed method invocation. Such **dyn** receivers arise in MonNom

¹We use the nomenclature of Nom [\[Muehlboeck and Tate 2017\]](#) with respect to optimistic and pessimistic subtyping.

$$\boxed{\vdash \mathcal{P}} \quad \frac{\vdash \mathcal{H} \quad \mathcal{H} \vdash \mathcal{S} : \mathcal{H} \quad \mathcal{H} \mid \mathcal{S} \vdash \mathcal{I} : \mathcal{S} \quad \mathcal{H} \mid \mathcal{S} \mid \emptyset \vdash e \downarrow \mathbb{B}}{\vdash \langle \mathcal{H} \mid \mathcal{S} \mid \mathcal{I} \mid e \rangle}$$

Fig. 7. Program Typing (selected judgements from the [technical report](#))

from more than just omitted type annotations because, as mentioned before, structural objects such as lambdas and records are assigned type `dyn` rather than a structural type.

Figure 7 specifies when a MonNom program is considered valid. In particular, the inheritance hierarchy must be valid ($\vdash \mathcal{H}$). Then the signature must provide method signatures for every class and interface as well as constructor and field signatures for every class in a manner that respects inheritance of methods ($\mathcal{H} \vdash \mathcal{S} : \mathcal{H}$). The implementation must provide, for every class, a constructor initializing every field as well as a body for each method in accordance with the signature ($\mathcal{H} \mid \mathcal{S} \vdash \mathcal{I} : \mathcal{S}$). And lastly, the main expression must have Boolean type. The definition of these respective judgements is deferred to the [technical report](#) because there is nothing surprising for readers familiar with Java-like languages (though some might find the details of method inheritance with gradual types interesting).

4 TRANSITION

Now that we have an overview of the syntax and type system the MonNom programmer works within, we can more formally discuss the primary goal of the paper: guaranteeing a transition path from untyped structural code to typed nominal code. Our calculus and language are designed to enable programmers to replace a record or lambda with a corresponding (new) class and be assured that the change will preserve the behavior of the program. Similarly, programmers should be able to define interfaces describing how objects are used and transparently replace dynamic types with those new interfaces.

As we discussed earlier, key to reasoning about this transition is the precision relation. The precision relation ($\sqsubseteq$) indicates the left component is more precise than the right component. The gradual guarantee essentially says that more precise components can always be relaxed to less precise variants and the only effect on behavior would be reduction in errors. The static gradual guarantee is specifically about compile-time behavior, whereas the dynamic gradual guarantee is specifically about run-time behavior. These guarantees ensure that no surprises happen when transitioning an imprecise program to a more precise variant; the only changes that can occur are errors that can arise from incorrectly describing the invariants (e.g. types) of the program, say by annotating a variable with a type that does not accurately describe the expressions/values that get assigned to that variable.

In most existing work, the precision relation is defined on types and trivially lifted to expressions. Since we also want to model transitioning from structural to nominal code, our notion of precision is more general. Because we make structural code untyped, the static gradual guarantee is relatively uninteresting for this work. As such, here we tend to focus on the dynamic gradual guarantee, which is what requires us to develop principled run-time interactions between structural and nominal code.

Recall the example in Figure 1. In that example, one transition replaces a record with an instance of a new class. This change adds a class to the hierarchy and changes an expression in the code in a way that is not semantically equivalent because class instances are more restricted in how they can be used. Yet we can still provide a gradual guarantee between these programs, meaning we can guarantee the less precise program using a record can be used wherever the more precise program using a new class *would* work. In particular, we can guarantee that such a replacement

$$\begin{array}{c}
\boxed{\vdash \mathcal{H} \sqsubseteq \mathcal{H}} \\
\hline
\vdash \emptyset \sqsubseteq \emptyset
\end{array}
\quad
\frac{\vdash \mathcal{H} \sqsubseteq \mathcal{H}'}{\vdash \mathcal{H}; N \leq I_1, \dots \sqsubseteq \mathcal{H}'}
\quad
\frac{\vdash \mathcal{H} \sqsubseteq \mathcal{H}' \quad \forall i'. \exists i. I_i = I_{i'}' \quad \forall i, I'. \mathcal{H} \vdash I_i \leq I' \wedge \mathcal{H}' \vdash I' \implies \exists i'. \mathcal{H}' \vdash I_{i'}' \leq I'}{\vdash \mathcal{H}; N \leq I_1, \dots \sqsubseteq \mathcal{H}'; N \leq I_1', \dots}$$

$$\boxed{\vdash \mathcal{P} \sqsubseteq \mathcal{P}}
\quad
\frac{\vdash \mathcal{H}_{\mathcal{P}} \sqsubseteq \mathcal{H}_{\mathcal{P}'} \quad \vdash \mathcal{S}_{\mathcal{P}} \sqsubseteq \mathcal{S}_{\mathcal{P}'} \quad \mathcal{P} \sqsubseteq \mathcal{H}_{\mathcal{P}'} \vdash \mathcal{I}_{\mathcal{P}} \sqsubseteq \mathcal{I}_{\mathcal{P}'} \quad \mathcal{P} \sqsubseteq \mathcal{H}_{\mathcal{P}'} \vdash e_{\mathcal{P}} \sqsubseteq e_{\mathcal{P}'}}{\vdash \mathcal{P} \sqsubseteq \mathcal{P}'}$$

Fig. 8. Program Precision (selected judgements from the [technical report](#))

works provided all the interfaces the record gets treated as having are explicitly inherited by the new class, and all values assigned to fields in the record match the types of the fields of the class.

There is a novel technical caveat regarding our repeated emphasis on the fact that the class is new. Technically speaking, the dynamic gradual guarantee requires that *reducing* precision does not change program behavior. So one would expect that replacing an allocation of a class C with a record allocation with the same methods and fields would not change behavior, but this is not quite true. In particular, whereas casting an instance of class C to the type C would succeed, casting a corresponding record would fail in MonNom because we only allow interfaces to be imposed upon structural objects. Of course, if the type C occurs nowhere in the less precise program, i.e. C is *new* in the more precise program, then this difference in behavior is unobservable. This means that it is critical to incorporate the fact that, as code is migrated, so is the nominal hierarchy, and precision needs to account for these changes in the nominal hierarchy even as it reasons about more local changes throughout the program. A key contribution of this work is showing how we can adapt the existing notion of precision to account for this new dimension of code migration.

4.1 Program Precision

As mentioned, reasoning about precision in MonNom first requires reasoning about the nominal hierarchy, in particular the inheritance hierarchy. The rules for precision between inheritance hierarchies are shown in Figure 8. These rules permit the more precise program to have more nominal types than the less precise program, while also ensuring that—when comparing two nominal types that *are* present in both programs—nominal subtyping coincides in both programs. Thus one can transition a program by introducing a new nominal type so long as one also simultaneously adds that nominal type to the relevant inheritance clauses. The simultaneity is important because, in MonNom, class instances cannot be cast to interfaces they do not explicitly declare, which in turn means MonNom can generate all interface-method handlers for the class at compile time (as is standard for typed object-oriented languages) and can reject subtypings not explicit in the hierarchy (which is important for efficient/decidable type-checking/casts).

The rules for precision between signatures and implementations are elided here because, besides the ability of the more precise signature/implementation to provide details for interfaces and classes that are not present in the less precise signature/implementation, precision is simply defined componentwise in the obvious manner. More importantly, the reader should observe that the precision relation for expressions (and consequently implementations) is parameterized by the more precise program and *just* the inheritance hierarchy of the less precise program. The latter parameter enables the precision relation to determine which classes in the more precise program are not in scope for the less precise expression, and the former parameter then enables the precision relation to determine which structural definitions in the less precise expression correspond to the implementations of new classes in the more precise program, as described in more detail next.

Extensibility $\chi ::= \text{fix} \mid \text{ext}$

$$\boxed{\mathcal{P} \sqsubseteq \mathcal{H} \vdash e \sqsubseteq e}$$

$$\frac{\forall i. \mathcal{H}' \vdash \tau'_i \quad \forall i. \vdash \tau_i \sqsubseteq \tau'_i \quad \forall i. \mathcal{P} \sqsubseteq \mathcal{H}' \vdash e_i \sqsubseteq e'_i \quad \mathcal{P} \sqsubseteq \mathcal{H}' \vdash e[x_1 \mapsto x'_1, \dots] \sqsubseteq e'}{\mathcal{P} \sqsubseteq \mathcal{H}' \vdash \text{let } \langle x_1 : \tau_1, \dots \rangle := \langle e_1, \dots \rangle \text{ in } e \sqsubseteq \text{let } \langle x'_1 : \tau'_1, \dots \rangle := \langle e'_1, \dots \rangle \text{ in } e'}$$

$$\frac{\begin{array}{c} x \text{ is not free in } b \\ \mathcal{P} \sqsubseteq \mathcal{H}' \vdash x := \{ \mid \lambda b \} \sqsubseteq e' : \chi' \end{array}}{\mathcal{P} \sqsubseteq \mathcal{H}' \vdash \text{new } \lambda \langle b \rangle \sqsubseteq e'} \quad \frac{\mathcal{P} \sqsubseteq \mathcal{H}' \vdash x := \{f_1 := e_1; \dots \mid m_1 b_1; \dots\} \sqsubseteq e' : \text{ext}}{\mathcal{P} \sqsubseteq \mathcal{H}' \vdash \text{new } x := \{f_1 := e_1; \dots \mid m_1 b_1; \dots\} \sqsubseteq e'}$$

$$\frac{\neg \mathcal{H}' \vdash C \quad \forall i. \mathcal{H}' \vdash \tau'_i \quad x : C(x_1 : \tau_1, \dots) \{f_1 := e_{f,1}; \dots \mid m_1 b_1; \dots\} \in \mathcal{I}_{\mathcal{P}} \quad \forall i. \vdash \tau_i \sqsubseteq \tau'_i \quad \forall i. \mathcal{P} \sqsubseteq \mathcal{H}' \vdash e_i \sqsubseteq e'_i \quad \mathcal{P} \sqsubseteq \mathcal{H}' \vdash x := \{f_1 := e_{f,1}; \dots \mid m_1 b_1[x_1 \mapsto x'_1, \dots]; \dots\} \sqsubseteq e' : \chi'}{\mathcal{P} \sqsubseteq \mathcal{H}' \vdash \text{new } C(e_1, \dots) \sqsubseteq \text{let } \langle x'_1 : \tau'_1, \dots \rangle := \langle e'_1, \dots \rangle \text{ in } e'}$$

$$\frac{\mathcal{H}' \vdash \tau' \quad \vdash \tau \sqsubseteq \tau' \quad \mathcal{H}_{\mathcal{P}} \vdash \tau \leq \tau_x \quad \mathcal{P} \sqsubseteq \mathcal{H}' \vdash e_x \sqsubseteq e'_x \quad \mathcal{P} \sqsubseteq \mathcal{H}' \vdash e[x \mapsto x'] \sqsubseteq e'}{\mathcal{P} \sqsubseteq \mathcal{H}' \vdash \text{let } \langle x_x : \tau \rangle := \langle e_x \rangle \text{ in let } \langle x : \tau_x \rangle := \langle x_x \rangle \text{ in } e \sqsubseteq \text{let } \langle x' : \tau' \rangle := \langle e'_x \rangle \text{ in } e'}$$

$$\boxed{\mathcal{P} \sqsubseteq \mathcal{H} \vdash x := \{f := e; \dots \mid mb; \dots\} \sqsubseteq e : \chi}$$

$$\frac{x \text{ is not free in } b \quad \mathcal{P} \sqsubseteq \mathcal{H}' \vdash b \sqsubseteq b'}{\mathcal{P} \sqsubseteq \mathcal{H}' \vdash x := \{ \mid \lambda b \} \sqsubseteq \text{new } \lambda \langle b' \rangle : \text{fix}}$$

$$\frac{\forall i. \mathcal{P} \sqsubseteq \mathcal{H}' \vdash e_i[x \mapsto x'] \sqsubseteq e'_i \quad \forall i. \mathcal{P} \sqsubseteq \mathcal{H}' \vdash b_i[x \mapsto x'] \sqsubseteq b'_i}{\mathcal{P} \sqsubseteq \mathcal{H}' \vdash x := \{f_1 := e_1; \dots \mid m_1 b_1; \dots\} \sqsubseteq \text{new } x' := \{f_1 := e'_1; \dots \mid m_1 b'_1; \dots\} : \text{ext}}$$

$$\frac{\mathcal{P} \sqsubseteq \mathcal{H}' \vdash x := \{f_1 := e_1; \dots \mid m_1 b_1; \dots\} \sqsubseteq e' : \text{ext} \quad \mathcal{P} \sqsubseteq \mathcal{H}' \vdash e_f \sqsubseteq e'_f}{\mathcal{P} \sqsubseteq \mathcal{H}' \vdash x := \{f_1 := e_1; \dots; f := e_f \mid m_1 b_1; \dots\} \sqsubseteq e'.f := e'_f : \text{ext}}$$

$$\frac{x \text{ is not free in } b \quad \mathcal{P} \sqsubseteq \mathcal{H}' \vdash x := \{f_1 := e_1; \dots; f := \text{new } \lambda \langle b \rangle \mid m_1 b_1; \dots\} \sqsubseteq e' : \chi'}{\mathcal{P} \sqsubseteq \mathcal{H}' \vdash x := \{f_1 := e_1; \dots \mid m_1 b_1; \dots; f b\} \sqsubseteq e' : \chi'}$$

Fig. 9. Expression Precision (selected rules from the [technical report](#))

4.2 Expression Precision

Our example in Figure 1 showed that code migration can change expressions beyond just type annotations. Some of these differences are within a paradigm and some are across paradigms. Figure 9 shows the five most novel rules of our precision relation that formally reasons about such migration patterns.

The first rule is mostly straightforward. The key detail to note is that the types that occur in the less precise let expression are required to be valid in the provided less precise hierarchy. In prior works, the set of valid types was the same across more precise and less precise programs, and as such if a type was valid in the more precise program it was necessarily valid in the less precise program, making this requirement hold automatically. But our guarantees explicitly reason about changes in the nominal hierarchy, and as such this requirement needs to be made explicit.

The three new rules make use of the judgement $\mathcal{P} \sqsubseteq \mathcal{H}' \vdash x := \{f_1 := e_1; \dots \mid m_1 b_1; \dots\} \sqsubseteq e' : \chi'$, which indicates primarily that e' is an expression that will allocate an object containing the given fields and methods (with less precise initializers and bodies). The variable x indicates what variable denotes **this** or **self** in the method bodies, and the extensibility χ' indicates whether the structural object defined by e' is extensible with additional fields after allocation. Its first selected rule indicates

that a lambda expression can be used whenever there are no fields, only a λ method (with no self-reference), and the value does not need to be extensible. Its second selected rule indicates that a record can always be used and will be extensible. Its third selected rule permits a field to be added *after* the object is created (provided the object is extensible). And its fourth selected rule permits a method (with no self-reference) to instead be represented by a field initialized to a functional structural object. Thus this judgement captures the flexibility of structural objects.

Using this judgement, the first new rule indicates that a lambda expression can be relaxed to any structural object with a corresponding λ method (including a record), and the second new rule indicates that a record expression can be relaxed to any extensible structural object with the corresponding fields and methods.

The final new rule is one of the main contributions of the paper. It says that a class-instance allocation can be relaxed to a structural-object allocation with the same fields and methods of the class *provided* the class is new, i.e. *not* found in the less precise nominal hierarchy. This indicates that, in MonNom, the transition `fooey-C` $\Rightarrow$ `fooey-D` in Figure 1—solidifying a structural record into a nominal class—preserves the behavior of the program (provided the extensibility of the record was no longer needed, and the fields and methods were ascribed the appropriate types). It also indicates that, in MonNom, any place where a class instance would work (aside from casts to that same class) will also work with a lambda or record with the same structure as the class instance *despite* their lack of nominal structure (and, in particular, their lack of declared nominal interfaces). Thus this rule, in combination with the gradual guarantee, captures both the cross-paradigm migration and interoperability guarantees of MonNom.

The final expression-precision rule basically encodes a theorem of the system with regards to nominal subtyping. Consider the case where τ and τ_x respectively equal τ' and τ'_x , so that this rule applies when τ is a nominal subtype of τ_x . With this precision rule, the static gradual guarantee implies that restricting the type of a variable to a nominal subtype that its expression necessarily belongs to—such as in the less precise program of the final rule—always improves the typeability of the program (i.e. fewer compile-time errors). Similarly, the dynamic gradual guarantee implies that doing so always improves the executability of the program (i.e. fewer run-time errors). We will refer to these properties respectively as *static subsumption* and *dynamic subsumption*, which are closely related to the Liskov substitution principle [Liskov and Wing 1994].

4.3 The Static Gradual Guarantee

MonNom ensures a compile-time guarantee about transitioning between structural and nominal code, which is our adaptation of the static gradual guarantee [Siek et al. 2015a].

THEOREM 4.1 (STATIC GRADUAL GUARANTEE). *For all programs satisfying $\vdash \mathcal{P} \sqsubseteq \mathcal{P}'$, if $\vdash \mathcal{P}, \vdash \mathcal{H}_{\mathcal{P}'}$, and $\mathcal{H}_{\mathcal{P}'} \vdash \mathcal{S}_{\mathcal{P}'} : \mathcal{H}_{\mathcal{P}'}$ hold, then so does $\vdash \mathcal{P}'$.*

This theorem (whose proof is in the [technical report](#)) ensures that if a program is well-typed then any less precise (i.e. more dynamically typed) program is necessarily also well-typed (provided at least its inheritance hierarchy and signature are well-formed²). In particular, according to our definition of precision, this implies that a program is well-typed if it is possible to add classes and interfaces, replace all records and lambdas with allocations of classes, and replace all occurrences of `dyn` with nominal types such that the resulting program would be well-typed. In other words, any program with a viable path towards being statically well-typed is guaranteed to be gradually well-typed, even if that path requires changing structural code to nominal code.

²The issue is that signatures of methods in sub-classes/interfaces and super-classes/interfaces need to be compatible with each other, so if one decides to relax a program by replacing some type in a method signature with `dyn` in some interface, then one must make sure to similarly relax the method's signature in inheriting classes and interfaces.

Observation $o ::= \text{false} \mid \text{true} \mid \infty \mid \text{error}$

$$\boxed{\vdash \mathcal{P} \twoheadrightarrow o}$$

$$\frac{
 \begin{array}{c}
 \vdash \mathcal{P} \rightsquigarrow \check{\mathcal{P}} \quad \check{e}_1 = \check{e}_{\check{\mathcal{P}}} \quad H_1 = \emptyset \\
 \forall i < n. \check{\mathcal{P}} \vdash H_i \mid \check{e}_i \rightarrow H_{i+1} \mid \check{e}_{i+1} \quad e_n = v \quad o = v
 \end{array}
 }{
 \vdash \mathcal{P} \twoheadrightarrow o
 }$$

$$\frac{
 \begin{array}{c}
 \vdash \mathcal{P} \rightsquigarrow \check{\mathcal{P}} \quad \check{e}_1 = \check{e}_{\check{\mathcal{P}}} \quad H_1 = \emptyset \\
 \forall i < n. \check{\mathcal{P}} \vdash H_i \mid \check{e}_i \rightarrow H_{i+1} \mid \check{e}_{i+1} \quad \check{\mathcal{P}} \vdash H_n \mid \check{e}_n \rightarrow \text{error}
 \end{array}
 }{
 \vdash \mathcal{P} \twoheadrightarrow \text{error}
 }$$

$$\frac{
 \begin{array}{c}
 \vdash \mathcal{P} \rightsquigarrow \check{\mathcal{P}} \quad \check{e}_1 = \check{e}_{\check{\mathcal{P}}} \quad H_1 = \emptyset \\
 \forall i \in \mathbb{N}. \check{\mathcal{P}} \vdash H_i \mid \check{e}_i \rightarrow H_{i+1} \mid \check{e}_{i+1}
 \end{array}
 }{
 \vdash \mathcal{P} \twoheadrightarrow \infty
 }$$

Fig. 10. Program Semantics

Lowered	Program	$\check{\mathcal{P}} ::= \langle \mathcal{H} \mid \mathcal{S} \mid \check{\mathcal{I}} \mid \check{e} \rangle$	Method Body	$\check{b} ::= (\Gamma) \mapsto \check{e} : \tau$
	Type Context	$\check{\Gamma} ::= \emptyset \mid \check{\Gamma}, x$	Implementation	$\check{\mathcal{I}} ::= \emptyset \mid \check{\mathcal{I}}; x : C(\check{\Gamma}) \{f := \check{e}; \dots \mid m\check{b}\}$
	Expression	$\check{e} ::= x \mid \text{let } \langle \check{\Gamma} \rangle := \langle \check{e}, \dots \rangle \text{ in } \check{e} \mid \text{false} \mid \text{true} \mid \check{e} == \check{e}$ $\mid \check{e}.f^\delta \mid \check{e}.f^\delta := \check{e} \mid \check{e}.m(\check{e}, \dots)^\delta$ $\mid \text{new } C(\check{e}, \dots) \mid \text{new } \lambda\langle \check{b} \rangle \mid \text{new } x := \{f := \check{e}; \dots \mid m\check{b}; \dots\}$ $\mid \ell \mid \langle \ell.f \rangle \mid C(v, \dots)\{\check{e}, \dots\} \mid \text{cast}^\gamma \check{e} \text{ to } \tau \mid \text{impose}^\gamma \ell.m \text{ on } \check{e}$		
	Location	ℓ	Value	$v ::= \text{false} \mid \text{true} \mid \ell \mid \langle \ell.f \rangle$
	Guard Mode	$\gamma ::= \emptyset \mid \text{dyn}$	Dispatch Mode	$\delta ::= \langle \tau \rangle$

Fig. 11. Lowered Grammar

5 SEMANTICS

The semantics of MonNom was carefully designed to simultaneously ensure our transition guarantee and to enable an efficient implementation. Before going into some of the more detailed rules, Figure 10 provides the overall structure of how we formalize our semantics. First, the program is lowered. Second, the program state is initialized to the lowering of the main expression of the program and to the empty heap. Lastly, the program state is repeatedly reduced until either arriving at some (Boolean) value, erring, or simply running ad infinitum, each of which we consider to be the observable semantics of the program.

5.1 Lowering

We define the semantics of MonNom by translating (surface) programs to lowered programs. The grammar of these lowered programs is shown in Figure 11. Note that there is no change to the inheritance hierarchy or the signature when lowering; the only changes are to expressions and components that depend on expressions (such as the implementation).

The primary change to expressions is in the second row, where we make field access and method invocation explicitly distinct (translating application to λ -invocation), and attach *dispatch modes* to each of the forms of type-directed object-access/mutation in order to keep track of the type of the receiver the program was compiled with. Our implementation uses this type information to determine how the invocation should be implemented: a v-table lookup, an interface-method lookup, or a structural-dictionary lookup. The dispatch mode is even semantically relevant, as it affects the casting behavior of the invocation, though MonNom’s dynamic-subsumption property guarantees that using the receiver’s synthesized type imposes the least-restrictive casts compared to any of its nominal supertypes.

$$\boxed{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \downarrow \tau \rightsquigarrow \check{e}} \quad \frac{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \rightsquigarrow \check{e}}{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \downarrow \tau \rightsquigarrow \text{cast}^{\text{dyn}} \check{e} \text{ to } \tau}$$

$$\boxed{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \rightsquigarrow \check{e}}$$

$$\frac{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \uparrow \tau \quad \mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \rightsquigarrow \check{e}}{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e.f \uparrow \check{e}.f^{(\tau)}} \quad \frac{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \uparrow \tau \quad \mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \rightsquigarrow \check{e}}{S \vdash \tau.f : \tau_f \quad \mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e.f \downarrow \tau_f \rightsquigarrow \check{e}_f}$$

$$\frac{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \uparrow \tau \quad \mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \rightsquigarrow \check{e} \quad S \vdash \tau.\lambda(\tau_1, \dots) : \tau_\lambda \quad \forall i. \mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e_i \downarrow \tau_i \rightsquigarrow \check{e}_i}{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e(e_1, \dots) \rightsquigarrow \check{e}.\lambda(\check{e}_1, \dots)^{(\tau)}} \quad \frac{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \uparrow \tau \quad \mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \rightsquigarrow \check{e} \quad S \vdash \tau.f(\tau_1, \dots) : \tau_f \quad \forall i. \mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e_i \downarrow \tau_i \rightsquigarrow \check{e}_i}{\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e.f(e_1, \dots) \rightsquigarrow \check{e}.f(\check{e}_1, \dots)^{(\tau)}}$$

$$\boxed{\vdash \mathcal{P} \rightsquigarrow \check{\mathcal{P}}} \quad \frac{\mathcal{H} \mid \mathcal{S} \vdash I \rightsquigarrow \check{I} \quad \mathcal{H} \mid \mathcal{S} \mid \emptyset \vdash e \downarrow \mathbb{B} \rightsquigarrow \check{e}}{\vdash \langle \mathcal{H} \mid \mathcal{S} \mid I \mid e \rangle \rightsquigarrow \langle \mathcal{H} \mid \mathcal{S} \mid \check{I} \mid \check{e} \rangle}$$

Fig. 12. Lowering (selected rules of $\mathcal{H} \mid \mathcal{S} \mid \Gamma \vdash e \rightsquigarrow \check{e}$ from the [technical report](#))

Heap	$H ::= \emptyset \mid H; \ell \mapsto h$	Mark	$\mu ::= \text{init} \mid \text{mut}$
Heap Value	$h ::= C(v, \dots)\{v, \dots\} \mid \lambda_i b \mid \{f \mapsto_\mu v; \dots \mid mb; \dots\}_i$	Imposition	$\iota ::= \emptyset \mid \iota, I$

Fig. 13. Heap Grammar

The second change is the various constructs added in the final row. Most of these arise during execution. The only one that arises during lowering is `cast`. The cast is labeled with a *guard mode* indicating whether or not the cast can err; unguarded ($\emptyset$) casts cannot err whereas guarded (**dyn**) casts can.

Figure 12 presents the most interesting rules for lowering programs. There are two judgements for lowering expressions: one corresponding to when the typing rules *checked* for a particular type, and one corresponding to when the typing rules *synthesized* the type from the expression. For lowering checked expressions, one simply inserts a guarded cast to the expected type. For lowering unchecked expressions, note that the synthesized type of the receiver is used as the dispatch mode of the object-access/mutation operations. Note, also, that some expressions can be lowered to either a field access followed by a λ -invocation or to simply a named method invocation. In MonNom, we ensure that the non-determinism of lowering does not affect the program's observable semantics. (The proof is in the [technical report](#), which also provides a more precise theorem that ensures specifically lowering is observationally deterministic without requiring reduction to be observationally deterministic.)

THEOREM 5.1 (DETERMINISM). *For all programs satisfying $\vdash \mathcal{P}$, any two observations satisfying $\vdash \mathcal{P} \rightsquigarrow o$ and $\vdash \mathcal{P} \rightsquigarrow o'$ are necessarily equal.*

5.2 The Heap

After lowering, during execution, MonNom uses a heap, both due to the stateful nature of its objects, and to implement gradual typing efficiently. As shown in Figure 13, the heap is a (partial) mapping of locations to heap values, which are class instances, lambda closures, or records. Note that record fields have a marking μ that indicates whether the field has its initial value or has been

$$\begin{array}{c}
\boxed{\check{\mathcal{P}} \mid H \vdash v.m \rightsquigarrow_Y \check{b}} \quad \frac{\ell \mapsto \lambda_i \check{b} \in H}{\check{\mathcal{P}} \mid H \vdash \ell.\lambda \rightsquigarrow_{\text{dyn}} \check{b}} \quad \frac{\ell \mapsto \{\dots \mid m_1 \check{b}_1; \dots\}_i \in H}{\check{\mathcal{P}} \mid H \vdash \ell.m_i \rightsquigarrow_{\text{dyn}} \check{b}_i} \\
\frac{\ell \mapsto C(v_1, \dots)\{\dots\} \in H \quad x : C(x_1, \dots)\{\dots \mid m_1 \check{b}_1; \dots\} \in \check{\mathcal{I}}_{\check{\mathcal{P}}}}{\check{\mathcal{P}} \mid H \vdash \ell.m_i \rightsquigarrow_{\emptyset} \check{b}_i[x \mapsto \ell, x_1 \mapsto v_1, \dots]} \quad \frac{\check{\mathcal{P}} \mid H \vdash \ell.f \rightsquigarrow_Y \check{b}}{\check{\mathcal{P}} \mid H \vdash \langle \ell.f \rangle.\lambda \rightsquigarrow_Y \check{b}} \\
\boxed{\check{\mathcal{P}} \mid H \rightarrow H \vdash \ell.m \rightsquigarrow_Y \check{b}} \quad \frac{\ell \mapsto \{f_1 \mapsto_{\mu_1} v_1; \dots; f \mapsto_{\text{init}} v; f'_1 \mapsto_{\mu'_1} v'_1; \dots \mid m_1 \check{b}_1; \dots\}_i \in H}{\check{\mathcal{P}} \mid H \vdash v.\lambda \rightsquigarrow_Y \check{b}} \\
\frac{\check{\mathcal{P}} \mid H \vdash \ell.m \rightsquigarrow_Y \check{b} \quad H' = H[\ell \mapsto \{f_1 \mapsto_{\mu_1} v_1; \dots; f \mapsto_{\text{init}} v; f'_1 \mapsto_{\mu'_1} v'_1; \dots \mid m_1 \check{b}_1; \dots\}_i]}{\check{\mathcal{P}} \mid H \rightarrow H' \vdash \ell.f \rightsquigarrow_{\text{dyn}} \check{b}}
\end{array}$$

Fig. 14. Heap Semantics (selected judgements from the [technical report](#))

mutated, the purpose of which we will explain in Section 5.6. More importantly, lambda closures and records have a list of “imposed” interfaces that is initially empty but is expanded as they get cast to interfaces during execution. This list guarantees the heap value has any λ -method expected by these interfaces (whereas named methods are checked on demand), and it affects how the values returned by methods of the heap value are cast.

The semantics use various judgements for operating on the heap. In Figure 14, we show only the two more interesting judgements. The first judgement is for fetching the body of a value’s method. The second judgement is for fetching the body of a (typed) location’s method (potentially from a field). The details of their rules will be discussed in Section 5.6.

5.3 Values

Our calculus has only four kinds of values, though one of those—locations—indirectly represents three kinds of heap values. Up until now we have not discussed $\langle \ell.f \rangle$, which denotes a bound method. These arise from the need for a class or record with a method needing to be more precise than a record with a field containing a functional value, combined with the need for an untyped field access followed by application to be semantically equivalent to an untyped method invocation. Thus the value $\langle \ell.f \rangle$ denotes the functional value resulting from accessing a method as if it were an (untyped) field (as shown in the first rule in Figure 17). In order for transitions to preserve the behavior of ($\Rightarrow$), it is important that two separate untyped field accesses of the same method return the same value, which is why we formalize bound methods as a value rather than a heap value. That said, in our implementation we implement these as a variant of lambda closures—using other techniques to ensure uniqueness—that are artificially restricted to disallow casting in order to be faithful to the calculus.

5.4 Casts

MonNom has two casting operators: $\text{cast}^Y \check{e}$ to τ and $\text{impose}^Y \ell.m$ on τ . The first operator uses the first judgement in Figure 15 to cast the value to the expected type, and the second operator looks up the impositions on the given ℓ , then uses the third judgement in Figure 15 to determine the applicable return types the imposed interfaces expect the method m to have, and finally casts the value according to those types using the first judgement (as shown in the relevant rule in Figure 17).

The first two rules of casting simply check to see if the value already has the expected types. The second two rules specify a form of *monotonic* casts [Siek et al. 2015b; Vitousek et al. 2017]—checking if the *location* is compatible with the expected type, and then modifying the *location*’s state in the

$$\begin{array}{c}
\boxed{\check{\mathcal{P}} \mid H \rightarrow H \vdash v : \vec{\tau}} \\
\hline
\check{\mathcal{P}} \mid H \rightarrow H \vdash v : \emptyset
\end{array}
\quad
\frac{\check{\mathcal{P}} \mid H \rightarrow H' \vdash v : \vec{\tau} \quad \check{\mathcal{P}} \mid H' \vdash v : \tau}{\check{\mathcal{P}} \mid H \rightarrow H' \vdash v : \vec{\tau}, \tau}$$

$$\frac{\check{\mathcal{P}} \mid H \rightarrow H' \vdash \ell : \vec{\tau} \quad \ell \mapsto \lambda_i \check{b} \in H' \quad H'' = H'[\ell \mapsto \lambda_{i,I} \check{b}]}{\check{\mathcal{P}} \mid H \rightarrow H'' \vdash \ell : \vec{\tau}, I}
\quad
\frac{\ell \mapsto \{f_1 \mapsto_{\mu_1} v_1; \dots \mid m_1 \check{b}_1; \dots\}_I \in H' \quad \forall s. \mathcal{S}_{\check{\mathcal{P}}} \vdash I.\lambda s \implies \exists i. m_i = \lambda \quad H'' = H'[\ell \mapsto \{f_1 \mapsto_{\mu_1} v_1; \dots \mid m_1 \check{b}_1; \dots\}_{i,I}]}{\check{\mathcal{P}} \mid H \rightarrow H'' \vdash \ell : \vec{\tau}, I}$$

$$\boxed{\check{\mathcal{P}} \mid H \vdash v : \tau}
\quad
\frac{\ell \mapsto C(\dots)\{\dots\} \in H \quad \mathcal{H}_{\check{\mathcal{P}}} \mid H \vdash \ell \mapsto \iota \quad \mathcal{H}_{\check{\mathcal{P}}} \vdash C \triangleleft \tau}{\check{\mathcal{P}} \mid H \vdash \ell : \tau}
\quad
\frac{\mathcal{H}_{\check{\mathcal{P}}} \vdash \iota \triangleleft \tau}{\check{\mathcal{P}} \mid H \vdash \ell : \tau}$$

$$\boxed{\mathcal{S} \vdash (\iota).m : \vec{\tau}}
\quad
\frac{\mathcal{S} \vdash (\iota).m : \vec{\tau} \quad \mathcal{S} \vdash I.m(\dots) : \tau}{\mathcal{S} \vdash (\iota, I).m : \vec{\tau}, \tau}
\quad
\frac{\mathcal{S} \vdash (\iota).m : \vec{\tau} \quad \nexists s. \mathcal{S} \vdash I.ms}{\mathcal{S} \vdash (\iota, I).m : \vec{\tau}}$$

Fig. 15. Cast Semantics (selected judgements from the [technical report](#))

heap to permanently impose the expected type upon it. Note that in MonNom only interface types can be imposed in this manner—structural objects cannot be cast to class types. Also note that MonNom does not impose interface types upon class instances—class instances are restricted to just the explicitly inherited interfaces, both enabling their implementations to be optimized for that specific set of interfaces and enabling casting of nominal values to be implemented using efficient nominal subtyping.

When imposing an interface upon a structural object, our semantics does not check for presence or compatibility of all methods expected by the interface. We make this choice for two reasons. First, besides migration, one application we aim to serve in this work is facilitating prototyping and testing, wherein it is common for the developer to want to implement only the subset of the methods expected by the interface that is actually needed for the focused task at hand (without having to provide tedious stubs for the unneeded methods). Second, eagerly checking all these methods is costly and offers no performance when using the method-invocation techniques we discuss in Section 6.2. However, we do eagerly check that if the interface provides a λ -method then so does the structural object. The distinguished nature of that name enables some performance optimizations by using an optimized memory layout for all values with such a method, and so to be memory safe our implementation needs to eagerly check that the value has a λ -method and therefore an optimized memory layout, as we will discuss in Section 6.4.

5.5 Errors and Safety

The MonNom calculus makes a distinction between getting stuck and erring. In particular, while an error does arise from getting stuck, it can only happen in specific expressions where the implementation knows to explicitly check for such conditions. We formalize this semantics for errors in Figure 16.

Notice that nearly every potentially erroneous redex involves either a dynamic dispatch mode or a dynamic guard mode, reflecting the fact that statically typed code and unguarded casts should not err. There is, however, one exception: invocation of *named* methods of *interfaces*. This is due to the fact that our casts of structural objects to interfaces do not check for the presence or compatibility

$$\begin{array}{c}
\text{Potentially Erroneous Redex } \varepsilon ::= v.f^{\langle \text{dyn} \rangle} \mid v.f^{\langle \text{dyn} \rangle} := v \mid \ell.f(v, \dots)^{\langle I \rangle} \mid v.m(v, \dots)^{\langle \text{dyn} \rangle} \\
\mid \text{cast}^{\text{dyn}} v \text{ to } N \mid \text{impose}^{\text{dyn}} \ell.m \text{ on } v \\
\\
\boxed{\check{\mathcal{P}} \vdash H \mid \check{e} \rightarrow \text{error}} \quad \frac{\nexists H', \check{e}'. \check{\mathcal{P}} \vdash H \mid \varepsilon \rightarrow H' \mid \check{e}'}{\check{\mathcal{P}} \vdash H \mid E[\varepsilon] \rightarrow \text{error}}
\end{array}$$

Fig. 16. Error Semantics (where Evaluation Contexts E are defined in the [technical report](#))

of named methods. Nonetheless, we can implement this without adding significant dynamic checks to typed invocation by using the technique described in Section 6.2.

By distinguishing erring from getting stuck, we can formalize the type safety of MonNom.

THEOREM 5.2 (SAFETY). *For any program satisfying $\vdash \mathcal{P}$, there exists an observation satisfying $\vdash \mathcal{P} \Rightarrow o$.*

This theorem (whose proof is in the [technical report](#)) indicates that well-typed programs only get stuck if they reach a potentially erroneous redex. We cannot ensure that well-typed programs do not err due to the presence of dynamic typing in MonNom, though one can show that the “statically typed” subset of MonNom would be free of errors.

5.6 Invocation

Being an object-oriented language, method invocation is a core feature of MonNom. It is important to remember that every invocation has two sides: the caller and the callee. In MonNom, lowering takes care of the caller side by determining the dispatch mode to use and inserting casts of the arguments to the parameter types expected by that dispatch mode. Thus the three rules for invocation shown in Figure 17 focus on the callee side.

Each of these three rules uses three different (but closely related) judgements to operate on the heap (recall Figure 14) in order to lookup the body of the given method. The first rule for untyped (i.e. $\langle \text{dyn} \rangle$) invocation simply looks up a body of a *method* directly provided by the receiver. The second rule for untyped invocation handles the case where the receiver directly provides a *field* of the given name, extracting the method body from the λ -method implementation provided by that field’s value. Together these two cases ensure that the two ways to invoke a method—directly, or applying to the result of a field access—work equivalently in *untyped* code. The rule for *typed* (i.e. $\langle N \rangle$) invocation uses a judgement that effectively combines these two cases, but with an extra step: if the implementation is provided by a (record) field, the field is checked to have never been mutated. Common implementations of and optimizations for method invocation in major nominal object-oriented languages assume methods are immutable, so this extra step dynamically asserts that those assumptions are valid for the structural interoperation at hand. In particular, whereas with the untyped expression $e.f(e_1, \dots)$ one determines the value of the field before evaluating the arguments, we found it useful for typed method invocation to lookup the method while simultaneously supplying its *already* evaluated arguments, and so for these to be equivalent—so that the dynamic gradual guarantee holds—we need the field to not have been mutated in the interim.

After the method body has been looked up, it is supplied the respective arguments, casting them to the types expected by the method body. For untyped invocation, these casts are always guarded, where for typed invocation they are unguarded if the receiver was a class instance. Furthermore, the value returned by the call is checked to be compatible with the expected return type. For untyped invocation, this check is a no-op since the caller is untyped (though still having the no-op in the

$$\boxed{\check{\mathcal{P}} \vdash H \mid \check{e} \xrightarrow{H} H \mid \check{e}} \quad \frac{\check{\mathcal{P}} \mid H \vdash \ell.f \rightsquigarrow_Y \check{b}}{\check{\mathcal{P}} \vdash H \mid \ell.f^{(\text{dyn})} \xrightarrow{\mathcal{O}} H \mid \langle \ell.f \rangle} \quad \frac{\check{\mathcal{P}} \mid H \vdash v.m \rightsquigarrow_Y (x_1 : \tau_1, \dots) \mapsto \check{e} : \tau}{\check{\mathcal{P}} \vdash H \mid v.m(v_1, \dots)^{(\text{dyn})} \xrightarrow{\mathcal{O}} H \mid \text{cast}^{\mathcal{O}} \text{let } \langle x_1, \dots \rangle := \langle \text{cast}^{\text{dyn}} v_1 \text{ to } \tau_1, \dots \rangle \text{ in } \check{e} \text{ to dyn}} \\
\frac{\check{\mathcal{P}} \mid H \vdash \ell.f \mapsto v \quad \check{\mathcal{P}} \mid H \vdash v.\lambda \rightsquigarrow_Y (x_1 : \tau_1, \dots) \mapsto \check{e} : \tau}{\check{\mathcal{P}} \vdash H \mid \ell.f(v_1, \dots)^{(\text{dyn})} \xrightarrow{\mathcal{O}} H \mid \text{cast}^{\mathcal{O}} \text{let } \langle x_1, \dots \rangle := \langle \text{cast}^{\text{dyn}} v_1 \text{ to } \tau_1, \dots \rangle \text{ in } \check{e} \text{ to dyn}} \\
\frac{\check{\mathcal{P}} \mid H \rightarrow H' \vdash \ell.m \rightsquigarrow_Y (x_1 : \tau_1, \dots) \mapsto \check{e} : \tau}{\check{\mathcal{P}} \vdash H \mid \ell.m(v_1, \dots)^{(\text{N})} \xrightarrow{\mathcal{O}} H' \mid \text{impose}^Y \ell.m \text{ on let } \langle x_1, \dots \rangle := \langle \text{cast}^Y v_1 \text{ to } \tau_1, \dots \rangle \text{ in } \check{e}} \\
\frac{\check{\mathcal{P}} \mid H \rightarrow H' \vdash v : \tau \quad \mathcal{H}_{\check{\mathcal{P}}} \mid H \vdash \ell \mapsto \iota \quad \mathcal{S}_{\check{\mathcal{P}}} \vdash (\iota).m : \bar{\tau} \quad \check{\mathcal{P}} \mid H \rightarrow H' \vdash v : \bar{\tau}}{\check{\mathcal{P}} \vdash H \mid \text{cast}^Y v \text{ to } \tau \xrightarrow{\mathcal{O}} H' \mid v} \quad \frac{\check{\mathcal{P}} \vdash H \mid \text{impose}^Y \ell.m \text{ on } v \xrightarrow{\mathcal{O}} H' \mid v}{\check{\mathcal{P}} \vdash H \mid \check{e} \xrightarrow{H''} H' \mid \check{e}' \quad \nexists \ell, h, h'. \ell \mapsto h \in H' \wedge \ell \mapsto h' \in H''} \\
\boxed{\check{\mathcal{P}} \vdash H \mid \check{e} \rightarrow H \mid \check{e}} \quad \frac{\check{\mathcal{P}} \vdash H \mid \check{e} \xrightarrow{H''} H' \mid \check{e}' \quad \nexists \ell, h, h'. \ell \mapsto h \in H' \wedge \ell \mapsto h' \in H''}{\check{\mathcal{P}} \vdash H \mid \check{e} \rightarrow H'; H'' \mid \check{e}'}$$

Fig. 17. Lowered-Expression Semantics (selected rules from the [technical report](#))

formal semantics facilitates the proof of the dynamic gradual guarantee). For typed invocation, if the receiver was structural then we need to check that the returned value matches the return types expected of all the interfaces that have been imposed upon the receiver. This necessarily includes the interface specified in dispatch mode, guaranteeing type safety. But it is important that we cast the value to *all* interfaces imposed on the location, as imposing *just* the interface the invocation happened to be typed with would fail to satisfy dynamic subsumption.

5.7 The Dynamic Gradual Guarantee

At last we can formally state our run-time guarantee about transitioning between structural and nominal code with a property known as the dynamic gradual guarantee [Siek et al. 2015a].

THEOREM 5.3 (DYNAMIC GRADUAL GUARANTEE). *For all programs satisfying $\vdash \mathcal{P}$, $\vdash \mathcal{P}'$, and $\vdash \mathcal{P} \sqsubseteq \mathcal{P}'$, any observation satisfying $\vdash \mathcal{P} \Rightarrow o$ either also satisfies $\vdash \mathcal{P}' \Rightarrow o$ or is **error**; and for any observation satisfying $\vdash \mathcal{P}' \Rightarrow o$, either $\vdash \mathcal{P} \Rightarrow o$ also holds or $\vdash \mathcal{P} \Rightarrow \mathbf{error}$ holds.*

This theorem (whose proof is in the [technical report](#)) ensures that well-typed more precise and less precise programs are observably equivalent except that the former can err when the latter does not. As mentioned earlier, statically typed MonNom is error free, which in turn means this theorem implies that a program does not err if it is possible to add interfaces, replace all records and lambdas with classes implementing those interfaces, and replace all occurrences of **dyn** with nominal types such that the resulting program statically type-checks. In other words, in combination with the static gradual guarantee, any program with a viable path towards being statically well-typed (and so necessarily never erring) is guaranteed to be gradually well-typed and to never err, even if that path requires changing structural code to nominal code.

Of course, this is great in theory, but in practice gradual typing has a history of significant issues with large overheads caused by the casts ensuring safety [Takikawa et al. 2016]. Next we demonstrate that the design of MonNom enables an efficient implementation.

6 IMPLEMENTATION

We have implemented MonNom with an ahead-of-time compiler using LLVM as the backend. Our implementation extends the calculus in many ways, including more primitives as well as a standard library for common data structures. Importantly, our implementation adds generics to MonNom. This makes implementing monotonic casts efficiently more challenging as we must account for type arguments in impositions (particularly in the function interfaces provided by our standard library and used in the relevant benchmarks). In the interest of space, though, here we discuss only the techniques that we think are most critical to achieving our performance for features of MonNom present in the calculus.

6.1 Value Representation

We represent primitive values differently depending on whether they are statically or dynamically typed. In particular, primitive values are unpacked when statically typed and packed or boxed when dynamically typed. For dynamically typed values, the lower two bits differentiate between (aligned) references and the two kinds of packed primitives—integers and floating-point numbers—adapting the packing technique from [Chambers et al. \[1989\]](#). If the low bits are 11, the packed value represents the signed integer resulting from arithmetic-right-shifting the value by 2. If the low bits are 01 or 10, the packed value represents the IEEE-754 floating-point number resulting from funnel-shifting the value right by 3 (moving the third bit to be the sign bit); the effect of this is that all floats with small-magnitude exponents (i.e. signed 10-bit) can be packed rather than boxed. (When boxing floats, we memoize the statically allocated values for positive and negative zero.) This packing scheme prevents the performance of both integer-intensive and floating-point-intensive programs from degrading severely due to heap allocations in untyped and mixed programs.

6.2 Typed Method Invocation and Call Tags

One of the key challenges to implementing MonNom is supporting typed method invocations on structural objects that were cast to interfaces they were not created to support, ideally without slowing down the fast path where the receiver is typed.

Originally, we implemented interface-method dispatch using interface tables. In this strategy, the object descriptor (which also provides, say, the v-table of class methods) provides the list of interfaces implemented by the object, and in each case providing a method table pointing to the implementations of each interface method. When we cast a structural object, we extended its interface table with a newly allocated method table for the interface that was filled with stubs that would be filled on demand. However, these allocations were costly, and we found ourselves employing speculation and heuristics in an attempt to precompute these tables, which we worried might not scale well and could lead to unpredictable performance cliffs.

More recently, we developed an extension of interface-method tables [[Alpern et al. 2001](#)] that bridges this gap—an extension we refer to as *call tags*. An interface-method table provides a fixed-sized array of code pointers, and every interface method is globally assigned an index where its implementation lies in this array. But an object can implement multiple interface methods assigned to the same index. To address this conflict, when one calls the code pointer at the corresponding index, in addition to the explicit arguments of the method one also passes the identifier of the interface method. That way, if the receiver has multiple implementations corresponding to that index, it can first switch on that identifier before doing anything else, which works even if those methods have different arities, signatures, and even calling conventions.

In typed languages, the type system guarantees that the passed identifier will be recognized by the switching code. But in the untyped setting, we cannot make such a guarantee, and matters get

very complicated when one realizes that even the size of the stack frame cannot be known without recognizing that identifier. To resolve this, we made it so that the method identifier—i.e. call tag—itself provides the code pointer to jump to when the call tag is not recognized—leaving the arguments and return address untouched. In this way, the “fall-back” handler for an interface method can convert all the arguments to their untyped representations, then invoke the corresponding untyped method on the argument specifying the receiver, and cast the returned value to the types expected by the receiver’s imposed interfaces.

By using call tags, we get an implementation strategy wherein the fast path for typed invocation on typed receivers is practically untouched, and the slow path has only small indirections employing the necessary coercions. Using this, we can simply have an object descriptor for each class and for each allocation site of a structural object—no need for speculation or heuristics.

6.3 Untyped Method Invocation

An untyped method invocation is semantically equivalent to a field access followed by an argument application. However, implementing one this way would be inefficient when the object directly provides a method implementation because it creates an intermediate bound-method value. But implementing it simply as a method invocation would be incorrect in the case where the method is provided by a field whose value is changed during evaluation of the arguments.

To address these problems, we implement untyped method invocation through two stages. The first stage is executed just before argument evaluation. Its job is to determine the call-tag handler *and* the receiver to use later. Then the arguments are evaluated. Afterwards, the second stage simply calls the previously determined call-tag handler with a call tag indicating the number of arguments, providing the previously determined receiver and the values of the arguments. Thus the first stage does the bulk of the work.

For the first stage, during compilation a hashtable is inlined directly into the object descriptor. This hashtable lists all of the named methods directly provided by the object as well as some of the fields directly provided by the object upon allocation. For named methods, the entry provides a pointer to the precompiled call-tag handler to use for the second stage. For selected fields, the entry provides the offset of the field within the object. Only fields with non-primitive types are listed, which both avoids fields that cannot provide a method implementation and ensures the values of listed fields are necessarily references if and only if their lower two bits are 00. Executing the first stage involves first looking up the corresponding entry in the receiver’s object descriptor. If a method entry is found, then the specified call-tag handler and the same receiver are forwarded to the second stage. If a field entry is found, then the field’s value is checked to be a reference with a λ -method, in which case that λ -method handler and that field’s value are forwarded to the second stage. If no entry is found and the object is a record, its additional-field dictionary is searched for the appropriate field and handled as with built-in fields. Otherwise, the invocation fails.

6.4 λ -Methods

Our implementation (and consequently casting semantics) employs special treatment for λ -methods. If an object implements a λ -method, room for a call-tag-handling function pointer is reserved at the head of the object itself, rather than in its descriptor. This prevents the need to load the object descriptor when invoking λ -methods.

On the callee side, for class instances the function switches on the corresponding call tags for all implemented interfaces, along with a special call tag for just the class, as well as the call tag of the appropriate arity that is used for untyped method invocation. Call tags proved to be helpful here because our implementation supports nominal subtypings such as `Int  $\leq$  Object` as well as generics, and as such compatible method signatures can still represent objects differently (e.g. boxed

vs. unboxed integers); having a call tag for each interface then enables quick conversions before jumping to the main implementation of the method. (Note that we explored the option of restricting method-signature compatibility in order to ensure consistent object representation and simply use standard functions, but our experiments found that standard functions offered no performance improvements over call tags.) For structural objects, the function only switches on the (untyped) call tag of the appropriate arity, but if no match is found it jumps to the fall-back handler provided by the call tag, just as with named methods.

On the caller side, typed invocations of λ -methods simply load the function pointer from the standard offset within the object (forgoing the object descriptor entirely) and call it with the appropriate call tag. Because our casting semantics for casting to interfaces with a λ -method eagerly checks that the structural object has some such method, this operation is guaranteed to be safe even for structural objects masquerading as their imposed interfaces. For untyped invocations though, we first have to check if such a λ -method exists. One way to do so would be to put a flag in the object itself, but that would waste memory. Another way to do so would be to put a flag in the object descriptor, but the whole point of the optimization is to avoid loading from the descriptor. So we encoded a flag in the address of the descriptor. In particular, we use memory alignment to ensure that the address of the descriptor has a specific low bit set if and only if the object has a λ -method and consequently the corresponding function pointer at the standard offset. (We could have also taken advantage of the fact that alignment ensures the low bits of any object-descriptor address are always unused, but this would require adding bit-masking operations to typed code.)

6.5 Fields

Typed field-access/mutation is implemented using predetermined offsets. For untyped field access/mutation, if the field is typed then we have to account for potential casts and change in representation, and if the field belongs to a record then we need mutations of the field to mark it appropriately. Rather than use an inlined data structure listing field information in the object descriptor and a corresponding generalized read/write operation, the object descriptor provides two function pointers accepting a receiver and a field identifier. The former returns the (packed) value of identified field. The latter accepts an additional (packed) value that it updates the identified field to. For classes and records, the implementations of these functions switch to the fields known at compile time. For records, if the identifier is not recognized then the additional-fields dictionary is searched for a corresponding entry.

7 EVALUATION

7.1 Methodology

Following Takikawa et al. [2016], it is important to evaluate the performance of gradually typed programs in different stages of program evolution to see if run-time checks introduce prohibitive worst- or average-case overheads. This is often done by taking a fully typed program and removing type annotations from modules one by one (in varying order), and then evaluating program performance on each of the intermediate configurations. In the case of the work presented here, the picture is a little more complicated as we have an additional axis of migration—structural vs. nominal. Furthermore, type annotations involving class types can only be added after the corresponding records have been converted to class allocations. This means both that there are more variables to toggle and that not all these toggles are independent. In particular, from a fully typed nominal program we generate different configurations by

- (1) removing type annotations from the fields, method signatures, and method bodies of a class (including static methods), except those necessary for inheriting typed methods,

- (2) removing a class (and all type annotations referring to that class) entirely and replacing uses of its constructor with record or lambda expressions (preferring the latter when possible) whose method signatures and bodies are typed according to the typing discipline of the code they occur within,
- (3) and/or removing an interface (and all type annotations referring to that interface) entirely.

We then run all the valid generated configurations and compare their running times. The expectation is that running times decrease while going from left to right, though the mixed configurations in the middle may suffer from overheads due to casts and cross-paradigm indirections.

The following experiments were run on an Intel® Core™ i7-8700 CPU with 16GB main memory running Windows 10 on minimal activity. For each configuration, the reported running time is the average over 10 runs (each of which were measured after first performing warm-up runs).

7.2 Benchmarks

MonNom is a new language with a unique set of features for which we built everything from the ground up. As such, there exists no corpus of programs in the language, and its standard library is rather small. We keep adding new code and benchmarks as we develop MonNom further, but for now, we present the results on three benchmark programs (the MonNom code for which can be found in the [technical report](#)): sieve, intersort, and float.

7.2.1 sieve. sieve is a key benchmark originally developed by [Takikawa et al. \[2016\]](#) to represent a worst-case scenario for higher-order casts. It originally consisted of two heavily interacting modules that may each be typed or untyped, resulting in four configurations. Mixed configurations of the program feature a particularly large number of interactions (and therefore casts) between typed and untyped code compared to the rest of the work of the program, making it an important benchmark for gradual-typing implementations. Due to its small size, it is feasible to increase the granularity to individual classes and interfaces and vary each according to the possible variations discussed above. One of the modules contains three lambdas that Nom [\[Muehlboeck and Tate 2017\]](#) had to replace with classes implementing a particular interface.³ The fully typed MonNom program is almost identical to the Nom program, except that it represents (nullary) functions with a *generic* (nullary-)function (standard-library) interface rather than an interface monomorphized to (nullary) functions returning integers. Furthermore, MonNom can also replace functional classes (implementing functional interfaces) with lambdas. Thus, in addition to the two main classes corresponding to the two original modules, we have three additional classes to generate various configurations with. Altogether, we consider 272 configurations of sieve for MonNom.

Due to its widespread use in the literature, sieve is also a good candidate to compare our results to related work, in particular with Grift [\[Kuhlenschmidt et al. 2019\]](#), Monotonic Grift [\[Kuhlenschmidt et al. 2018\]](#), Typed Racket [\[Tobin-Hochstadt and Felleisen 2006\]](#) (CS 8.0), HiggsCheck [\[Richards et al. 2017\]](#), Transient⁴ Reticulated Python [\[Vitousek et al. 2014, 2017\]](#) (PyPy 7.3.5), and Nom [\[Muehlboeck and Tate 2017\]](#). There are substantial differences between these languages and MonNom, and as such not all MonNom configurations have a good corresponding configuration in each of these languages. So, for each language, we selected MonNom configurations for which there were good

³In the original benchmark, there was also a temporary pair value, which MonNom could implement as either a record or class. However, we found the version of this benchmark included in the artifact for Grift did not use this pair value, instead accessing the values contained in it more directly. Through profiling, we found that the original version of this benchmark caused the many implementations to spend most of their time collecting garbage due to these temporarily created pairs. That issue is unrelated to gradual typing, so we went with the version in Grift's suite that focused more on gradual-typing-specific performance considerations. Note that the modified version is still memory-intensive, and as such memory-management techniques—particularly for short-lived objects—do have notable effects on absolute performance.

⁴We did not evaluate Monotonic Reticulated Python because we could not get it to support our benchmarks.

Fig. 18. Measurements for sieve

corresponding configurations that were also representative of how these implementations had previously been evaluated. We did the same for some major industry object-oriented languages: Java (HotSpot build 13+13), C# (CoreCLR 6.0.100-preview.7.21379.14), and JavaScript (Node.js 10.19.0).

Figure 18 shows the results of our measurements for sieve. The upper left-hand plot shows the absolute running times as a scatter-plot in which we group configurations by the number of steps they are away from the fully typed configuration, plotting them from left (untyped structural) to right (typed nominal). For each configuration we evaluated in other languages, we plot its measurement in the same column as its corresponding MonNom configuration.⁵⁶

The upper right-hand plot shows the style of usability diagram that is common in the literature, showing the percentage of configurations that incur less than some amount of overhead compared to the fully untyped program. We see that Typed Racket has improved substantially over the years since Takikawa et al. [2016] observed more than 100x overhead on a mixed configuration of this benchmark. However, both it and Grift still incur overheads of several times the running time of the fully untyped configuration, while Nom’s and MonNom’s worst-case overheads are measured in percentages.

The lower left-hand plot of Figure 18 shows the numbers for the MonNom experiments in more detail. Some configurations are shown in grey while others are shown in black. The black configurations are those that follow the following “recommended” migration strategy:

⁵Though we evaluated sieve in HiggsCheck, its measurements are not shown. We found that the translation of sieve used by Richards et al. [2017] did not faithfully recreate the quantity or quality of type-boundary crossings that Takikawa et al. [2016] designed the benchmark to evaluate. When we ran HiggsCheck on our own translation, it failed to complete in any configuration due to resource exhaustion, which is why HiggsCheck is not represented in the plot even though Node.js is.

⁶Transient Reticulated Python’s measurements are not shown because, though there was little relative variation across configurations, their absolutes far exceeded the upper bound of the current figure across all configurations.

Fig. 19. Measurements for intersort

- (1) First, if a record’s fields are frequently accessed by *other* objects’ methods, turn the record into a class (though not necessarily with typed fields or methods) and add the appropriate class type annotations to the other objects’ methods.
- (2) Second, if a class’s methods call methods of another class, add type annotations to the other class before adding type annotations to this class.

The plot—and its counterparts for the other benchmarks—illustrates that this simple migration strategy for MonNom avoids the most significant pitfalls of gradual-typing overhead, and even typically leads one towards performance *improvements* proportionate to their migration effort.

The lower right-hand plot shows four sets of lines. Two lines show what percentage of “recommended” configurations exhibit a given amount of overhead, whereas another two lines show what percentage of all configurations exhibit a given amount of overhead. The blue lines plot overhead relative to the fully untyped program, i.e. the dashed blue line in the lower left-hand plot, which is the traditional metric established by Takikawa et al. [2016]. However, we believe this metric does not match one’s expectation that programs should generally get faster as more types are added. As such, the green lines plot overhead relative to a linear baseline between the fully typed and fully untyped program, i.e. the dotted green line in the lower left-hand plot, which punishes failure to obtain optimizations in proportion to the migration effort put in.

7.2.2 intersort. While sieve represents a stress test for casting lambdas, intersort is a benchmark we designed to model a more typical scenario of specifically object-oriented code migration by sorting a data structure using just its interface methods. It consists of slightly more components than sieve, which we therefore group into four modules:

- (1) generic interfaces for lists and bidirectional iterators,
- (2) generic classes implementing those interfaces with doubly-linked lists,

Fig. 20. Measurements for float

- (3) an implementation of Quicksort using bidirectional iterators,
- (4) and a main class generating 100,000 pseudo-random integers and then running Quicksort.

The latter two consist of only static methods and as such cannot be turned into structural objects, but the list interfaces can be removed from the program, and the classes implementing doubly-linked lists can be turned into records. We transition each module as a whole unit, resulting in 36 possible configurations.

Our measurements for intersort are shown in Figure 19. They illustrate that MonNom’s performance generally improves proportionate to migration effort—especially when using our recommended migration strategy—despite the heavy use of classes and higher-order interfaces.

7.2.3 float. float is a benchmark we have taken from the set of benchmarks used to evaluate Nom and transient Reticulated Python [Vitousek et al. 2014, 2017]. The benchmark generates a large list of triples of floats (“points in 3D space”), iterates the list to normalize them, and then folds the list. This benchmark does not make heavy use of methods, but does make heavy use of floating-point numbers and of field accesses.

Our measurements for float are shown in Figure 20. For both MonNom and (both versions of) Grift, we see two cleanly separated sets of configurations. The distinguishing characteristic is whether the 3D-point data-structure is typed or not, with a notable loss when it is untyped. However, this loss is much more drastic for Grift than for MonNom. Despite being evaluated on primarily floating-point-intensive benchmarks, Grift boxes all floating-point numbers onto the heap, whereas MonNom only boxes floating-point numbers with large-magnitude exponents. We suspect this choice was critical to MonNom’s success on this floating-point-intensive benchmark.

7.3 Threats to Validity

7.3.1 Overlooked Configurations. While we believe the operations above represent reasonable and realistic chunks of work that can and will often be done at once, it is possible that we missed plausible configurations with far worse running times.

7.3.2 Small Corpus. As with every new language and custom compiler/runtime, there is not any existing code that we could run and the standard library is minimal. In addition, with the larger variety of program variations and cross-program constraints on which variations are permissible, creating the benchmarks themselves is a larger task than the usual combination of fully annotated and fully unannotated files. As such, we thought it important to at least include a known worst-case benchmark (sieve) and a realistic benchmark (intertsort, and possibly float), but our particular choices might not be sufficiently representative.

7.3.3 Usual Experimental Issues. Although benchmarks were run with consistent settings and without being affected by other concurrent processes, it is possible we failed to control for other external factors we were unaware of.

8 CONCLUSION

Nominality is a well-known device towards supporting efficient implementation, even in the context of gradual typing [Muehlboeck and Tate 2017; Wrigstad et al. 2010], and nominal class-based object-oriented systems are also a popular organizing principle in industrial programming languages. Yet both nominality and static type-checking are often seen as too onerous, in particular for the early stages of developing a program. Thus, the idea of allowing programmers to start in an untyped, structural setting and later support transforming it into a typed, nominal setting is not a new one, dating back to at least Anderson and Drossopoulou [2003]. The present paper represents the most advanced version of this idea in terms of the supported language features and overall semantic guarantees, building on years of work on gradual typing [Garcia et al. 2016; Siek and Taha 2007; Siek et al. 2015a; Tobin-Hochstadt and Felleisen 2006], while also recognizing the implementation challenges of (sound) gradual typing [Bauman et al. 2017; Greenman and Felleisen 2018; Kuhlenschmidt et al. 2019; Muehlboeck and Tate 2017; Richards et al. 2017; Roberts et al. 2019; Takikawa et al. 2016]. We presented preliminary evidence that our design can be implemented efficiently, and we showed that gradual-typing research techniques can be adapted to transition not only types but even paradigms while still providing strong guarantees and good performance.

DATA AVAILABILITY STATEMENT

Our artifact provides the source code for MonNom, the source code of each configuration used to evaluate each language, the data for each plot, and a virtual machine with everything compiled and scripts set up to reproduce all measurements [Muehlboeck and Tate 2021a].

ACKNOWLEDGMENTS

We thank the reviewers for their valuable suggestions towards improving the paper. We also thank Mae Milano and Adrian Sampson, as well as the members of the Programming Languages Discussion Group at Cornell University and of the Programming Research Laboratory at Northeastern University, for their helpful feedback on preliminary findings of this work.

This material is based upon work supported in part by the National Science Foundation (NSF) through grant CCF-1350182 and the Austrian Science Fund (FWF) through grant Z211-N23 (Wittgenstein Award). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF or the FWF.

REFERENCES

- Bowen Alpern, Anthony Cocchi, Stephen Fink, and David Grove. 2001. Efficient Implementation of Java Interfaces: Invokeinterface Considered Harmless. In *OOPSLA*. ACM, New York, NY, USA, 108–124. <https://doi.org/10.1145/504282.504291>
- Christopher Anderson and Sophia Drossopoulou. 2003. BabyJ: From Object Based to Class Based Programming via Types. *Electronic Notes in Theoretical Computer Science* 82, 8 (2003), 53–81. [https://doi.org/10.1016/S1571-0661\(04\)80802-8](https://doi.org/10.1016/S1571-0661(04)80802-8)
- WOOD.
- Spenser Bauman, Carl Friedrich Bolz-Tereick, Jeremy Siek, and Sam Tobin-Hochstadt. 2017. Sound Gradual Typing: Only Mostly Dead. *PACMPL* 1, OOPSLA, Article 54 (2017), 24 pages. <https://doi.org/10.1145/3133878>
- Gavin Bierman, Erik Meijer, and Mads Torgersen. 2010. Adding Dynamic Types to C#. In *ECOOP*. Springer Berlin Heidelberg, Berlin, Heidelberg, 76–100. https://doi.org/10.1007/978-3-642-14107-2_5
- John Peter Campora, Sheng Chen, Martin Erwig, and Eric Walkingshaw. 2017. Migrating Gradual Types. *PACMPL* 2, POPL, Article 15 (2017), 29 pages. <https://doi.org/10.1145/3158103>
- Craig Chambers, David Ungar, and Elgin Lee. 1989. An Efficient Implementation of SELF, a Dynamically-Typed Object-Oriented Language Based on Prototypes. In *OOPSLA*. Association for Computing Machinery, New York, NY, USA, 49–70. <https://doi.org/10.1145/74877.74884>
- Daniel Feltey, Ben Greenman, Christophe Scholliers, Robert Bruce Findler, and Vincent St-Amour. 2018. Collapsible Contracts: Fixing a Pathology of Gradual Typing. *PACMPL* 2, OOPSLA, Article 133 (2018), 27 pages. <https://doi.org/10.1145/3276503>
- Ronald Garcia, Alison M. Clark, and Éric Tanter. 2016. Abstracting Gradual Typing. In *POPL*. ACM, New York, NY, USA, 429–442. <https://doi.org/10.1145/2837614.2837670>
- Ben Greenman and Matthias Felleisen. 2018. A Spectrum of Type Soundness and Performance. *PACMPL* 2, ICFP, Article 71 (2018), 32 pages. <https://doi.org/10.1145/3236766>
- Jessica Gronski, Kenneth Knowles, Aaron Tomb, Stephen N. Freund, and Cormac Flanagan. 2006. Sage: Hybrid Checking for Flexible Specifications. *Scheme and Functional Programming Workshop* 6 (2006), 93–104. <http://scheme2006.cs.uchicago.edu/06-freund.pdf>
- David Herman, Aaron Tomb, and Cormac Flanagan. 2010. Space-Efficient Gradual Typing. *Higher Order Symbol. Comput.* 23, 2 (2010), 167–189. <https://doi.org/10.1007/s10990-011-9066-z>
- Andre Kuhlenschmidt, Deyaaeldean Almahallawi, and Jeremy G Siek. 2018. An Efficient Compiler for the Gradually Typed Lambda Calculus. *Scheme and Functional Programming Workshop* 18 (2018), 19 pages. http://www.schemeworkshop.org/2018/Kuhlenschmidt_Almahallawi_Siek.pdf
- Andre Kuhlenschmidt, Deyaaeldean Almahallawi, and Jeremy G. Siek. 2019. Toward Efficient Gradual Typing for Structural Types via Coercions. In *PLDI*. ACM, New York, NY, USA, 517–532. <https://doi.org/10.1145/3314221.3314627>
- Barbara H. Liskov and Jeannette M. Wing. 1994. A Behavioral Notion of Subtyping. *ACM Trans. Program. Lang. Syst.* 16, 6 (Nov. 1994), 1811–1841. <https://doi.org/10.1145/197320.197383>
- Jacob Matthews and Robert Bruce Findler. 2007. Operational Semantics for Multi-Language Programs. In *POPL*. Association for Computing Machinery, New York, NY, USA, 3–10. <https://doi.org/10.1145/1190216.1190220>
- Cameron Moy, Phúc C. Nguyễn, Sam Tobin-Hochstadt, and David Van Horn. 2021. Corpe Reviver: Sound and Efficient Gradual Typing via Contract Verification. *PACMPL* 5, POPL, Article 53 (2021), 28 pages. <https://doi.org/10.1145/3434334>
- Fabian Muehlboeck and Ross Tate. 2017. Sound Gradual Typing is Nominally Alive and Well. *PACMPL* 1, OOPSLA, Article 56 (2017), 30 pages. <https://doi.org/10.1145/3133880>
- Fabian Muehlboeck and Ross Tate. 2021a. Transitioning from Structural to Nominal Code with Efficient Gradual Typing: Artifact. <https://doi.org/10.5281/zenodo.5518181>
- Fabian Muehlboeck and Ross Tate. 2021b. Transitioning from Structural to Nominal Code with Efficient Gradual Typing: Technical Report. <https://www.cs.cornell.edu/~ross/publications/monnom/>
- Max S. New, Daniel R. Licata, and Amal Ahmed. 2019. Gradual Type Theory. *PACMPL* 3, POPL, Article 15 (2019), 31 pages. <https://doi.org/10.1145/3290328>
- Gregor Richards, Ellen Artica, and Alexi Turcotte. 2017. The VM Already Knew That: Leveraging Compile-Time Knowledge to Optimize Gradual Typing. *PACMPL* 1, OOPSLA, Article 55 (2017), 27 pages. <https://doi.org/10.1145/3133879>
- Richard Roberts, Stefan Marr, Michael Homer, and James Noble. 2019. Transient Typechecks Are (Almost) Free. In *ECOOP*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, Article 5, 28 pages. <https://doi.org/10.4230/LIPIcs.ECOOP.2019.5>
- Jeremy Siek and Walid Taha. 2007. Gradual Typing for Objects. In *ECOOP*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2–27. https://doi.org/10.1007/978-3-540-73589-2_2
- Jeremy G Siek and Walid Taha. 2006. Gradual Typing for Functional Languages. *Scheme and Functional Programming Workshop* 6 (2006), 81–92. <http://scheme2006.cs.uchicago.edu/13-siek.pdf>
- Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, and John Tang Boyland. 2015a. Refined Criteria for Gradual Typing. In *SNAPL*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 274–293. <https://doi.org/10.4230/LIPIcs>

SNAPL.2015.274

- Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, Sam Tobin-Hochstadt, and Ronald Garcia. 2015b. Monotonic References for Efficient Gradual Typing. In *ESOP*. Springer Berlin Heidelberg, Berlin, Heidelberg, 432–456. https://doi.org/10.1007/978-3-662-46669-8_18
- Asumu Takikawa, Daniel Feltey, Ben Greenman, Max S. New, Jan Vitek, and Matthias Felleisen. 2016. Is Sound Gradual Typing Dead?. In *POPL*. ACM, New York, NY, USA, 456–468. <https://doi.org/10.1145/2837614.2837630>
- Sam Tobin-Hochstadt and Matthias Felleisen. 2006. Interlanguage Migration: From Scripts to Programs. In *OOPSLA*. ACM, New York, NY, USA, 964–974. <https://doi.org/10.1145/1176617.1176755>
- Sam Tobin-Hochstadt, Matthias Felleisen, Robert Findler, Matthew Flatt, Ben Greenman, Andrew M. Kent, Vincent St-Amour, T. Stephen Strickland, and Asumu Takikawa. 2017. Migratory Typing: Ten Years Later. In *SNAPL*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, Article 17, 17 pages. <https://doi.org/10.4230/LIPICs.SNAPL.2017.17>
- Michael M. Vitousek, Andrew M. Kent, Jeremy G. Siek, and Jim Baker. 2014. Design and Evaluation of Gradual Typing for Python. In *DLS*. ACM, New York, NY, USA, 45–56. <https://doi.org/10.1145/2661088.2661101>
- Michael M. Vitousek, Cameron Swords, and Jeremy G. Siek. 2017. Big Types in Little Runtime: Open-World Soundness and Collaborative Blame for Gradual Type Systems. In *POPL*. ACM, New York, NY, USA, 762–774. <https://doi.org/10.1145/3009837.3009849>
- Tobias Wrigstad, Francesco Zappa Nardelli, Sylvain Lebesne, Johan Östlund, and Jan Vitek. 2010. Integrating Typed and Untyped Code in a Scripting Language. In *POPL*. ACM, New York, NY, USA, 377–388. <https://doi.org/10.1145/1706299.1706343>