Solver-Based Gradual Type Migration

LUNA PHIPPS-COSTIN, University of Massachusetts Amherst, United States
CAROLYN JANE ANDERSON, Wellesley College, United States
MICHAEL GREENBERG, Stevens Institute of Technology, United States
ARJUN GUHA, Northeastern University, United States

Gradually typed languages allow programmers to mix statically and dynamically typed code, enabling them to incrementally reap the benefits of static typing as they add type annotations to their code. However, this type migration process is typically a manual effort with limited tool support. This paper examines the problem of *automated type migration*: given a dynamic program, infer additional or improved type annotations.

Existing type migration algorithms prioritize different goals, such as maximizing type precision, maintaining compatibility with unmigrated code, and preserving the semantics of the original program. We argue that the type migration problem involves fundamental compromises: optimizing for a single goal often comes at the expense of others. Ideally, a type migration tool would flexibly accommodate a range of user priorities.

We present TypeWhich, a new approach to automated type migration for the gradually-typed lambda calculus with some extensions. Unlike prior work, which relies on custom solvers, TypeWhich produces constraints for an off-the-shelf MaxSMT solver. This allows us to easily express objectives, such as minimizing the number of necessary syntactic coercions, and constraining the type of the migration to be compatible with unmigrated code.

We present the first comprehensive evaluation of GTLC type migration algorithms, and compare TypeWhich to four other tools from the literature. Our evaluation uses prior benchmarks, and a new set of "challenge problems." Moreover, we design a new evaluation methodology that highlights the subtleties of gradual type migration. In addition, we apply TypeWhich to a suite of benchmarks for Grift, a programming language based on the GTLC. TypeWhich is able to reconstruct all human-written annotations on all but one program.

CCS Concepts: • **Software and its engineering → Multiparadigm languages**.

Additional Key Words and Phrases: gradual typing, type inference

ACM Reference Format:
Luna Phipps-Costin, Carolyn Jane Anderson, Michael Greenberg, and Arjun Guha. 2021. Solver-Based Gradual Type Migration. *Proc. ACM Program. Lang.* 5, OOPSLA, Article 111 (October 2021), 27 pages. https://doi.org/10.1145/3485488

1 INTRODUCTION

Gradually typed languages allow programmers to freely mix statically and dynamically typed code. This enables users to add static types gradually, providing the benefits of static typing without requiring the entirety of a codebase to be overhauled at once [Siek and Taha 2006; Tobin-Hochstadt and Felleisen 2006]. Over the past decade, gradually typed dialects of several mainstream languages, such as JavaScript, Python, and Ruby, have become established in industry. However, the process of *migrating* an untyped program to use gradual types has largely remained a labor-intensive manual

Authors' addresses: Luna Phipps-Costin, Amherst, University of Massachusetts Amherst, United States; Carolyn Jane Anderson, Wellesley, Wellesley College, United States; Michael Greenberg, Hoboken, Stevens Institute of Technology, United States; Arjun Guha, Boston, Northeastern University, United States.

This work is licensed under a Creative Commons Attribution 4.0 International License.
© 2021 Copyright held by the owner/author(s).
2475-1421/2021/10-ART111
https://doi.org/10.1145/3485488

effort. Just as type inference facilitates static typing, type migration tools have the potential to make gradual typing easier to use.

However, automating type migration is a challenging problem. Even if we consider a small language, such as the *gradually typed lambda calculus* (GTLC) [Siek and Taha 2006], and limit ourselves to modifying existing type annotations, a single program may have many possible migrations. Existing approaches either produce a single migration [Flanagan et al. 1996; Henglein and Rehof 1995; Rastogi et al. 2012; Siek and Vachharajani 2008; Wright and Cartwright 1997], or a menu of possible migrations without guidance on which to select [Campora et al. 2018b; Migeed and Palsberg 2020]. How should we choose among the migrations produced by various approaches?

In this paper, we present a new approach to gradual type migration for the GTLC, and the first comprehensive evaluation of prior work in this area. We first illustrate the tension between *precise type migrations* that produce informative type annotations, *safe type migrations* that do not introduce new dynamic errors, and *compatible type migrations* that preserve interoperability with other, unmigrated code. We show that prior approaches navigate these tradeoffs in different ways. Some favor making types as precise as possible, even if the increased precision compromises safety or compatibility. Others favor compatibility and safety at the expense of precision. Finally, many approaches statically reject programs that may have dynamic type errors.

We present TypeWhich, a new approach to type migration that navigates these tradeoffs as follows. (1) TypeWhich does not statically reject any closed programs, including programs that obviously crash with a dynamic error. (2) TypeWhich always produces safe type migrations, thus does not introduce new dynamic errors. (3) TypeWhich is the first approach that allows the user to chose between precise migrations and compatible migrations, and we show that there are different situations in which the user may prefer one or the other.

Under the hood, TypeWhich differs from prior work in two key ways. First, whereas prior work relies on custom constraint solvers, TypeWhich generates constraints for an off-the-shelf MaxSMT solver [Bjørner et al. 2015]. This makes it easy to add constraints and language features, as we demonstrate by extending the GTLC in several ways and supporting the Grift gradually typed language [Kuhlenschmidt et al. 2019].

Second, using a general-purpose solver is what allows TypeWhich to readily support multiple kinds of migrations. We accomplish this by using the MaxSMT solver in a two-stage process. We first formulate a MaxSMT problem with an objective function that synthesizes precise types. The reconstructed type of the program may not be compatible with all contexts, but it reveals the (potentially higher-order) interface of the program. We then formulate new constraints on the type of the program to enforce compatibility, and use the MaxSMT solver a second time to produce a new solution.

Our evaluation compares TypeWhich to four other type migration approaches using a two-part evaluation suite: a set of existing benchmarks by Migeed and Palsberg [2020], and a new set of "challenge problems" that we devise. We also design an evaluation methodology that reflects the subtleties of type migration. Although different approaches to type migration prioritize different goals, TypeWhich performs as well or better than existing tools on all prior benchmark suites. We design our "challenge problems" to highlight the strengths and weaknesses of all approaches, including TypeWhich. Finally, we apply TypeWhich to a suite of Grift programs from Kuhlenschmidt et al. [2019], and find that it reproduces all hand-written type annotations, with one exception.

Limitations. TypeWhich focuses on type migration for the core GTLC, which merely extends simple types with an unknown ($\star$). This allows us to directly compare TypeWhich to other GTLC type migration algorithms (§6). But, it does limit the scope of our work and our conclusions:

(1) This paper does *not* consider union types, polymorphism, subtyping, recursive types, and other features that are necessary to build more complete programming languages.
(2) Our evaluation on the GTLC uses small, artificial programs. These benchmarks illustrate tricky cases where different approaches produce different results, but they do not represent the structure of real-world programs. TypeWhich has a frontend for the Grift programming language, which we use to evaluate on the benchmarks presented by Kuhlenschmidt et al. [2019], but these are also small.
(3) This paper focuses on type migration for the guarded semantics of gradual types. Many gradually typed languages, including TypeScript, use alternative semantics, which we briefly discuss in §7.

Contributions. Our key contributions are as follows:

(1) We illustrate the tension between the many goals of type migration, and argue that type migration tools should give users the ability to choose between different goals (§2 and §3).
(2) We present the TypeWhich approach to type migration, which formulates constraints for an off-the-shelf MaxSMT solver (§4). TypeWhich supports the GTLC and additional language features required to support the Grift gradually typed language (§5).
(3) We present a new set of type migration "challenge problems" that illustrate the strengths and weaknesses of different approaches to type migration (§6).
(4) We present a comprehensive comparison of five approaches to type migration (including ours), using a new evaluation methodology. For this comparison, we implement a unified framework for running, evaluating, and validating type migration algorithms.
(5) Finally, we contribute re-implementations of the type migration algorithms from Campora et al. [2018b] and Rastogi et al. [2012]. Ours is the first publicly available implementation of Rastogi et al. [2012].

Artifact. The artifact for this paper is available at https://doi.org/10.5281/zenodo.5141479.

2 WHAT MATTERS FOR TYPE MIGRATION?

When designing a type migration tool, we must consider several important questions:

(1) A key goal of type migration is to improve the precision of type annotations. However, there are often multiple ways to improve type precision [Migeed and Palsberg 2020] that induce different run-time checks. For any given type migration system, we must therefore ask the question, *Can a user choose between several alternative migrations?*
(2) When the migrated code is only a fragment of a larger codebase, increasing type precision can introduce type errors at the boundaries between migrated and unmigrated code [Rastogi et al. 2012]. Thus we must ask, *Does the migrated code remain compatible with other, unmigrated code?*
(3) A type migration tool may also uncover potential run-time errors. However, these errors may be unreachable, or only occur in certain configurations or on certain platforms. Thus we must ask, *Should a migration turn (potential) run-time errors into static type errors?*
(4) Finally, safe gradually typed languages introduce checks that enforce type safety at run-time. Making a type more precise can alter these checks, affecting run-time behavior. Thus we must ask, *Does the migrated program preserve the behavior of the original program?*

This section explores these questions with examples from the gradually-typed lambda calculus (GTLC) with some modest extensions. We write programs in an OCaml-like syntax with explicit type annotations. The type $\star$ is the *unknown type* (also known as the dynamic type or the any type),

which is compatible with all types. Under the hood, converting to and from the ★ type introduces coercions [Henglein 1994]; these coercions can fail at run-time with a dynamic type error.

Type migration can introduce new static errors. Figure 1 shows a function that uses its ★-typed argument first as a number and then as a function. Since ★ is compatible with all types, the function is well-typed, but guaranteed to produce a dynamic type error when applied. In this case, it seems harmless for a type migration tool to turn this dynamic type error into a static type error.

```
1 let A (x : ★) =
2   let _ = x + 10 in
3   x ()
```

Fig. 1. Reachable error.

However, it is also possible for the crashing expression to be unreachable. Figure 2 wraps the same dynamic error in the unused branch of a conditional. In this case, improving the type annotation would lead to a spurious error: the migrated program would fail even though the original ran without error. Although this example is contrived, programs in untyped languages often have code whose reachability is environment-dependent (e.g., JavaScript web programs that support multiple browsers, Python programs that can be run in Python 2 and 3). The flexibility of gradual typing is particularly valuable in these cases, but reasoning about safety and precision in tandem is subtle.

```
1 let B (x : ★) =
2   if false then
3     let _ = x + 10 in
4     x ()
5   else
6     x ()
```

Fig. 2. Unreachable error.

Type migration can restrict the context of a program. There are many cases where it is impractical to migrate an entire program at once. For example, the programmer may not be able to modify the source code of a library; they may be migrating a library that is used by others; or it may just be unacceptable to change every file in a large software project. In these situations, the type migration question is even trickier.

Figure 3 shows a higher-order function `C` that calculates $1 + f(x)$ when x is greater than zero. We could migrate `C` to require f to be an integer function, which precisely captures how `C` uses f. However, this migration makes some calls to `C` ill-typed. For example, `C 0 0` evaluates to `42` before migration, but is ill-typed after migration.

```
1  int -> int        int
2
3
   let C (f : ★) (x : ★) =
4    if x > 0 then
5      1 + f x
6    else
7      42
```

Fig. 3. Context restriction.

Figure 4 illustrates another subtle interaction between type-migrated code and its context. The function `D` receives f and expects it to be a function over numbers. Unlike the previous example, `D` always calls f, so it may appear safe to annotate f with the type `int -> int`. However, `D` also returns f back to its caller, so this migration changes the return type of `D` from ★ to `int -> int`. For example, when f is the identity function, $D(f)$ returns the identity function before migration, but after migration $D(f)$ is restricted to only work on `int`s.

```
1 int -> int
2
3 let D (f : ★) =
4   f 100 + 10;
5   f              int -> int
6
7 let id : ★ = D(fun (x: ★) . x)
```

Fig. 4. Context restriction.

To summarize, there is a fundamental trade-off between making types precise in migrated code, and maintaining compatibility with unmigrated code.

Type migration can introduce new dynamic errors. So far, we have looked at migrations that introduce static type errors. However, there is a more insidious problem that can occur: a migration can introduce new dynamic type errors. Figure 5 shows a program that runs without error: `E` receives the identity function and applies it to two different types. However, since `E`'s argument has type ★, which is compatible with all types,

```
1 let E(id : ★) =
2   id 2;
3   id true          int
4
5 E(fun (x : ★) . x);
```

Fig. 5. Dynamic type error.

Base types

$B \coloneqq \mathsf{int} \mid \mathsf{bool}$

Types and contexts

$S, T \coloneqq B$	Base type
$\mid S \to T$	Function type
$\mid \star$	Unknown type
$\Gamma \coloneqq \cdot \mid \Gamma, x : T$	

Constants

$b \coloneqq \mathsf{true} \mid \mathsf{false}$	Boolean literal
$n \coloneqq \cdots$	Integer literal
$c \coloneqq b \mid n$	

Expressions

$e \coloneqq x$	Identifier
$\mid c$	Literal
$\mid \mathsf{fun}(x : T).e$	Function
$\mid e_1\ e_2$	Application
$\mid e_1 \times e_2$	Multiplication

Type Consistency $\boxed{T \sim T}$

$$\frac{}{\star \sim T} \quad \frac{}{T \sim \star} \quad \frac{}{T \sim T} \quad \frac{}{B \sim B} \quad \frac{S_1 \sim S_2 \qquad T_1 \sim T_2}{S_1 \to T_1 \sim S_2 \to T_2}$$

Typing Literals $\boxed{ty : c \to B}$

$ty(n) = \mathsf{int} \quad ty(b) = \mathsf{bool}$

Typing $\boxed{\Gamma \vdash e : T}$

$$\frac{\Gamma(x) = T}{\Gamma \vdash x : T} \quad \frac{}{\Gamma \vdash c : ty(c)} \quad \frac{\Gamma, x : S \vdash e : T}{\Gamma \vdash \mathsf{fun}(x : S).e : S \to T}$$

$$\frac{\Gamma \vdash e_1 : \star \quad \Gamma \vdash e_2 : T}{\Gamma \vdash e_1\ e_2 : \star} \quad \frac{\Gamma \vdash e_1 : S \to T \quad \Gamma \vdash e_2 : S' \quad S \sim S'}{\Gamma \vdash e_1\ e_2 : T}$$

$$\frac{\Gamma \vdash e_1 : S \quad \Gamma \vdash e_2 : T \quad S \sim \mathsf{int} \quad T \sim \mathsf{int}}{\Gamma \vdash e_1 \times e_2 : \mathsf{int}}$$

Fig. 6. The Gradually Typed Lambda Calculus (GTLC): surface syntax and typing.

the program is well-typed even if we migrate the identity function to require an integer argument. Gradual typing will wrap the function to dynamically check that it only receives integers. So the program runs without error before migration, but produces a dynamic type error after migration. Strictly speaking, although this migration introduces a new dynamic error, its static types are more precise. When evaluating migrations, it is not enough to consider just the types or interfaces: it is important to understand which run-time checks will be inserted.

In summary, there are several competing concerns that we must consider when choosing an approach to type migration. TypeWhich prioritizes preserving the behavior of the original program: it produces types that do not introduce new static or dynamic errors in the migrated code. However, this objective leaves the question of context unanswered. Should TypeWhich produce the most precise type it can? This may make the migrated code incompatible with unmigrated code. So, should TypeWhich instead produce a type that is compatible with all untyped code? This would mean discarding a lot of useful information, e.g., the types of function arguments. Or, should TypeWhich strike a compromise between precision and compatibility? We think the right answer depends on the context in which the type migration tool is being used. Instead of making an arbitrary decision, TypeWhich allows the programmer to choose between several migrations that prioritize different properties.

3 FORMALIZING THE TYPE MIGRATION PROBLEM

We now formally define the type migration problem. We first briefly review the *gradually typed lambda calculus* (GTLC) [Siek et al. 2015b], which is a core calculus for mixing typed and untyped code. We then present several definitions of type migration for the GTLC.

3.1 The Gradually Typed Lambda Calculus

The Gradually Typed Lambda Calculus (GTLC) extends the typed lambda calculus with base types (integers and booleans) and the *unknown type* $\star$. Figure 6 shows its syntax and typing rules.

Type checking relies on the *type consistency* relation, $S \sim T$. Type consistency determines whether an S-typed expression may appear in a T-typed context. Two types are consistent if they are structurally equal up to any unknown ($\star$) types within them; the $\star$-type is consistent with all types and any expression may appear in a $\star$-typed context. The type consistency relation is

Ground types

$G ::= B \mid \text{fun}$

Coercions

$k ::= G?$	Untag
$\mid G!$	Tag
$\mid \text{wrap}(k_1, k_2)$	Wrap function
$\mid k_1; k_2$	Sequence
$\mid \text{id}_T$	Identity

Expressions

$e ::= \cdots \mid [k]\ e$	Apply coercion

Untagged values

$u ::= c \mid \text{fun}(x : T).e$

Values

$v ::= u \mid \text{box}(G, u)$

Evaluation Contexts

$E ::= [] \mid E\ e \mid v\ E \mid [k]\ E$

Active Expressions

$ae ::= (\text{fun}(x : T).e)\ v \mid [k]\ v$

Evaluation $\boxed{\vdash e \hookrightarrow e}$

$$(\text{fun}(x : T).e)\ v \hookrightarrow e[x/v]$$
$$[\text{id}]\ v \hookrightarrow v$$
$$[G!]\ (u) \hookrightarrow \text{box}(G, u)$$
$$[G?]\ (\text{box}(G, u)) \hookrightarrow u$$
$$[\text{wrap}(k_1, k_2)]\ v \hookrightarrow \text{fun}(x : \star).[k_2]\ (v\ ([k_1]\ x))$$
$$[k_1; k_2]\ v \hookrightarrow [k_2]\ ([k_1]\ v)$$
$$\frac{ae \hookrightarrow e'}{E[ae] \hookrightarrow E[e']}$$

$$\text{coerce}(T, T) = \text{id}_T$$
$$\text{coerce}(\star, B) = B?$$
$$\text{coerce}(B, \star) = B!$$
$$\text{coerce}(\star, \star \to \star) = \text{fun?}$$
$$\text{coerce}(\star \to \star, \star) = \text{fun!}$$
$$\text{coerce}(S_1 \to S_2, T_1 \to T_2) = \text{wrap}(\text{coerce}(T_1, S_1), \text{coerce}(S_2, T_2))$$
$$\text{coerce}(\star, T_1 \to T_2) = \text{fun?}; \text{wrap}(\text{coerce}(T_1, \star), \text{coerce}(\star, T_2))$$
$$\text{coerce}(T_1 \to T_2, \star) = \text{wrap}(\text{coerce}(\star, T_1), \text{coerce}(T_2, \star)); \text{fun!}$$
$$\text{coerce}(S, T) = \text{coerce}(S, \star); \text{coerce}(\star, T)$$

Coercion Insertion $\boxed{\Gamma \vdash e \Rightarrow e, T}$

$$\frac{\Gamma(x) = T}{\Gamma \vdash x \Rightarrow x, T} \qquad \frac{}{\Gamma \vdash c \Rightarrow c, ty(c)}$$

$$\frac{\Gamma, x : S \vdash e \Rightarrow e', T}{\Gamma \vdash \text{fun}(x : S).e \Rightarrow \text{fun}(x : S).e', S \to T}$$

$$\frac{\Gamma \vdash e_1 \Rightarrow e_1', S \to T \quad \Gamma \vdash e_2 \Rightarrow e_2', S'}{\Gamma \vdash e_1\ e_2 \Rightarrow e_1'\ ([\text{coerce}(S', S)]\ e_2'), T}$$

$$\frac{\Gamma \vdash e_1 \Rightarrow e_1', T \quad T \neq T_1 \to T_2 \quad \Gamma \vdash e_2 \Rightarrow e_2', S}{\Gamma \vdash e_1\ e_2 \Rightarrow ([\text{coerce}(T, \star \to \star)]\ e_1')\ ([\text{coerce}(S, \star)]\ e_2'), \star}$$

Fig. 7. Coercion insertion and evaluation for the GTLC.

reflexive and symmetric, but not transitive: int and bool are both consistent with $\star$ but not with each other.

The typing rules for identifiers, literals, and functions are straightforward, but there are two function application rules: (1) If the expression in function position has type $\star$, then the argument may have any type, and the result of the application has type $\star$. (2) When the type of the function expression is an arrow type ($S \to T$), the result has type T. The type of the argument must be consistent with—but not necessarily equal to—the type of argument the function expects ($S' \sim S$).

We add a built-in multiplication operator that requires the types of its operands to be consistent with int (i.e., an operand may have type $\star$). We choose multiplication because the "+" operator is overloaded in many untyped languages: we add addition in §5, where we discuss overloading.

3.2 Ground Types and Coercion-Based Semantics

Programs in the GTLC are not run directly, but are first compiled to an intermediate representation where static type consistency checks are turned into dynamic checks if necessary. There are two well-known mechanisms for describing these dynamic checks: casts and coercions. We use coercions, following Henglein [1994], as they most closely match the type-tagging and tag-checking operations used at run-time in dynamic languages.[1]

The *ground types* (G in Figure 7) are the types that are dynamically observable, and include all base types and a ground type fun for all functions. The two basic coercions (k) tag a value with a

[1]The two approaches are inter-translatable [Greenberg 2013; Herman et al. 2011] with full abstraction [Siek et al. 2015a].

Type Precision $\boxed{T \sqsubseteq T}$ **Expression Precision** $\boxed{e \sqsubseteq e}$

$$\frac{}{\star \sqsubseteq T} \quad \frac{}{T \sqsubseteq T} \quad \frac{S_1 \sqsubseteq S_2 \qquad T_1 \sqsubseteq T_2}{S_1 \to T_1 \sqsubseteq S_2 \to T_2} \quad \frac{}{x \sqsubseteq x} \quad \frac{}{c \sqsubseteq c} \quad \frac{e_1 \sqsubseteq e_1' \qquad e_2 \sqsubseteq e_2'}{e_1\ e_2 \sqsubseteq e_1'\ e_2'} \quad \frac{T \sqsubseteq T' \qquad e \sqsubseteq e'}{\mathsf{fun}(x : T).e \sqsubseteq \mathsf{fun}(x : T').e'}$$

Fig. 8. Type and expression precision.

ground type ($G!$) and untag a value after checking that it has a particular ground type ($G?$). Both of these operations can fail: an already-tagged value cannot be re-tagged, and untagging succeeds only if the value has the expected ground type. There are three additional coercions: identity coercions, which exist only to simplify certain definitions; a sequencing coercion ($k_1; k_2$); and a *function proxy* wrap that lifts coercions to functions.

To see how the coercion system works, consider a case where we have a $\star$-typed value f that we want to treat as a function of type $\mathsf{int} \to \mathsf{int}$. To do so, we apply f to a coercion as follows:

$$[\mathsf{fun?}; \mathsf{wrap}(\mathsf{int!}, \mathsf{int?})]f$$

The sequence evaluates from left to right: it first checks that f is a function (fun?), and then wraps f in a function proxy that will tag its int argument (since f expects a $\star$ value) and will untag its result (since f returns a $\star$, but we expect an int).

The values of the language (v) include constants, functions, and values tagged with a ground type. We define tagged values ($\mathsf{box}(G, u)$) so that a tag can only be placed on an untagged value (u).

The coercion insertion rules are analogous to typing, but produce both a type and an equivalent expression with explicit coercions. They rely on the *coerce* metafunction that translates a static consistency check $S \sim T$ into a corresponding coercion that is dynamically checkable. When two types are identical, *coerce* produces the identity coercion, which can be safely removed. The final case of *coerce* addresses inconsistencies ($S \nsim T$). Instead of rejecting programs with inconsistent checks, we produce a coercion that is *doomed to fail*. Gradual typing systems sometimes reject programs that demand casts between incompatible types. However, doing so violates the desired property that migrations should preserve the behavior of the original program when possible. If we rejected these programs, a user would need to excise all incompatibilities, whether or not they are in live code branches, at the onset of migration.

3.3 Type Migration

All formulations of the type migration problem rely on defining *type precision*, where $\star$ is the least precise type. The type precision relation (Figure 8), written $S \sqsubseteq T$, is a partial order that holds when S is less precise than T (or S and T are identical). We use type precision to define expression precision in the obvious way: an expression is more precise than its structural equivalent if its type annotations are more precise according to the type precision relation.

Migeed and Palsberg [2020] define a type migration as an expression that has more precise type annotations, and use this definition to study the decidability and computational complexity of several problems, such as finding migrations that cannot be made more precise.

Definition 3.1 (Type Migration). Given $\vdash e : T$ and $\vdash e' : T'$, e' is a *type migration* of e if $e \sqsubseteq e'$ and $T \sqsubseteq T'$.

However, as we argued in §2, improving type precision is one of several competing goals for type migration. Another important goal is to avoid introducing new errors into the program. To reason about this, we must reformulate the definition of a type migration to relate the values produced by the original expression and its migration. We propose the following definition:

Definition 3.2 (Safe Type Migration). Given $\vdash e : T$ and $\vdash e' : T'$, e' is a *safe type migration* of e if:

(1) $e \sqsubseteq e'$;
(2) $T \sqsubseteq T'$; and
(3) $e \hookrightarrow^* v$ if and only if $e' \hookrightarrow^* v'$ with $v \sqsubseteq v'$.

This definition of type migration relates the values of the two expressions. However, it is too weak. For one thing, it does not say anything about programs that produce errors or do not terminate. But there is a more serious problem: it is too permissive for function types. For example, given the identity function with type $\star \to \star$, this definition allows a type migration that changes its type to $\mathsf{int} \to \mathsf{int}$, which will produce a dynamic type error if the function is applied to non-integers.

To address this issue, the definition of type migration must take into account the contexts in which the migrated expression may be used. We define a well-typed program context as a context with a hole that can be filled with a well-typed open expression to get a well-typed closed expression.

Definition 3.3 (Well-Typed Program Context). A program context C is well typed, written $C : (\Gamma \vdash S) \Rightarrow T$ if for all expressions e where $\Gamma \vdash e : S$ we have $\vdash C[e] : T$.

We now define a *context-restricted type migration* as a more precisely-typed expression that is equivalent to the original expression in all contexts that can be filled with an expression of a given type S. Note that the type expected by the context (S) must be consistent (but not identical) with the types of both the original and the migrated expression.

Definition 3.4 (Context-restricted Type Migration). Given $\vdash e : T$, $\vdash e' : T'$, and a type S where $S \sim T$ and $S \sim T'$, e' is a *context-restricted type migration* of e at type S if:

(1) $e \sqsubseteq e'$;
(2) $T \sqsubseteq T'$; and
(3) For all $C : (\cdot \vdash S) \Rightarrow U$, either a) $C[e] \hookrightarrow^* v$ and $C[e'] \hookrightarrow^* v'$ with $v \sqsubseteq v'$; b) both $C[e]$ and $C[e']$ get stuck at a failed coercion;[2] or c) both $C[e]$ and $C[e']$ do not terminate.

We call a context-restricted type migration at type $\star$ a *compatible type migration.*

At the limit, the context's expected type S could be $\star$, in which case the definition is essentially equivalent to that of Rastogi et al. [2012, Theorem 3.22]. However, this is a very strong requirement that rules out many informative migrations (§2). If the programmer is comfortable making assumptions about how the rest of the program will interact with the migrated expression, they may choose a more precise S, and allow a wider range of valid type migrations.

We present these definitions to describe the type migration problem that we seek to address in TypeWhich. However, we do not prove that TypeWhich produces a context-restricted type migration. Instead, this paper presents empirical results to show the effectiveness of TypeWhich on the GTLC, when compared to other type migration tools.

4 THE TYPEWHICH APPROACH TO TYPE MIGRATION

We now present TypeWhich, an approach to type migration that differs in two ways from previous work. (1) Instead of relying on a custom constraint solver, TypeWhich produces constraints and an objective function for the Z3 MaxSMT solver [Bjørner et al. 2015]. (2) Instead of producing a single migration, or several migrations without guidance on which to choose, TypeWhich allows the user to choose between migrations that prioritize type precision or compatibility with untyped code. Moreover, the TypeWhich migration algorithm handles these different scenarios in a uniform way. This section presents TypeWhich's type migration algorithm for the core GTLC. §5

[2] This definition collapses all errors to stuck states. If the GTLC were extended with exception handling, then this definition would have to be adjusted.

Types

$T \coloneqq \cdots$
| $\alpha, \beta, \gamma, \delta$ Type metavariables

Coercions

$k \coloneqq \cdots$
| $\underline{coerce}(T_1, T_2)$ Coercion from T_1 to T_2

Type Representation

```
(declare-datatypes ()
  ((Typ (star) (int) (bool)
        (arr (in Typ) (out Typ)))))
```

Constraints

$\phi \coloneqq T_1 = T_2$ Type equality
| w Boolean variable (weight)
| $\phi_1 \wedge \phi_2$ Conjunction
| $\phi_1 \vee \phi_2$ Disjunction
| $\neg\phi$ Negation

Constraint Metafunctions

$ground \in T \to \phi$
$ground(T) = T \in B \vee T = \star \to \star$

Fig. 9. The type constraint language, and language extensions for constraint generation.

extends TypeWhich with additional language features, including some that have not been precisely described in prior work.

4.1 The Language of Type Constraints

For the purpose of constraint generation, we make two additions to the GTLC (Figure 9):

(1) We extend types with type metavariables (α).
(2) We introduce a new coercion, $\underline{coerce}(S, T)$, which represents a suspended call to the *coerce* metafunction (Figure 7). The type arguments to $\underline{coerce}$ may include type metavariables. After constraint solving, we substitute any type metavariables with concrete types and use the *coerce* metafunction to get a primitive coercion (k).

Both of these are auxiliary and do not appear in the final program.

The constraints (ϕ) that we generate are boolean-sorted formulas for a MaxSMT solver that supports the theory of algebraic datatypes [Barrett et al. 2007]. In addition to the usual propositional connectives, our constraints involve equalities between types ($T_1 = T_2$) and auxiliary boolean variables (w). We use these boolean variables to define soft constraints that guide the solver towards solutions with fewer non-trivial coercions.

Using Z3's algebraic datatypes, we define a new sort (`Typ`) that encodes all types (T) except type metavariables. Constraint generation defines a `Typ`-sorted constant for every metavariable that occurs in a type. For example, we can solve the type constraint $\alpha \to \mathsf{int} = \beta$ with the following commands to the solver:

```
(declare-const alpha Typ)
(declare-const beta Typ)
(assert (= (arr alpha int) beta))
```

This example is satisfiable, and the model assigns `alpha` and `beta` to metavariable-free types (represented as `Typ`). If σ is such a model, we write $\textsc{Subst}(\sigma, \beta)$ to mean the metavariable-free type assigned to β, i.e., the closure of substituting with the model σ. In this example, α is unconstrained, so there are several possible models: $\sigma = \{\alpha \mapsto \mathsf{int}, \beta \mapsto \alpha \to \mathsf{int}\}$ is one, as is $\sigma' = \{\alpha \mapsto \star, \dots\}$. We have $\textsc{Subst}(\sigma, \beta) = \mathsf{int} \to \mathsf{int}$, while $\textsc{Subst}(\sigma', \beta) = \star \to \mathsf{int}$.

Finally, for succinctness, we define $ground(T)$, which produces a constraint that is satisfiable when T is a ground type. At the moment, the only ground types are base types and dynamic function types ($\star \to \star$). §5 extends the language with additional types and augments the definition of *ground*.

$$\boxed{\Gamma \vdash e \Rightarrow e, T, \phi}$$

$$\textsc{Id}\ \frac{\phi = (\alpha = \Gamma(x) \wedge w) \vee (\alpha = \star \wedge \neg w) \quad \alpha, w \text{ is fresh}}{\Gamma \vdash x \Rightarrow [\underline{coerce}(\Gamma(x), \alpha)]x, \alpha, \phi} \qquad \textsc{Const}\ \frac{\phi = (\alpha = ty(c) \wedge w) \vee (\alpha = \star \wedge \neg w) \quad \alpha, w \text{ is fresh}}{\Gamma \vdash c \Rightarrow [\underline{coerce}(ty(c), \alpha)]c, \alpha, \phi}$$

$$\textsc{Fun}\ \frac{\Gamma, x : \alpha \vdash e \Rightarrow e', T, \phi_1 \quad \beta, w \text{ fresh} \quad \phi_2 = (\beta = \alpha \to T \wedge w) \vee (\beta = \star \wedge \mathit{ground}(\alpha \to T) \wedge \neg w)}{\Gamma \vdash \mathsf{fun}(x : \alpha).e \Rightarrow [\underline{coerce}(\alpha \to T, \beta)]\mathsf{fun}(x : \alpha).e', \beta, \phi_1 \wedge \phi_2}$$

$$\textsc{App}\ \frac{\begin{array}{c}\Gamma \vdash e_1 \Rightarrow e_1', T_1, \phi_1 \qquad \Gamma \vdash e_2 \Rightarrow e_2', T_2, \phi_2 \qquad \alpha, \beta, \gamma, w_1, \text{and } w_2 \text{ are fresh} \\ \phi_3 = (T_1 = \alpha \to \beta \wedge w_1) \vee (T_1 = \alpha = \beta = \star \wedge \neg w_1) \qquad \phi_4 = (T_2 = \alpha) \qquad \phi_5 = (\beta = \gamma \wedge w_2) \vee (\gamma = \star \wedge \neg w_2)\end{array}}{\Gamma \vdash e_1\ e_2 \Rightarrow [\underline{coerce}(\beta, \gamma)](([\underline{coerce}(T_1, \alpha \to \beta)]e_1')\ e_2'), \gamma, \phi_1 \wedge \phi_2 \wedge \phi_3 \wedge \phi_4 \wedge \phi_5}$$

$$\textsc{Mul}\ \frac{\begin{array}{c}\Gamma \vdash e_1 \Rightarrow e_1', T_1, \phi_1 \qquad \Gamma \vdash e_2 \Rightarrow e_2', T_2, \phi_2 \qquad w_1, w_2, \text{and } w_3 \text{ are fresh} \\ \phi_3 = (T_1 = \mathsf{int} \wedge w_1) \vee (T_1 = \star \wedge \neg w_1) \quad \phi_4 = (T_2 = \mathsf{int} \wedge w_2) \vee (T_2 = \star \wedge \neg w_2) \\ \phi_5 = (\alpha = \mathsf{int} \wedge w_3) \vee (\alpha = \star \wedge \neg w_3)\end{array}}{\Gamma \vdash e_1 \times e_2 \Rightarrow [\underline{coerce}(\mathsf{int}, \alpha)]([\underline{coerce}(T_1, \mathsf{int})]e_1' \times [\underline{coerce}(T_2, \mathsf{int})]e_2'), \alpha, \phi_1 \wedge \phi_2 \wedge \phi_3 \wedge \phi_4 \wedge \phi_5}$$

Fig. 10. Constraint generation for GTLC

4.2 Generating Type Constraints

We now present constraint generation for the GTLC. To simplify the presentation, we assume that all bound variables have type $\star$. Constraint generation is a two-step process:

(1) We replace every $\star$ annotation in the input program with a fresh metavariable. The solution to the constraints maps these metavariables to types, which may be more precise than $\star$.
(2) We generate constraints by applying deterministic, syntax-directed inference rules.

Since the first step is straightforward, we focus on constraint generation. The constraint generation rules are of the form $\Gamma \vdash e \Rightarrow e', T, \phi$: the inputs are the type environment (Γ) and the expression (e), and the outputs are as follows:

(1) An output expression (e') that is equivalent to the input expression, but with explicit coercions.
(2) A type (T), which is the type of the expression, and may include metavariables.
(3) A constraint (ϕ) with type-sorted and boolean-sorted free variables.

When formulating constraint generation, there are several requirements to keep in mind. First, the constraint ϕ may be satisfiable in several ways. We will eventually use soft constraints to choose among solutions, but we design the constraint generation process so that *all models of ϕ correspond to valid migrations*. Second, as argued in §2, we do not want to reject any programs. We therefore set up constraint generation so that we *do not introduce new static errors*. Our final goal is to *favor informative types*. We do this via soft constraints that penalize the number of non-trivial, syntactic coercions. Note that this is not the same as minimizing the number of coercions performed during evaluation, which is a harder problem (but see Campora et al. [2018a]).

Constraint Generation Rules. Constraint generation is syntax directed (Figure 10), albeit we assume we can generate fresh names. As a general principle, we allow all expressions to be coerced to $\star$: this enables us to migrate all programs, even though it may generate coercions that are doomed to fail if they are ever run. This property is critical to ensure that models exist for all programs (Theorem 4.2).[3]

[3]We have also implemented a version of TypeWhich that uses an alternative constraint generation rule for identifiers that enforces rigid types together with a modified version of the function application rule that can coerce the function argument. This leads to a loss of type precision, but produces type annotations that are more robust to code-refactoring. Both approaches are sound and safe at the generated types (§4.3).

Following this principle, the rule for identifiers (Id) introduces a coercion that is either the identity coercion (when α is T, the type of the identifier in the environment), or a coercion to $\star$ (when α is $\star$). At a later step (§4.3), we produce a soft constraint favoring w over $\neg w$, which guides the solver towards solutions that avoid the non-trivial coercions when possible.

Similarly, the rule for constants (Const) generates two new variables: α and a fresh weight variable w. The rule constrains the type α to either be the type of the constant, or the $\star$ type (i.e, to avoid rejecting true $\times$ 1). In the former case, we constrain w to be true, and in the latter, to false.

The rule for functions (Fun) assumes that the argument is annotated with a unique metavariable (α) and recurs into the function body, which produces some type T. The rule gives the function the type β (a fresh metavariable), and constrains it to be the type of the function ($\alpha \to T$) or the $\star$ type. In the latter case, we also constrain the type of the function to be the ground type ($\star \to \star$). We use a weight w to prefer the former case without rejecting expressions like $1 \times (\mathsf{fun}(x : \star).x)$.

The rule for function applications (App) produces a constraint that is a conjunction of five clauses: ϕ_1 and ϕ_2 are the constraints that arise when recurring into the two sub-expressions of the application; ϕ_3 constrains the type of the function; ϕ_4 constrains the type of the argument; and ϕ_5 constrains the type of the result. Together, ϕ_3 and ϕ_4 capture the two ways in which applications can be typed in the GTLC: the function may be of type $\star$, in which case it is coerced to the function ground type, $\star \to \star$ and w_1 is false, or the function already has a function type, and w_1 is true. In either case, the argument type is constrained to be the function input type α. The final constraint allows the result type, β, to be coerced to $\star$; w_2 is true only if this is a non-trivial coercion.

The rule for multiplication (Mul) produces a five-part conjunction: ϕ_1 and ϕ_2 are the constraints produced by its operands; ϕ_3 and ϕ_4 constrain each operand to either be int or $\star$ and use weights to prefer the former; and ϕ_5 constrains the type of the result to either be int or $\star$, with a weight that prefers for the former; again, this is necessary to avoid rejecting programs.

Example 1: Types for the Identity Function. Consider the following program, which applies the identity function to 42 and true, and has the least precise type annotations:[4]

$$(\mathsf{fun}(id : \star).(\mathsf{fun}(n : \star).id\ \mathsf{true})(id\ 42))\ (\mathsf{fun}(x : \star).x)$$

First, consider how we might manually migrate the program. One approach is to change the type of x to int (underlined below), and leave the other annotations unchanged:

$$(\mathsf{fun}(id : \star).(\mathsf{fun}(n : \star).id\ \mathsf{true})\ (id\ 42))\ (\mathsf{fun}(x : \underline{\mathsf{int}}).x)$$

Note that this program is well-typed and has a more precise type than the original. However, it produces a run-time type error on *id true*, whereas the original program does not. Fortunately, constraint generation rules out this migration: the outermost application coerces the argument type to $\star$. However, the argument type (int $\to$ int) is not a ground type, which App also requires.

The following type migration, also constructed manually, is the most precise migration that does not introduce a run-time error (changes to the original program are underlined):

$$(\mathsf{fun}(id : \underline{\star \to \star}).(\mathsf{fun}(n : \underline{\mathsf{int}}).id\ \mathsf{true})\ (id\ 42))\ (\mathsf{fun}(x : \star).x)$$

However, concluding that n has type int requires reasoning about the flow of values through the identity function. Our constraint generation rules can't find this solution. Instead, the most precise type allowed by our constraints gives *id* the type $\star \to \star$ and leaves n and x at type $\star$:

$$(\mathsf{fun}(id : \underline{\star \to \star}).(\mathsf{fun}(n : \star).id\ \mathsf{true})\ (id\ 42))\ (\mathsf{fun}(x : \star).x)$$

This example illustrates an important principle that we follow in constraint generation: if we generate a new coercion around an expression e to type $\star$, then we must also *constrain* the type of

[4] This is a variation of the example in Figure 5.

```
▷ The only annotations in e are ★
function PreciseMigrate(e)
    e₁ ← IntroduceMetavars(e)                          ▷ Replace every ★ with a fresh αs
    · ⊢ e₁ ⇒ e′, T₁, ϕ                                 ▷ Generate constraints and objectives
    for α ∈ ϕ do                                       ▷ The set of type metavariables in ϕ
        (declare-const α Typ)
    for w ∈ ϕ do                                       ▷ The set of weight variables in ϕ
        (declare-const w Bool)
        (assert-soft w 1)
    (check-sat ϕ)
    σ ← (get-model)                                    ▷ Model mapping type metavariables to types
    return Subst(σ, e′)                                ▷ Migrated program with explicit coercions
```

Fig. 11. Precise Type Migration.

e to be a ground type. As we grow the language with more types, the set of ground types will grow. When this happens, we update the definition of the *ground* predicate, but the rest of constraint generation remains unchanged.

The following theorem establishes that all models that satisfy our constraint generation rules produce well-typed expressions.

Theorem 4.1 (Type Migration Soundness). *If* $\Gamma \vdash e \Rightarrow e', T, \phi$ *and* σ *is a model for* ϕ*, then* $\textit{Subst}(\sigma, \Gamma) \vdash \textit{Subst}(\sigma, e') : \textit{Subst}(\sigma, T)$.

4.3 Solving Constraints for Precise Type Migration

Our formulation of constraint generation produces a constraint (ϕ) that may have multiple models, all of which encode valid type migrations of varying precision. Our goal in this section is to find as precise a migration as possible. To do this, we rely on the MaxSMT solver's ability to define *soft constraints*. The solver prefers solutions that obey these constraints, but can violate them when necessary to produce a model.

Our constraint generation rules adhere to the following recipe: every rule that introduces a coercion also introduces a fresh boolean variable (w) that is true when the coercion is trivial ($\textit{coerce}(T, T)$) and false otherwise. The Fun rule introduces one boolean variable, while the App rule introduces two, since it may introduce two non-trivial coercions.

We use the algorithm sketched in Figure 11. For each boolean variable, we produce a soft constraint asserting that w should hold (the corresponding coercion should be trivial if possible). Given these soft constraints, we check that the formula ϕ is satisfiable and get a model (σ) that assigns type metavariables to types. We then substitute metavariables with concrete types accordingly.

Example 2: A migration that is too precise. Consider the following as an input to our algorithm:

$$F_1 \triangleq \mathsf{fun}(f : \star).\mathsf{fun}(g : \star).(f\ 1) \times (g\ f)$$

The algorithm produces the following migration, which has the most precise types possible:

$$F_2 \triangleq \mathsf{fun}(f : \mathsf{int} \to \mathsf{int}).\mathsf{fun}(g : (\mathsf{int} \to \mathsf{int}) \to \mathsf{int}).(f\ 1) \times (g\ f)$$

But is the most precise type really the best type? The answer depends on how the original function was used. For example, in the following context F_2 is not substitutable for F_1:

Before (produces zero)	After (static type error)
$\underline{F_1}$ (fun($x : \star$).0) (fun(k : bool $\to \star$).k true)	$\underline{F_2}$ (fun($x : \star$).0) (fun(k : bool $\to \star$).k true)

The left-hand side type-checks and evaluates to 0, while the right-hand side has a static type error: the bool type in the (unmigrated) context is inconsistent with the migrated type int.

We might reason that it is acceptable to generate this static error. But there is a second, more serious problem: in a gradually typed language, it is possible to turn static type errors into run-time type errors. Consider the following variation where the annotation on k in the unmigrated version is less precise:

Before (produces zero)	After (dynamic type error)
$\underline{F_1}$ (fun$(x : \star).0$) (fun$(k : \star).k$ true)	$\underline{F_2}$ (fun$(x : \star).0$) (fun$(k : \star).k$ true)

Both programs above are well-typed. However, the static error from the previous example is now a dynamic error. As we argued in §2, making types more precise in a portion of a program can introduce run-time errors at the (higher-order) boundary between migrated and unmigrated code.

Perhaps we can address this problem by producing a different migration of F_1:

$$F_3 \triangleq \mathsf{fun}(f : \star \to \mathsf{int}).\mathsf{fun}(g : (\star \to \mathsf{int}) \to \mathsf{int}).(f\ 1) \times (g\ f)$$

This migration is less precise than F_2: although f and g must still be functions, they are not required to consume integers. It is therefore equivalent to F_1 in our unmigrated context.

Before (produces zero)	After (also produces zero)
$\underline{F_1}$ (fun$(x : \star).0$) (fun$(k : \star).k$ true)	$\underline{F_3}$ (fun$(x : \star).0$) (fun$(k : \star).k$ true)

Unfortunately, there are other contexts that lead to errors in F_3 that do not occur with F_1. For instance, the following program produces an error with F_3 but not F_1.

$$F_3\ (\mathsf{fun}(x : \star).x)\ (\mathsf{fun}(id : \star).(\mathsf{fun}(b : \star).0)\ (id\ \mathsf{true}))$$

We can address this problem with a migration with even lower precision:

$$F_4 \triangleq \mathsf{fun}(f : \star \to \star).\mathsf{fun}(g : (\star \to \star) \to \mathsf{int}).(f\ 1) \times (g\ f)$$

This expression does not produce the same error as the previous example, and is compatible with all our examples. However, we have lost a lot of information about how F_1 uses its arguments. To summarize, we have seen a series of migrations for F_1 in decreasing order of precision:

$$F_1 \sqsubseteq F_4 \sqsubseteq F_3 \sqsubseteq F_2$$

Our algorithm produces F_2, but the other, less precise migrations are compatible with more contexts. So, which migration is best? The answer depends on the context of use for the program. If the programmer is generating documentation, they may prefer the more precise migration. On the other hand, if they are adding types to a library and cannot make assumptions about the function's caller, they may desire the migration that is compatible with more contexts.

4.4 Choosing Alternative Migrations

Although the algorithm presented above produces the most precise migration that the TypeWhich constraints encode, we can also use TypeWhich to infer alternative migrations that prioritize other properties, such as contextual compatibility.

At first glance, it seems straightforward to weaken the more precise type inferred in the preceding section. Suppose the algorithm produces a migration e with type T, and we want a less precise type S ($S \sqsubseteq T$). It seems that we could simply wrap e in a coercion: $[coerce(T, S)]e$. Unfortunately, this purported solution is no different from the adversarial contexts presented above. The expression has the desired weaker type S, but gradual typing ensures that it behaves the same as the stronger type T at run-time, including producing the same run-time errors! Instead, we need to alter the type annotations that are *internal* to e.

```
1:  ▷ The only annotations in e are ★
2:  function MIGRATE(WEAKEN, e)
3:      e₁ ← INTRODUCEMETAVARS(e)        ▷ Replace ★s with fresh αs
4:      · ⊢ e₁ ⇒ e′, T₁, ϕ
5:      for α ∈ ϕ do                     ▷ The type metavariables in ϕ
6:          (declare-const α Typ)
7:      for w ∈ ϕ do                     ▷ The weight variables in ϕ
8:          (declare-const w Bool)
9:          (assert-soft w 1)
10:     (check-sat ϕ)
11:     σ ← (get-model)
12:     T₂ ← SUBST(σ, T₁)                ▷ The most precise type
13:     ϕ′ ← WEAKEN(T₂, T₁)
14:     (check-sat ϕ ∧ ϕ′)
15:     σ′ ← (get-model)
16:     return SUBST(σ′, e′)
```

$$T_1 = \star \text{ [Rastogi et al. 2012]}$$

$$P(T_1 \to T_2, \phi_T, b) \triangleq P(T_1, \texttt{(arr-in } \phi_T\texttt{)}, \neg b) \wedge P(T_2, \texttt{(arr-out } \phi_T\texttt{)}, b)$$

$$P(B, \phi_T, \text{true}) \triangleq \text{true}$$

$$P(B, \phi_T, \text{false}) \triangleq \phi_T = \star$$

$$P(\star, \phi_T, b) \triangleq \text{true}$$

$$\text{WEAKEN}(T, \phi_T) \triangleq P(T, \phi_T, \text{true})$$

Fig. 12. The Type Migration Algorithm.

Weakening Migrations. TYPEWHICH employs a two-step approach to type migration. We first generate constraints and calculate the most precise type possible (T_2), as described earlier (lines 3–12 of Figure 12; identical to Figure 11). We then apply the WEAKEN metafunction, which identifies all the base types in negative position in T_2 (following Rastogi et al. [2012]). The second argument to WEAKEN is a Typ-sorted formula that represents the type of the program (T_1). In WEAKEN's helper function P, we use this formula to identify portions of the output type in negative position, and constrain them to be equal to ★. The result is a constraint (ϕ') that weakens the program type.

Once we have the weakening constraint, we must update the type annotations in the migrated program and calculate the new weaker type. To do so, we run the solver once more with the added constraint (line 14). This produces a new model (line 15), which we use to substitute type metavariables and produce a fully annotated program.

It is worth reflecting on why a two-stage procedure is necessary. The first stage produces the most precise type that we can. This is necessary to discover a type skeleton that is as precise as possible; otherwise, we might miss some of the structure, e.g., by failing to predict arrow types. The second stage is necessary in order to propagate the constraints on the program's type back through the migrated program, which may involve arbitrary changes to internal type annotations.

Critically, the new set of constraints ϕ' must not impose unnecessary conditions on the type of the program. For example, suppose the original program e has a precise type int $\to$ int. Since this type only allows the context to provide int-arguments to e, we might conclude that a better type for e is $\star \to$ int. But this may be impossible: for instance, if e is the identity function, its argument and result types must be the same. On the other hand, if the body of e is a multiplication, then making the input type ★ does not affect the output type: it can remain int. By adding the constraint and re-solving, TYPEWHICH is able to distinguish between these two scenarios.

We note that there are several possible variations for WEAKEN. When migrating higher-order functions, it is useful to use a definition that turns base-typed inputs in negative position to ★, but preserves arrow types in the input. An alternative is to turn all input types to ★ to maximize compatibility, similar to Rastogi et al. [2012]. Our implementation of TYPEWHICH supports both of these and could be easily extended to other variations as well.

Our two-stage approach to contextual safety highlights the key trade-off between precision and compatibility in type migration. Our first-pass discovers the most precise types that we can;

Ground types **Base Types**
$G ::= \cdots \mid \text{ref}$ $B ::= \cdots \mid \text{unit}$
Constants **Types**
$c ::= \cdots \mid \text{unit}$ $T ::= \cdots \mid \text{ref}\ T$
Expressions
$e ::= \cdots$
$\mid \text{ref}\ e$ Create cell
$\mid\ !e$ Read cell
$\mid\ e_1 := e_2$ Write cell

Constraint Metafunctions
$ground \in T \to \phi$
$ground(T) = T \in B \vee T = \star \to \star \vee T = \text{ref}\ \star$

Type Representation

```
(declare-datatypes ()
  ((Typ (star) (int) (bool)
        (ref (to Typ))
        (arr (in Typ) (out Typ)))))
```

$$\textsc{If}\ \frac{\Gamma \vdash e_1 \Rightarrow e_1', T_1, \phi_1 \quad \Gamma \vdash e_2 \Rightarrow e_2', T_2, \phi_2 \quad \Gamma \vdash e_3 \Rightarrow e_3', T_3, \phi_3 \quad w_1, w_2, \alpha \text{ are fresh} \quad \phi_4 = ((T_1 = \text{bool} \wedge w_1) \vee (T_1 = \star \wedge \neg w_1)) \wedge ((T_2 = T_3 = \alpha \wedge w_2) \vee (\alpha = \star \wedge ground(T_2) \wedge ground(T_3) \wedge \neg w_2))}{\Gamma \vdash \text{if}\ e_1\ \text{then}\ e_2\ \text{else}\ e_3 \Rightarrow \text{if}\ [\underline{coerce}(T_1, \text{bool})]e_1'\ \text{then}\ [\underline{coerce}(T_2, \alpha)]e_2'\ \text{else}\ [\underline{coerce}(T_3, \alpha)]e_3', \alpha, \phi_1 \wedge \phi_2 \wedge \phi_3 \wedge \phi_4}$$

$$\textsc{Add}\ \frac{\Gamma \vdash e_1 \Rightarrow T_1, e_1'\phi_1 \quad \Gamma \vdash e_2 \Rightarrow T_2, e_2', \phi_2 \quad w, \alpha \text{ are fresh} \quad \phi_3 = ((\alpha = \text{int} \vee \alpha = \text{str}) \wedge \alpha = T_1 = T_2 \wedge w) \vee (\alpha = \star \wedge ground(T_1) \wedge ground(T_2) \wedge \neg w)}{\Gamma \vdash e_1 + e_2 \Rightarrow ([\underline{coerce}(T_1, \alpha)]e_1') + ([\underline{coerce}(T_2, \alpha)]e_2'), \alpha, \phi_1 \wedge \phi_2 \wedge \phi_3}$$

$$\textsc{Ref}\ \frac{\Gamma \vdash e \Rightarrow T, e', \phi_1 \quad \alpha \text{ and } w \text{ are fresh} \quad \phi_2 = (\alpha = \text{ref}\ T \wedge w) \vee (\alpha = \star \wedge ground(\text{ref}\ T) \wedge \neg w)}{\Gamma \vdash \text{ref}\ e \Rightarrow [\underline{coerce}(\text{ref}\ T, \alpha)]\text{ref}\ e', \alpha, \phi_1 \wedge \phi_2}$$

$$\textsc{Deref}\ \frac{\Gamma \vdash e \Rightarrow T, e', \phi_1 \quad \alpha, w \text{ are fresh} \quad \phi_2 = (T = \text{ref}\ \alpha \wedge w) \vee (T = \alpha = \star \wedge \neg w)}{\Gamma \vdash\ !e \Rightarrow\ !([\underline{coerce}(T, \text{ref}\ \alpha)]e'), \alpha, \phi_1 \wedge \phi_2}$$

$$\textsc{SetRef}\ \frac{\Gamma \vdash e_1 \Rightarrow T_1, e_1', \phi_1 \quad \Gamma \vdash e_2 \Rightarrow T_2, e_2', \phi_2 \quad \alpha \text{ and } w \text{ are fresh} \quad \phi_3 = (T_1 = \text{ref}\ \alpha \wedge T_2 = \alpha \wedge w) \vee (\alpha = \star \wedge ground(T_1) \wedge ground(T_2) \wedge \neg w)}{\Gamma \vdash e_1 := e_2 \Rightarrow ([\underline{coerce}(T_1, \text{ref}\ \alpha)]e_1') := [coerce(T_2, \alpha)]e_2', \text{Unit}, \phi_1 \wedge \phi_2 \wedge \phi_3}$$

Fig. 13. Extensions to the GTLC.

our second-pass sacrifices some of this precision to provide compatibility with a wider range of contexts.

Theorem 4.2 (Type Migration Completeness). *Every well scoped dynamic program e has a migration, i.e., there exists e', T, and ϕ such $\cdot \vdash e \Rightarrow e', T, \phi$ such that ϕ is satisfiable in some model σ.*

5 LANGUAGE EXTENSIONS

We now extend the GTLC and TypeWhich to support several common language features. These new features affect our constraint generation rules, but they do not change the migration algorithm.

Conditionals. Retrofitted type checkers for untyped languages employ a variety of techniques to give precise types to conditional expressions (§7). The GTLC-based languages (e.g., Kuhlenschmidt et al. [2019]) use a simpler approach: (1) the type of the test must be consistent with bool, and (2) the type of the expression is the least upper bound of the types of either branch.

The If rule in Figure 13 shows constraint generation for conditionals. The generated constraint (ϕ_4) has two conjunctions that 1) constrain the type of the condition to bool or $\star$, and 2) constrain the types of each branch to be identical types or distinct ground types (in which case, both are coerced to the unknown type).

Overloaded Operators. Many languages have overloaded built-in operators: for instance, the "+" operator is frequently used for addition and string concatenation. To support this, the run-time system has three operators available: (1) primitive addition, (2) primitive string concatenation, and

Name	Expression
FArg-Mismatch	$(\text{fun}(f : \star).f\ \text{true})\ (\text{fun}(x : \star).x + 1)$
Rank2-Poly-Id	$(\text{fun}(i : \star).(\text{fun}(a : \star).(i\ \text{true}))\ (i\ 5))\ (\text{fun}(x : \star).x)$
Unreachable-Err	$(\text{fun}(b : \star).b\ (\text{fun}(c : \star).(\text{fun}(x : tdyn).x\ x)\ 5\ 5)\ (\text{fun}(d : \star).0))\ (\text{fun}(t : \star).\text{fun}(f : \star).f)$
F-In-F-Out*	$(\text{fun}(f : \star).(\text{fun}(y : \star).f)\ (f\ 5))\ (\text{fun}(x : \star).10 + x)$
Order3-Fun*	$\text{fun}(f : \star).\text{fun}(x : \star).x\ (f\ x)$
Order3-IntFun*	$\text{fun}(f : \star).\text{fun}(g : \star).f\ g\ ((g\ 10) + 1)$
Double-F*	$\text{fun}(f : \star).f\ (f\ \text{true})$
Outflows*	$(\text{fun}(x : \star).x\ 5 + x)\ 5$
Precision-Relation*	$(\text{fun}(f : \star).f\ \text{true} + (\text{fun}(g : \star).g\ 5)\ f)\ (\text{fun}(x : \star).5)$
If-Tag	$\text{fun}(tag : \star).\text{fun}(x : \star).\text{if}\ tag\ \text{then}\ x + 1\ \text{else if}\ x\ \text{then}\ 1\ \text{else}\ 0$

Fig. 14. Our Type Migration Challenge.

(3) a complex operation whose behavior depends on the run-time types of its arguments. Type migration can reveal the type at which an overloaded operator is used, which can help programmers understand their code and improve run-time performance. The constraint generation rule for "+" in Figure 13 introduces a boolean-sorted variable (w) that is true when the operands both have type int or str; when the variable is false, the constraint requires the two arguments to have type $\star$. Thus, it favors solutions that do not employ $\star$ when possible.

Mutable Data Structures. TypeWhich supports ML-style mutable references and mutable vectors. There are several ways to add mutable references to the GTLC [Herman et al. 2011; Siek and Taha 2006; Siek et al. 2015c]. However, all approaches share the following property: in untyped code, where all mutable cells contain $\star$-typed values, the *only* reason that reading or writing fails is when the expression in reference position is not a reference. In constraint generation, we are careful to avoid solutions that may introduce other kinds of errors.

The least precise reference type is a reference to the unknown type (ref $\star$), so we add this to the set of ground types (Figure 13). In the constraint generation rule for writes, we require that either (1) the type of value written is exactly the referenced type, or (2) both the reference and the value written are ground types. The restriction to ground types is necessary because, as in the function case, once a reference is coerced to $\star$, we have no way to recover its original type; allowing non-ground types to be coerced to $\star$ can introduce run-time errors. TypeWhich also supports mutable vectors implemented along the same lines.

Other language features. The implementation of TypeWhich supports a variety of other language features, including tuples, `let`, and a `fix` construct. Many of these are necessary to support the Grift programming language, which we use in our evaluation. Constraint generation rules for these extensions can be found in Phipps-Costin et al. [2021]

6 EVALUATION

This section presents the first comprehensive comparison of several type migration algorithms from the literature (along with TypeWhich). We first compare five type migration tools on a suite of 22 programs, including several new benchmarks. We also evaluate TypeWhich using the Grift benchmarks from Kuhlenschmidt et al. [2019] to show that TypeWhich can reconstruct hand-written type annotations in Grift.

6.1 Gradual Type Migration Benchmarks

We evaluate type migration tools using a two-part benchmark suite: a suite of benchmarks from Migeed and Palsberg [2020], and a new suite of *challenge programs* designed to illustrate the

strengths and weaknesses of different approaches to type migration. Our proposed challenge suite is presented in Figure 14. We describe the ten programs below. Although TypeWhich supports several extensions to the GTLC (§5), we do not use them in the challenge suite so that we can run as many tools as possible. (The final If-Tag benchmark is an exception.)

(1) FArg-Mismatch: crashes at run-time, because the functional argument f expects an integer, but is applied to a boolean.
(2) Rank2-Poly-Id (based on Figure 5): defines the identity function and applies it to a number and a boolean. It uses a Church encoding of let-binding and sequencing that would require rank-2 polymorphism in an ML dialect.
(3) Unreachable-Err (based on Figure 2): has a crashing expression similar to FArg-Mismatch, but it is unreachable. The example encodes a conditional as a Church boolean.
(4) F-In-F-Out: defines a local function f that escapes.
(5) Order3-Fun: a higher-order function that receives two functions f and x. Moreover, the body calculates $f\ x$, so f must be a higher-order function itself.
(6) Order3-IntFun: similar to Order3-Fun, but the program uses operations that force several types to be int.
(7) Double-F: calculates $f\ (f\ \mathsf{true})$. The inner application suggests that f's argument must be bool. However, that would rule out $\mathsf{fun}(x : \star).0$ as a possible value for f.
(8) Outflows: defines a function that uses its argument as two different types. However, the function receives an integer.
(9) Precision-Relation: names a function f that must receive $\star$, since f is applied twice to two different types. However, the second application re-binds f to g, thus g may have a more precise type.
(10) If-Tag: receives a boolean and uses its value to determine the type of x. Conditionals are not in the core GTLC and not supported by all the tools that we consider. However, it is essential to think through conditionals, since they induce a type constraint between both branches, and a Church encoding incurs a significant loss of precision.

Some of these programs (marked with an asterisk in Figure 14) can be given types using Hindley-Milner type inference via translation into OCaml or Haskell. Doing so reveals important differences between conventional static types and the GTLC. For example, the most general type of Order3-Fun is a type scheme with two type variables. The GTLC does not support polymorphism, so a type migration must use $\star$ rather than the more precise type. In contrast, the type of f in Double-F is bool $\to$ bool. However, f can have other types in the GTLC.

6.2 Benchmarked Type Migration Tools

We evaluate the performance of the following tools, which have a variety of different goals, which we describe below:

(1) TypeWhich: our tool, which we run in two modes: (a) TypeWhich-P produces a safe migration, and (b) TypeWhich-C produces a compatible migration. In both modes, TypeWhich maximizes precision, and migrates all closed programs.
(2) Gtubi: *gradual typing with unification-based inference* [Siek and Vachharajani 2008] is the earliest work on gradual type migration. It produces safe migrations.
(3) InsAndOuts: our implementation of the algorithm in Rastogi et al. [2012]. The algorithm produces compatible migrations.
(4) MaxMigrate: Migeed and Palsberg [2020] presents algorithms for several migration problems. We use the *maximal migration* tool, which produces a migration that cannot be made more precise. The tool searches for migrations by building types up to some depth (we use depth

Tool	Migrations (% of programs)	Safe Migrations (% of programs)	Compatible Migrations (% of programs)	Improved Type Annotations (% of annotations)
Gtubi	0.36	0.36	0.32	0.76
InsAndOuts	0.91	0.91	0.91	0.43
MGT	1.00	1.00	0.86	0.48
MaxMigrate	0.77	0.64	0.32	0.73
TypeWhich-C	1.00	1.00	1.00	0.31
TypeWhich-P	1.00	1.00	0.86	0.57

Fig. 15. Concise performance metrics for automated type migration tools (higher numbers are better).

five as in the paper). A single program may have several maximal migrations; we take the first migration the tool produces. We halt with no output if no migration is found.

(5) MGT: our implementation of the algorithm in Campora et al. [2018b] for migrating untyped or partially typed programs. We start from untyped code (all functions annotated with $\star$), and take the first migration it produces.

6.3 Concise Evaluation

Using our suite of benchmarks, Figure 15 shows how the aforementioned tools perform on the axes of safety, compatibility, and precision. The first column of numbers reports the percentage of programs that are successfully migrated, and it is important to take this column into consideration when interpreting the other columns. For example, the final column suggests that the Gtubi outperforms MaxMigrate on type precision. However, the first column shows that Gtubi only migrates half as many programs as MaxMigrate.

6.4 How Should Type Migration Tools Be Evaluated?

The concise evaluation masks many subtle issues that arise in type migration. For instance, using the total number of type annotations improved is a good metric for type precision, but reporting only precision obscures the fact that not all improvements are alike: some change the behavior of the original program, while others preserve its semantics. We have also illustrated how type precision can come at the expense of compatibility with unmigrated code. This sacrifice may sometimes be warranted, but when a function is migrated, it should remain usable with at least *some arguments*. This seems like a trivial point, but consider the following migration:

Original Program	Migrated Program
$\mathsf{fun}(f : \star).\mathsf{fun}(x : \star).fxx$	$\mathsf{fun}(f : \mathsf{int} \to \mathsf{bool} \to \star).\mathsf{fun}(x : \star).fxx$

The migrated program has types that cannot be made more precise. However, the type of f requires x to be both an integer and a boolean, and thus renders the function unusable.

We propose a multi-stage evaluation process for automated type migration tools. For each tool, (1) we start with the full suite of programs and ask, *How many programs does the tool reject with static errors?* (2) We take the *remaining programs* and ask, *How many migrated programs crash with a new dynamic type error?* (3) We take the *remaining programs* and ask, *How many migrated programs are functions that are rendered unusable?* (4) We take the remaining programs and ask two final questions: (a) *How many migrated programs are functions with types that are incompatible with some untyped contexts?* and (b) *How many type annotations, counted across all remaining programs, are not improved by migration?*

For our evaluation, we partially automate the multi-step process described above. To trigger errors, benchmarks that are functions require an input, which we construct manually. For step 3, we inspect every program that crashes on some input, to determine if there is any other input that

Tool	The tool rejects the program, e.g., 1 + true Rejected / Total Programs	(fun(id : ⋆).id 1) (fun(x : bool).x) crashes New Dynamic Errors / Remaining Programs	(fun(f : int → bool → ⋆).fun(x : ⋆).$f\,x\,x$ is unusable Unusable Functions / Remaining Programs	fun(x : int).x is restricted Restricted Functions / Remaining Programs	Not Improved / Total ⋆
Gtubi	14 / 22	0 / 8	0 / 8	1 / 8	4 / 17
InsAndOuts	2 / 22	0 / 20	0 / 20	0 / 20	24 / 42
MGT	0 / 22	0 / 22	0 / 22	3 / 22	30 / 58
MaxMigrate	5 / 22	3 / 17	3 / 14	4 / 11	6 / 18
TypeWhich-C	0 / 22	0 / 22	0 / 22	0 / 22	40 / 58
TypeWhich-P	0 / 22	0 / 22	0 / 22	3 / 22	25 / 58

Fig. 16. Summary of type migration results. TypeWhich-C favors compatibility with unmigrated code, and TypeWhich-P favors precision. Above each column, we show an example of the kind of migrated program we count in that column. **Warning:** This figure requires careful interpretation. For example, the *Restricted Functions* column shows that both InsAndOuts and TypeWhich-C produce zero restricted migrations. However, the denominator is not the same: InsAndOuts rejects two programs in the first column and only 20 programs remain. In contrast, TypeWhich-C runs on the full suite of 22 programs. Thus, TypeWhich-C is arguably better than InsAndOuts since it does not reject any programs. Furthermore, consider Gtubi and MaxMigrate, both of which can produce restricted migrations. The penultimate column shows that Gtubi only produces one restricted program, whereas MaxMigrate produces three, which suggests that Gtubi is better. However, Gtubi statically rejects far more programs in the first column, thus more programs remain for MaxMigrate.

will make it not crash. In step 4a, to label a benchmark as incompatible with some contexts, we manually provide a context that leads to a dynamic error, and the benchmarking framework verifies that the error definitely occurs. Conversely, to label a benchmark as compatible with all contexts, we provide a hand-written, compatible migration and the benchmarking framework verifies that the migration returned by the tool is less precise than the hand-written migration.

Note that the denominator (potentially) decreases at each stage: if a tool fails to migrate a program, then it is impossible to assess whether the migrated program crashes with a dynamic error. Moreover, we do not want to give a system credit for increasing the precision of a type if the refinement triggers a new dynamic error (i.e., it was an unsafe migration).

6.5 Comprehensive Evaluation

The results of our evaluation illustrate the various strengths and weaknesses of different approaches to automated type migration. Before diving into the details of the complete results, we present a bird's-eye view of the evaluation scheme proposed in §6.4.

How to Interpret Figure 16. The summary table in Figure 16 must be interpreted carefully. Every row in the table shows the results of a single tool. The columns can be interpreted as follows:

- The rightmost column counts the number of type annotations that are *not improved* (i.e., lower is better), out of the total number of annotations in the migrated programs that are either context-restricted, or compatible with all contexts. Thus we do not count type annotation improvements that lead to immediate crashes or render the function unusable. Thus, *if safety and precision are the primary concern, this column is most informative.*
- The first three columns count programs that fail to migrate safely in three ways: programs that are rejected statically, programs that crash after migration with new dynamic errors, and programs that become unusable functions after migration.

- In the penultimate column, the denominator counts programs that migrate safely without the errors mentioned above, and the numerator counts programs that still produce errors in some contexts. Therefore, if *safety is the only concern*, then the denominator of the penultimate column (higher values are better) is the most informative. On the other hand, if *safety and compatibility are both significant*, then the difference between the denominator and numerator is most informative (a higher value is better).

Thus this figure has a deliberate safety bias: it discounts rejected programs, and increasing type precision, when doing so introduces errors. However, the figure makes it possible to compare tools on three axes: precision (conditioned on safety), compatibility (which is meaningless without safety), and safety alone.

Discussion of Figure 16. We include results from running TypeWhich in two different modes: TypeWhich-P prioritizes precision, while TypeWhich-C prioritizes contextual compatibility. By design, TypeWhich does not produce static or dynamic errors. When it is configured for type precision (TypeWhich-P), it does restrict the inputs of three functions. However, even in this mode, the remaining 19 programs remain compatible with all callers. On the other hand, when it is configured to prioritize contextual compatibility (TypeWhich-C), no programs are restricted, but fewer types are improved.

Like TypeWhich, MGT restricts some functions, but does not produce static or dynamic errors. Gtubi rejects several programs statically and restricts the behavior of some functions. However, it does not introduce any dynamic errors. MaxMigrate rejects a few programs: some do not have maximal migrations, on others it cannot find a migration within its search space, and one of our programs uses a conditional, which is unsupported. In addition, the tool introduces run-time errors in some programs, and makes some functions unusable. InsAndOuts rejects two programs.[5] On the remaining programs, it produces migrations that are compatible with arbitrary unmigrated code as intended. In fact, when we prioritize compatibility with unmigrated code, InsAndOuts outperforms all other approaches.

The right-most column of the table reports the number of type annotations that are *not* improved, and this must be interpreted very carefully. The point of gradual typing is that ⋆ serves as an "escape hatch" for programs that cannot be given more precise types. Our suite includes programs that *must* have some ⋆s, so every tool will have to leave some ⋆s unchanged. We naturally want a tool to improve as many types as possible, so we may prefer a tool that has the fewest number of unimproved types. However, notice that the denominator varies considerably. For example, TypeWhich-P cannot improve about half the annotations, but it does not introduce any errors. In contrast, the oldest tool—Gtubi–only leaves a small fraction of annotations unimproved, but it statically rejects the majority of programs.

Challenge Set Results. We now examine performance on the challenge set in more detail. Figure 17 shows the migrated challenge programs produced by these tools. We present and discuss results produced by running TypeWhich to prioritize precision (TypeWhich-P); results from TypeWhich-C can be found in the appendix.

Examining the detailed output on the challenge set programs reveals interesting differences in the migrations inferred by the various type migration tools, reflecting their differing priorities.

(1) FArg-Mismatch: InsAndOuts produces the most precise and informative result, showing that x is a boolean next to $x + 100$, which helps locate the error in the program.

[5] These are two programs from the Migeed and Palsberg [2020] benchmarks. From correspondence with the authors of Rastogi et al. [2012], our implementation seems faithful to the presentation in the paper, and the original implementation for Adobe ActionScript is no longer accessible.

FArg-Mismatch	MGT	(fun f : $\star$.f true)(fun x : int.x + 100)
	MaxMigrate	(fun f : bool $\rightarrow$ int.f true)(fun x : $\star$.x + 100)
	InsAndOuts	(fun f : bool $\rightarrow$ int.f true)(fun x : bool.[bool!]x + 100)
	TypeWhich-P	(fun f : $\star$ $\rightarrow$ int.f true)(fun x : $\star$.x + 100)
	Gtubi	constraint solving error
Rank2-Poly-Id	MGT	no improvement
	MaxMigrate	(fun i : $\star$ $\rightarrow$ $\star$.(fun a : int.i true)(i 5))(fun x : bool.x)
	InsAndOuts	identical to TypeWhich
	TypeWhich-P	(fun i : $\star$ $\rightarrow$ $\star$.(fun a : $\star$.i true)(i 5))(fun x : $\star$.x)
	Gtubi	constraint solving error
Unreachable-Err	MGT	identical to TypeWhich-P
	MaxMigrate	No maximal migration
	InsAndOuts	(fun b : (α $\rightarrow$ α) $\rightarrow$ ($\star$ $\rightarrow$ int) $\rightarrow$ $\star$ $\rightarrow$ int.b (fun c : $\star$.(fun x : int.x x) 5 5) (fun d : $\star$.0)) (fun t : α $\rightarrow$ α.fun f : $\star$ $\rightarrow$ int.f)
	TypeWhich-P	(fun b : ($\star$ $\rightarrow$ $\star$) $\rightarrow$ ($\star$ $\rightarrow$ int) $\rightarrow$ $\star$ $\rightarrow$ int.b (fun c : $\star$.(fun x : $\star$.x x) 5 5) (fun d : $\star$.0)) (fun t : $\star$ $\rightarrow$ $\star$.fun f : $\star$ $\rightarrow$ int.f)
	Gtubi	constraint solving error
F-In-F-Out	MGT	identical to TypeWhich
	MaxMigrate	(fun f : int $\rightarrow$ $\star$.(fun y : bool.f)(f 5))(fun x : int.10 + x)
	InsAndOuts	(fun f : $\star$ $\rightarrow$ int.(fun y : int.f)(f5))(fun x : $\star$.10 + x)
	TypeWhich-P	(fun f : int $\rightarrow$ int.(fun y : int.f)(f5))(fun x : int.10 + x)
	Gtubi	identical to TypeWhich
Order3-Fun	MGT	fun f : ($\star$ $\rightarrow$ $\star$) $\rightarrow$ $\star$.fun x : $\star$ $\rightarrow$ $\star$.x(f x)
	MaxMigrate	fun$f$: int $\rightarrow$ int.fun$x$: $\star$.x(f x)
	InsAndOuts	no improvement
	TypeWhich-P	identical to MGT
	Gtubi	fun f : (α $\rightarrow$ β) $\rightarrow$ α.fun x : α $\rightarrow$ β.x(f x)
Order3-IntFun	MGT	identical to TypeWhich
	MaxMigrate	fun$f$: int $\rightarrow$ int $\rightarrow$ int.fun g : $\star$.f g(g10 + 1)
	InsAndOuts	no improvement
	TypeWhich-P	fun f : (int $\rightarrow$ int) $\rightarrow$ int $\rightarrow$ $\star$.fun g : int $\rightarrow$ int.f g(g10 + 1)
	Gtubi	fun f : (int $\rightarrow$ int) $\rightarrow$ int $\rightarrow$ α.fun g : int $\rightarrow$ int.f g(g10 + 1)
Double-F	MGT	identical to TypeWhich
	MaxMigrate	fun f : $\star$ $\rightarrow$ int.f($f$true)
	InsAndOuts	no improvement
	TypeWhich-P	fun f : bool $\rightarrow$ bool.f($f$true)
	Gtubi	identical to TypeWhich
Outflows	MGT	no improvement
	MaxMigrate	no improvement
	InsAndOuts	(fun$x$: int.[int!]x 5 + x) 5
	TypeWhich-P	identical to InsAndOuts
	Gtubi	constraint solving error
Precision-Relation	MGT	(fun f : $\star$.(f true) + ((fun g : int $\rightarrow$ $\star$.g 5)f))(fun x : $\star$.5)
	MaxMigrate	(fun f : bool $\rightarrow$ int.($f$true) + ((fun g : $\star$ $\rightarrow$ int.g5)f))(fun x : bool.5)
	InsAndOuts	(fun f : $\star$ $\rightarrow$ int.(f true) + ((fun g : int $\rightarrow$ int.g 5)f))(fun x : $\star$.5)
	TypeWhich-P	(fun f : $\star$ $\rightarrow$ int.(f true) + ((fun g : $\star$ $\rightarrow$ int.g 5)f))(fun x : $\star$.5)
	Gtubi	constraint solving error
If-Tag	MGT	Identical to TypeWhich-P
	MaxMigrate	conditionals unsupported
	InsAndOuts	no improvement
	TypeWhich-P	fun *tag* : bool.fun x : bool.if *tag* then ([bool!]x) + 1 else if x then 1 else 0
	Gtubi	conditionals unsupported

Fig. 17. Migrations of the challenge set with TypeWhich in precise mode.

(2) Rank2-Poly-Id: InsAndOuts and TypeWhich produce the best result that does not introduce a run-time error. MaxMigrate produces the most precise static type, but has a dynamic type error.
(3) Unreachable-Err: TypeWhich, MGT, and InsAndOuts are the only tools that produce a result. The erroneous and unreachable portion gets the type $\star$ in TypeWhich; whereas InsAndOuts produces a type variable. The rest of the program has informative types.
(4) F-In-F-Out: MGT, Gtubi, and TypeWhich produce the most precise result. MaxMigrate produces an alternative, equally precise type, but introduces a dynamic type error.
(5) Order3-Fun: Gtubi produces the best result. Its result has type variables, thus is a type scheme. However, in a larger context, these variables would unify with concrete GTLC types. MGT and TypeWhich produce a similar result, but with $\star$. MaxMigrate produces $\mathsf{int} \to \mathsf{int}$ as the type of f, which is maximal, but introduces a subtle problem: $(f\ x)$ requires x to be an integer, but $x\ (f\ x)$ requires x to be a function.
(6) Order3-IntFun: the results are similar to Order3-Fun, with Gtubi again doing the best. However, since the program forces certain types to be int, TypeWhich and MGT now produce the same result.
(7) Double-F: MaxMigrate produces the best result. The most informative annotation on f that is compatible with all contexts is $\star \to \star$; no tool produces this type.
(8) Outflows: InsAndOuts and TypeWhich produce the best result. This program requires x to have two different types and thus crashes. Because the function receives an integer for x, Rastogi et al. [2012] and TypeWhich give x the type int. The other tools are not capable of reasoning in this manner. In a modification of this example where x is used with different types in each branch of a conditional, all tools would likely produce similar results.
(9) Precision-Relation: InsAndOuts produces the most precise type that does not introduce a run-time type inconsistency. TypeWhich does not give g the most precise type; MGT does not improve the type of f; and MaxMigrate finds a maximal migration that constrains f's argument to bool.
(10) If-Tag: Gtubi and MaxMigrate do not support conditionals. TypeWhich-P and MGT produce an unusual result that restricts the type of the argument x to bool and turns the $x + 1$ into $([\mathsf{bool}!]x) + 1$. If we were migrating a larger program that had this function as a sub-expression, and this function were actually applied to values of x with different types, the type of x would be $\star$.

Migeed and Palsberg [2020] Benchmarks. Migeed and Palsberg [2020] compare their maximal migration tool to the type migration tool in Campora et al. [2018b]. We extend the comparison to include TypeWhich, Gtubi, and InsAndOuts. The artifact that accompanies this paper includes the complete suite of benchmark results, and we include all of these benchmarks in our summary (Figure 16).

Summary. Our type migration challenge suite is designed to highlight the strengths and weaknesses of different algorithms. As discussed in §2, the competing goals of the type migration problem lead to a range of compromises; we do not claim that any one approach is best, since each approach reflects a different weighting of priorities. Because our challenge programs are synthetic, it would be possible to build a large set of programs that favor one tool at the expense of others. Our goal has been instead to curate a small set that illustrates a variety of weaknesses in every tool. In addition, our challenge programs are unlikely to be representative of real-world type migration problems. A more thorough evaluation would require scaling type migration tools to a widely-used language with a corpus of third-party code, which is beyond the scope of this paper.

6.6 Grift Performance Benchmarks

Kuhlenschmidt et al. [2019] present a benchmark suite to evaluate the performance of Grift programs (running time and space efficiency). Grift extends the GTLC with floating-point numbers, characters, loops, recursive functions, tuples, mutable references, vectors, and several primitive operators. Each benchmark has two versions: an untyped version and a fully-typed, hand-annotated version. We use TypeWhich (in precise mode) to migrate every untyped benchmark, and compare the result to the human type annotations. TypeWhich supports all Grift features except equirecursive types. However, because Grift's equirecursive types do not introduce new expression forms, TypeWhich can still be run on all programs: it just fails to improve annotations that require them.

TypeWhich performs as follows on the Grift benchmarks:

- **On 9 of 11 benchmarks**, TypeWhich produces exactly the same type annotations as the hand-typed version.
- **N-body** defines a number of unused functions over vectors. Since they are under-constrained, TypeWhich makes some arbitrary choices. On the reachable portion of the benchmark, we produce exactly the same type annotations as the hand-typed version.
- **Sieve** includes a library for stream processing, and the typed version of the benchmark gives streams a recursive type: $\mu\ s.\mathsf{Int} \times (\rightarrow s)$. TypeWhich migrates streams to type $\star$, which forces the stream elements to have type $\star$, due to ground constraints.

6.7 Implementation and Performance

The TypeWhich tool is open-source and written in approximately 12,000 lines of Rust. This code includes our new migration algorithm, implementations of the migration algorithms from Rastogi et al. [2012] and Campora et al. [2018b], and a unified evaluation framework that supports all the third-party tools that we use in our evaluation. The evaluation framework is designed to automatically validate the evaluation results we report. For example, to report that a migrated function is not compatible with all untyped contexts, our framework requires an example of a context that distinguishes between the migrated and original program, and runs both programs in the given context to verify that they differ. The framework also ensures that migrated programs are well-typed and structurally identical to the original program.

We perform all our experiments on on a virtual machine with 4 CPUs and 8 GB RAM, running on an AMD EPYC 7282 processor. The full suite consists of 892 LOC and 33 programs. TypeWhich produces migrations for our entire suite of benchmarks in under three seconds.

7 RELATED WORK

There is a growing body of work on automating gradual type migration and related issues. Our work is most closely related to the four algorithms we evaluate in §6. Siek and Vachharajani [2008] substitutes metavariables that appear in type annotations with concrete types, using a variation on unification. Rastogi et al. [2012] builds a type inference system for ActionScript. Their system ensures that inference never fails and produces types that are compatible with all untyped contexts. Campora et al. [2018b] uses variational typing to heuristically tame the exponential search space of types [Chen et al. 2014]. Migeed and Palsberg [2020] present decidability results for several type migration problems, including finding maximally precise migrations.

The aforementioned work relies on custom constraint solving algorithms. A key contribution of this paper is an approach to gradual type migration using an off-the-shelf MaxSMT solver, which makes it easier to build a type migration tool. In addition, we present a comprehensive evaluation comparing all five approaches. As part of this effort, we have produced new, open-source implementations of the algorithms presented in Rastogi et al. [2012] and Campora et al. [2018b].

Henglein [1994] introduces the theory of coercions that we use; Henglein and Rehof [1995] present an efficient compiler from Scheme to ML that inserts coercions when necessary. This work also uses a custom constraint solver and a complex graph algorithm. The latter defines a *polymorphic safety* criterion, which is related to our notion of a context-restricted type migration (Definition 3.4). Coercions are equivalent to casts [Siek et al. 2015a]; both are what Vitousek et al. [2014] calls the *guarded* approach to runtime type enforcement. Vitousek et al. introduces the *transient* approach, which only checks ground types/type tags at runtime rather than wrapping/proxying used in the guarded semantics. The transient approach is more efficient on the "mixed programs" we generate but offers weaker guarantees and can mask errors [Greenman and Felleisen 2018; Greenman and Migeed 2018]. TypeWhich is based on the guarded model; we could reframe TypeWhich to match transient by only coercing between ground types. TypeWhich-P infers types based on particular elimination forms—just like transient. Additionally, TypeWhich-C makes no assumptions about program contexts, just like transient's "open world soundness" [Vitousek et al. 2017].

Garcia and Cimini [2015] extend Siek and Vachharajani's work to infer principal types. Since we focus on monomorphic types, we do not directly compare against their algorithm. Miyazaki et al. [2019] build on Garcia and Cimini's work, discussing the coherence issues what we point out in §2: types induce run-time checks that can affect program behavior. However, while we migrate all programs, Miyazaki et al. use *dynamic type inference* to discover type inconsistencies and report them as run-time errors. Castagna et al. [2019] propose another account of gradual type inference that supports many features (let-polymorphism, recursion, and set-theoretic types). They do not consider run-time safety. Finally, Campora et al. [2018a] extend their previous work [2018b] with a cost model for selecting migrations. Like us, they discuss trade-offs in type migration, although they focus on type precision and performance, rather than semantics preservation.

Tobin Hochstadt and Felleisen [2008], Guha et al. [2011], Chugh et al. [2012], and Vekris et al. [2015] are examples of retrofitted type checkers for untyped languages that feature flow-sensitivity. These tools require programmers to manually migrate their code, while we focus on automatic type migration. However, they go beyond our work by considering flow-sensitivity.

Anderson et al. [2005] present type inference for a representative fragment of JavaScript. However, the approach is not designed for gradual typing, where portions of the program may be untyped. Similarly, Chandra et al. [2016] infer types for JavaScript programs with the goal of compiling them to run efficiently on low-powered devices; their approach is not gradual by design and deliberately rejects certain programs.

Pavlinovic et al. [2014] formulate a MaxSMT problem to localize OCaml type errors. We also use MaxSMT and encode types in a similar manner. However, both the form of our constraints and the role of the MaxSMT solver are very different. In error localization, the MaxSMT problem helps isolate type errors from well-typed portions of the program. In our work, the entire program must be well-typed. Moreover, our constraints allow several typings, and we use soft constraints to guide the MaxSMT solver towards solutions with fewer coercions.

Soft Scheme [Wright and Cartwright 1997] infers types for Scheme programs. However, its type system is significantly different from the GTLC, which hinders comparisons to contemporary type migration tools for the GTLC. Flanagan [1997, p. 41]'s discussion of how Soft Scheme's sophistication can lead to un-intuitive types inspired work on set-based analysis of Scheme programs: Flanagan et al. [1996] map program points to sets of abstract values, rather than types.

Thorn [Bloom et al. 2009] presents an approach to gradual typing where ordinary typed expressions cannot have runtime type errors, thus do not require runtime checks. Instead, the programmer must use *like types* at the interface between typed and untyped code, where runtime type errors may occur. The GTLC does not make this distinction manifest, but it is essential for a safe type migration: types introduced by a tool must not introduce new runtime failures.

TypeWhich relies on a traditional approach to constraint generation: we carefully write constraint generation rules by hand. It is possible to complement hand-written rules with additional sources of information to get significantly better results. Some work uses run-time profiling to guide type inference [An et al. 2011; Furr et al. 2009; Saftoiu 2010]. More recent work uses programmer-supplied heuristics to guide type inference to produce more readable results [Kazerounian et al. 2020; Ren and Foster 2016]. These approaches preserve type soundness. It is also possible to produce type migrations using supervised machine learning [Hellendoorn et al. 2018; Malik et al. 2019; Pradel et al. 2020; Wei et al. 2020].

8 CONCLUSION

We present TypeWhich, a new approach to type migration for the GTLC that is more flexible than previous approaches in two key ways. First, we formulate constraints for an off-the-shelf MaxSMT solver rather than building a custom constraint solver, which makes it easier to extend TypeWhich. We demonstrate this flexibility by adding support for several language features beyond the core GTLC. Second, TypeWhich can produce alternative migrations that prioritize different goals, such as type precision and compatibility with unmigrated code. This makes TypeWhich a more flexible approach, suitable for migration in multiple contexts.

We also contribute to the evaluation of type migration algorithms. We define a multi-stage evaluation process that accounts for multiple goals of type migration. We present a "type migration challenge set": a benchmark suite designed to illustrate the strengths and weaknesses of various type migration algorithms. We evaluate TypeWhich alongside four existing type migration systems. Toward this end, we contribute open-source implementations of two existing algorithms from the literature, which we incorporate into a unified framework for automated type migration evaluation. We hope these evaluation metrics, new benchmarks, and benchmarking framework will aid future work by illuminating the differences among the many approaches to gradual type migration.

ACKNOWLEDGEMENTS

We thank the OOPSLA reviews for their thoughtful feedback. We thank Aseem Rastogi and Zeina Migeed for helpful discussions about their work. We thank Matthias Felleisen and Shriram Krishnamurthi for reading early drafts, and discussing their experience with Soft Scheme and MrSpidey. We thank Laurence Tratt for help with `grmtools` [Diekmann and Tratt 2020], which TypeWhich uses significantly. This work is partially supported by the National Science Foundation under grants CCF-2102288 and CCF-2129344.

REFERENCES

Jong-hoon David An, Avik Chauduri, Jeffrey S. Foster, and Michael Hicks. 2011. Dynamic Inference of Static Types for Ruby. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*.

Christopher Anderson, Paola Giannini, and Sophia Drossopoulou. 2005. Towards Type Inference for JavaScript. In *European Conference on Object-Oriented Programming (ECOOP)*.

Clark Barrett, Iger Shikanian, and Cesare Tinelli. 2007. An Abstract Decision Procedure for a Theory of Inductive Data Types. *Journal on Satisfiability, Boolean Modeling and Computation* 3, 1–2 (2007), 21–46.

Nikolaj Bjørner, Anh-Dung Phan, and Lars Fleckenstein. 2015. *νZ*: An Optimizing SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*.

Bard Bloom, John Field, Nathaniel Nystrom, Johan Östlund, Gregor Richards, Rok Strniša, Jan Vitek, and Tobias Wrigstad. 2009. Thorn: Robust, Concurrent, Extensible Scripting on the JVM. In *ACM SIGPLAN Conference on Object Oriented Programming, Systems, Languages and Applications (OOPSLA)*.

John Peter Campora, Sheng Chen, Martin Erwig, and Eric Walkingshaw. 2018b. Migrating Gradual Types. *Proceedings of the ACM on Programming Languages (PACMPL)* 2, POPL (2018).

John Peter Campora, Sheng Chen, and Eric Walkingshaw. 2018a. Casts and Costs: Harmonizing Safety and Performance in Gradual Typing. *Proceedings of the ACM on Programming Languages (PACMPL)* 2, ICFP (2018).

Giuseppe Castagna, Victor Lanvin, Tommaso Petrucciani, and Jeremy G. Siek. 2019. Gradual Typing: A New Perspective. *Proceedings of the ACM on Programming Languages (PACMPL)* 3, POPL (2019).

Satish Chandra, Colin S. Gordon, Jean-Baptiste Jeannin, Cole Schlesinger, Manu Sridharan, Frank Tip, and Young-Il Choi. 2016. Type inference for static compilation of JavaScript. In *ACM SIGPLAN Conference on Object Oriented Programming, Systems, Languages and Applications (OOPSLA)*.

Sheng Chen, Martin Erwig, and Eric Walkingshaw. 2014. Extending Type Inference to Variational Programs. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 36, 1 (2014).

Ravi Chugh, Patrick M. Rondon, and Ranjit Jhala. 2012. Nested Refinements for Dynamic Languages. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*.

Lukas Diekmann and Laurence Tratt. 2020. Don't Panic! Better, Fewer, Syntax Errors for LR Parsers. In *European Conference on Object-Oriented Programming (ECOOP)*.

Cormac Flanagan. 1997. *Effective Static Debugging via Componential Set-based Analysis.* Ph.D. Dissertation. Rice University.

Cormac Flanagan, Matthew Flatt, Shriram Krishnamurthi, Stephanie Weirich, and Matthias Felleisen. 1996. Catching Bugs in the Web of Program Invariants. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*.

Michael Furr, Jong-hoon David An, and Jeffrey S. Foster. 2009. Profile-Guilding Static Typing for Dynamic Scripting Languages. In *ACM SIGPLAN Conference on Object Oriented Programming, Systems, Languages and Applications (OOPSLA)*.

Ronald Garcia and Matteo Cimini. 2015. Principal Type Schemes for Gradual Programs. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*.

Michael Greenberg. 2013. *Manifest Contracts.* Ph.D. Dissertation. University of Pennsylvania.

Ben Greenman and Matthias Felleisen. 2018. A Spectrum of Type Soundness and Performance. 2, ICFP (2018).

Ben Greenman and Zeina Migeed. 2018. On the Cost of Type-Tag Soundness. In *ACM SIGPLAN Workshop on Partial Evaluation and Program Manipulation (PEPM)*.

Arjun Guha, Claudiu Saftoiu, and Shriram Krishnamurthi. 2011. Typing Local Control and State Using Flow Analysis. In *European Symposium on Programming (ESOP)*.

Vincent J. Hellendoorn, Christian Bird, Earl T. Barr, and Miltiadis Allamanis. 2018. Deep Learning Type Inference. In *ACM SIGSOFT Symposium on the Foundations of Software Engineering (FSE)*.

Fritz Henglein. 1994. Dynamic typing: syntax and proof theory. *Science of Computer Programming* 22, 3 (1994), 197–230.

Fritz Henglein and Jakob Rehof. 1995. Safe polymorphic type inference for a dynamically typed language: Translating Scheme to ML. In *International Conference on Functional Programming Languages and Computer Architecture (FPCA)*.

David Herman, Aaron Tomb, and Cormac Flanagan. 2011. Space-efficient gradual typing. *Higher-Order and Symbolic Computation (HOSC)* 23, 2 (2011), 167–189.

Milod Kazerounian, Brianna M. Ren, and Jeffrey S. Foster. 2020. Sound, Heuristic Type Annotation Inference for Ruby. In *Dynamic Languages Symposium (DLS)*.

Andre Kuhlenschmidt, Deyaaeldeen Almahallawi, and Jeremy G. Siek. 2019. Toward Efficient Gradual Typing for Structural Types via Coercions. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*.

Rabee Sohail Malik, Jibesh Patra, and Michael Pradel. 2019. NL2Type: Inferring JavaScript Function Types from Natural Language Information. In *International Conference on Software Engineering (ICSE)*.

Zeina Migeed and Jens Palsberg. 2020. What is Decidable about Gradual Types? *Proceedings of the ACM on Programming Languages (PACMPL)* 4, POPL (2020).

Yusuke Miyazaki, Taro Sekiyama, and Atsushi Igarashi. 2019. Dynamic Type Inference for Gradual HindleyâĂŞMilner Typing. *Proceedings of the ACM on Programming Languages (PACMPL)* 3, POPL (2019).

Zvonimir Pavlinovic, Tim King, and Thomas Wies. 2014. Finding Minimum Type Error Sources. In *ACM SIGPLAN Conference on Object Oriented Programming, Systems, Languages and Applications (OOPSLA)*.

Luna Phipps-Costin, Carolyn Jane Anderson, Michael Greenberg, and Arjun Guha. 2021. Solver-based Gradual Type Migration. arXiv:2109.05049 [cs.PL]

Michael Pradel, Georgios Gousios, Jason Liu, and Satish Chandra. 2020. TypeWriter: Neural Type Prediction with Search-Based Validation. In *Joint Meeting of the European Software Engineering Conference (ESEC) and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (FSE)*.

Aseem Rastogi, Avik Chaudhuri, and Basil Hosmer. 2012. The Ins and Outs of Gradual Type Inference. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*.

Brianna M. Ren and Jeffrey S. Foster. 2016. Just-in-Time Static Type Checking for Dynamic Languages. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*.

Claudiu Saftoiu. 2010. *JSTrace: Run-time type discovery for JavaScript.* Master's thesis. Brown University.

Jeremy Siek, Peter Thiemann, and Philip Wadler. 2015a. Blame and Coercion: Together Again for the First Time. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*.

Jeremy Siek, Michael Vitousek, Matteo Cimini, and John Boyland. 2015b. Refined Criteria for Gradual Typing. In *Summit oN Advances in Programming Languages (SNAPL)*.

Jeremy G. Siek and Walid Taha. 2006. Gradual Typing for Functional Languages. In *Scheme and Functional Programming Workshop (SW)*.

Jeremy G. Siek and Manish Vachharajani. 2008. Gradual Typing with Unification-based Inference. In *Dynamic Languages Symposium (DLS)*.

Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, Sam Tobin-Hochstadt, and Ronald Garcia. 2015c. Monotonic References for Efficient Gradual Typing. In *European Symposium on Programming (ESOP)*.

Sam Tobin-Hochstadt and Matthias Felleisen. 2006. Interlanguage Migration: From Scripts to Programs. In *Dynamic Languages Symposium (DLS)*.

Sam Tobin Hochstadt and Matthias Felleisen. 2008. The Design and Implementation of Typed Scheme. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*.

Panagiotis Vekris, Benjamin Cosman, and Ranjit Jhala. 2015. Trust, but Verify: Two-Phase Typing for Dynamic Languages. In *European Conference on Object-Oriented Programming (ECOOP)*.

Michael M. Vitousek, Jeremy G. Siek, and Jim Baker. 2014. Design and Evaluation of Gradual Typing for Python. In *Dynamic Languages Symposium (DLS)*.

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 *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*.

Jiayi Wei, Maruth Goyal, Greg Durrett, and Isil Dillig. 2020. LambdaNet: Probabilistic Type Inference using Graph Neural Networks. In *International Conference on Learning Representations (ICLR)*.

Andrew K. Wright and Robert Cartwright. 1997. A Practical Soft Type System for Scheme. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 19, 1 (1997), 87–152.