

Solver-based Gradual Type Migration

LUNA PHIPPS-COSTIN, University of Massachusetts Amherst, United States

CAROLYN JANE ANDERSON, Wellesley College, United States

MICHAEL GREENBERG, Pomona College, 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 an extension of the gradually-typed lambda calculus. Unlike prior work, which relies on custom solvers, `TYPEWHICH` produces constraints that can be solved by 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.

1 INTRODUCTION

Gradually typed languages allows 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 [??]. 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 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* (GTL) [?], and limit ourselves to modifying existing type annotations, a single program may have many possible migrations. Existing approaches either produce a single migration [?????], or a menu of possible migrations without selection guidance [?]. How should we choose among the migrations produced by various approaches?

This paper argues that there is a fundamental tension between type migrations that produce precise or “informative” type annotations, and those that preserve the behavior of the original program. In fact, in many GTLC programs, making types more precise can introduce new dynamic errors. Making types more precise can also introduce static and dynamic errors at the (higher-order) boundary between migrated and unmigrated code, a serious concern when migrating a library or a fragment of a larger program. A general-purpose type migration tool would allow programmers to make an informed choice between multiple migrations depending on their context of use.

With these design constraints in mind, we present `TYPEWHICH`, a type migration tool that is novel in two key ways. First, unlike prior systems that rely on custom constraint solvers, `TYPEWHICH` generates constraints for an off-the-shelf MaxSMT solver [?]. 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 [?].

Second, using a general-purpose solver allows `TYPEWHICH` to support multiple migrations with different properties: the user can prioritize migrations with the most informative types, or migrations that maximize compatibility with unmigrated code, or something in between. 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 inferred program type 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.

This paper also presents the first comprehensive evaluation of different type migration approaches. We compare `TYPEWHICH` to four other type migration approaches using a two-part evaluation suite: a set of existing benchmarks by ?, 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 well in our evaluation. An advantage of `TYPEWHICH` over most existing work is that it does not reject any untyped programs. Finally, we apply `TYPEWHICH` to a suite of Grift programs from ?, and find that it reproduces all hand-written type annotations except in one case.

Contributions. Our key contributions are as follows:

- (1) We characterize the many goals of type migration, illustrate their inherent competition, and argue that type migration tools should allow users to make informed decisions about their own priorities (?? and ??).
- (2) We present the `TYPEWHICH` approach to type migration, which formulates constraints for an off-the-shelf MaxSMT solver (??). `TYPEWHICH` supports the GTLC and additional language features required to support the Grift gradually typed language (??).
- (3) We present a new set of type migration “challenge problems” that illustrate the strengths and weaknesses of different approaches to type migration (??).
- (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 ? and ?. Ours is the first publicly available implementation of ?.

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 [?] 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 small portion of a larger system, increasing precision can introduce type errors at the boundaries between migrated and unmigrated code [?]. 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 \star type introduces coercions [?]; these coercions can fail at run-time with a dynamic type error.

Type migration can introduce new static errors. ?? shows a function that uses its \star -typed argument first as a number and then as a function. Since \star 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 :  $\star$ ) =
2   let _ = x + 10 in
3   x ()

```

Fig. 1. Reachable error.

However, it is also possible for the crashing expression to be unreachable. ?? 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 :  $\star$ ) =
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.

?? 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 valid calls to c ill-typed. For example, $c \ 0$ evaluates to 42 before migration, but is ill-typed after migration.

```

1 int -> int int
2
3
4 let C (f :  $\star$ ) (x :  $\star$ ) =
5   if x > 0 then
6     1 + f x
7   else
8     42

```

Fig. 3. Context restriction.

?? 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 $\text{int} \rightarrow \text{int}$. However, D also returns f back to its caller, so this migration changes the return type of D from \star to $\text{int} \rightarrow \text{int}$. 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 ints .

```

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

<p>Base types $B := \text{int} \mid \text{bool}$</p> <p>Types and contexts $S, T := B$ Base type $\mid S \rightarrow T$ Function type $\mid \star$ Unknown type</p> <p>$\Gamma := \cdot \mid \Gamma, x : T$</p> <p>Constants $b := \text{true} \mid \text{false}$ Boolean literal $n := \dots$ Integer literal $c := b \mid n$</p> <p>Expressions $e := x$ Identifier $\mid c$ Literal $\mid \text{fun}(x : T).e$ Function $\mid e_1 e_2$ Application $\mid e_1 \times e_2$ Multiplication</p>	<p>Type Consistency $\boxed{T \sim T}$</p> $\frac{}{\star \sim T} \quad \frac{}{T \sim \star} \quad \frac{}{T \sim T} \quad \frac{}{B \sim B} \quad \frac{S_1 \sim S_2 \quad T_1 \sim T_2}{S_1 \rightarrow T_1 \sim S_2 \rightarrow T_2}$ <p>Typing Literals $\boxed{ty : c \rightarrow B}$</p> $ty(n) = \text{int} \quad ty(b) = \text{bool}$ <p>Typing $\boxed{\Gamma \vdash e : T}$</p> $\frac{\Gamma(x) = T}{\Gamma \vdash x : T} \quad \frac{}{\Gamma \vdash c : ty(c)} \quad \frac{\Gamma, x : S \vdash e : T}{\Gamma \vdash \text{fun}(x : S).e : S \rightarrow T}$ $\frac{\Gamma \vdash e_1 : \star \quad \Gamma \vdash e_2 : T}{\Gamma \vdash e_1 e_2 : \star} \quad \frac{\Gamma \vdash e_1 : S \rightarrow T \quad \Gamma \vdash e_2 : S' \quad S \sim S'}{\Gamma \vdash e_1 e_2 : T}$ $\frac{\Gamma \vdash e_1 : S \quad \Gamma \vdash e_2 : T \quad S \sim \text{int} \quad T \sim \text{int}}{\Gamma \vdash e_1 \times e_2 : \text{int}}$
---	---

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

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. ?? shows a program that runs without error: \mathbb{E} receives the identity function and applies it to two different types. However, since \mathbb{E} 's argument has type \star , which is compatible with all types, 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) [?], which is a core calculus for mixing typed and untyped code. We then present several definitions of type migration for the GTLC.

```

1 let E(id :  $\star$ ) =
2   id 2;
3   id true               $\boxed{\text{int}}$ 
4
5 E(fun (x :  $\star$ ) . x);

```

Fig. 5. Dynamic type error.

<p>Ground types $G := B \mid \text{fun}$</p> <p>Coercions $k := G?$ Untag $G!$ Tag $\text{wrap}(k_1, k_2)$ Wrap function $k_1; k_2$ Sequence id_T Identity</p> <p>Expressions $e := \dots \mid [k] e$ Apply coercion</p> <p>Untagged values $u := c \mid \text{fun}(x : T).e$</p> <p>Values $v := u \mid \text{box}(G, u)$</p> <p>Evaluation Contexts $E := [] \mid E e \mid v E \mid [k] E$</p> <p>Active Expressions $ae := (\text{fun}(x : T).e) v \mid [k] v$</p> <p>Evaluation $\boxed{\Gamma \vdash e \hookrightarrow e'}$ $(\text{fun}(x : T).e) v \hookrightarrow e[x/v]$ $[id] v \hookrightarrow v$ $[G!](u) \hookrightarrow \text{box}(G, u)$ $[G?](\text{box}(G, u)) \hookrightarrow u$ $[\text{wrap}(k_1, k_2)] v \hookrightarrow \text{fun}(x : \star).[k_2](v([k_1] x))$ $[k_1; k_2] v \hookrightarrow [k_2]([k_1] v)$ $ae \hookrightarrow e'$ $\hline E[ae] \hookrightarrow E[e']$</p>	<p>$\text{coerce}(T, T) = \text{id}_T$ $\text{coerce}(\star, b) = b?$ $\text{coerce}(b, \star) = b!$ $\text{coerce}(\star, \star \rightarrow \star) = \text{fun?}$ $\text{coerce}(\star \rightarrow \star, \star) = \text{fun!}$ $\text{coerce}(S_1 \rightarrow S_2, T_1 \rightarrow T_2) = \text{wrap}(\text{coerce}(T_1, S_1), \text{coerce}(S_2, T_2))$ $\text{coerce}(\star, T_1 \rightarrow T_2) = \text{fun?}; \text{wrap}(\text{coerce}(\star, T_1, \star), \text{coerce}(\star, T_2))$ $\text{coerce}(T_1 \rightarrow T_2, \star) = \text{wrap}(\text{coerce}(\star, T_1), \text{coerce}(T_2, \star)); \text{fun!}$ $\text{coerce}(S, T) = \text{coerce}(S, \star); \text{coerce}(\star, T)$</p> <p>Coercion Insertion $\boxed{\Gamma \vdash e \Rightarrow e', T}$</p> $\frac{\Gamma(x) = T}{\Gamma \vdash x \Rightarrow x, T} \quad \frac{}{\Gamma \vdash c \Rightarrow c, \text{ty}(c)}$ $\frac{\Gamma, x : S \vdash e \Rightarrow e', T}{\Gamma \vdash \text{fun}(x : S).e \Rightarrow \text{fun}(x : S).e', S \rightarrow T}$ $\frac{\Gamma \vdash e_1 \Rightarrow e'_1, S \rightarrow T \quad \Gamma \vdash e_2 \Rightarrow e'_2, S'}{\Gamma \vdash e_1 e_2 \Rightarrow e'_1([\text{coerce}(S', S)] e'_2), T}$ $\frac{\Gamma \vdash e_1 \Rightarrow e'_1, T \quad T \neq T_1 \rightarrow T_2 \quad \Gamma \vdash e_2 \Rightarrow e'_2, S}{\Gamma \vdash e_1 e_2 \Rightarrow ([\text{coerce}(T, \star \rightarrow \star)] e'_1) ([\text{coerce}(S, \star)] e'_2), \star}$
---	--

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

3.1 The Gradually Typed Lambda Calculus

The Gradually Typed Lambda Calculus (GTL) extends the typed lambda calculus with base types (integers and booleans) and the *unknown type* \star . ?? 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 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 ($S \rightarrow T$), 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 ??, where we discuss overloading.

$$\begin{array}{c}
\text{Type Precision} \quad \boxed{T \sqsubseteq T} \\
\frac{}{\star \sqsubseteq T} \quad \frac{}{T \sqsubseteq T} \quad \frac{S_1 \sqsubseteq S_2 \quad T_1 \sqsubseteq T_2}{S_1 \rightarrow T_1 \sqsubseteq S_2 \rightarrow T_2} \\
\text{Expression Precision} \quad \boxed{e \sqsubseteq e} \\
\frac{}{x \sqsubseteq x} \quad \frac{}{c \sqsubseteq c} \quad \frac{e_1 \sqsubseteq e'_1 \quad e_2 \sqsubseteq e'_2}{e_1 e_2 \sqsubseteq e'_1 e'_2} \quad \frac{T \sqsubseteq T' \quad e \sqsubseteq e'}{\text{fun}(x : T).e \sqsubseteq \text{fun}(x : T').e'}
\end{array}$$

Fig. 8. Type and expression precision.

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 $?$, 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 $?$) 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 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 $\text{int} \rightarrow \text{int}$. To do so, we apply f to a coercion as follows:

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

This coercion 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 ($\text{box}(G, u)$) 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 ($S \not\sim T$). 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 ($??$), 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.

¹The two approaches are inter-translatable $[[?]]$ with full abstraction $[[?]]$.

? 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 ??, improving type precision is one of several competing goals for type migration. Another important goal is to preserve the behavior of the original 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 of value-restricted type migration:

Definition 3.2 (Value-restricted Type Migration). Given $\vdash e : T$ and $\vdash e' : T'$, e' is a *restricted 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, this definition 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 $\star \rightarrow \star$, this definition allows a type migration that changes its type to $\text{int} \rightarrow \text{int}$, 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 is used. We define a well-typed program context C 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; or c) both $C[e]$ and $C[e']$ do not terminate.

At the limit, the context's expected type S could be \star , in which case the definition is essentially equivalent to that of ?, Theorem 3.22. However, this is a very strong requirement that rules out many informative migrations (??). 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.

<p>Types</p> $T := \dots$ $\alpha, \beta, \gamma, \delta$ Type metavariables <p>Coercions</p> $k := \dots$ $\text{coerce}(T_1, T_2)$ Coercion from T_1 to T_2 <p>Type Representation</p> $\begin{array}{l} 1 \text{ (declare-datatypes ()} \\ 2 \text{ ((Typ (star) (int) (bool)} \\ 3 \text{ (arr (in Typ) (out Typ))))}) \end{array}$	<p>Constraints</p> $\phi := T_1 = T_2$ Type equality w Boolean variable (weight) $\phi_1 \wedge \phi_2$ Conjunction $\phi_1 \vee \phi_2$ Disjunction $\neg \phi$ Negation <p>Constraint Metafunctions</p> $\text{ground} \in T \rightarrow \phi$ $\text{ground}(T) = T \in B \vee T = \star \rightarrow \star$
--	---

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

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 [?]. (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, compatibility with untyped code, or other properties. Moreover, the **TYPEWHICH** migration algorithm handles these different scenarios in a uniform, type-directed way. This section presents **TYPEWHICH**'s type migration algorithm for the core GTLC. **??** 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 (??):

- (1) We extend types with type metavariables (α).
- (2) We introduce a new coercion, $\text{coerce}(S, T)$, which represents a suspended call to the coerce metafunction (??). The type arguments to 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 [?]. In addition to the usual propositional connectives, our constraints involve equalities between types ($T_1 = T_2$), predicates over types (e.g., to check if a type is an arrow type), 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 $\alpha \rightarrow \text{int} = \beta$ 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 $\text{SUBST}(\sigma, \beta)$ 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,

$$\begin{array}{c}
 \boxed{\Gamma \vdash e \Rightarrow e, T, \phi} \\
 \text{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{\text{coerce}}(\Gamma(x), \alpha)]x, \alpha, \phi} \quad \text{CONST} \frac{\phi = (\alpha = \text{ty}(c) \wedge w) \vee (\alpha = \star \wedge \neg w) \quad \alpha, w \text{ is fresh}}{\Gamma \vdash c \Rightarrow [\underline{\text{coerce}}(\text{ty}(c), \alpha)]c, \alpha, \phi} \\
 \\
 \text{FUN} \frac{\Gamma, x : \alpha \vdash e \Rightarrow e', T, \phi_1 \quad \beta, w \text{ fresh} \quad \phi_2 = (\beta = \alpha \rightarrow T \wedge w) \vee (\beta = \star \wedge \text{ground}(\alpha \rightarrow T) \wedge \neg w)}{\Gamma \vdash \text{fun}(x : \alpha).e \Rightarrow [\underline{\text{coerce}}(\alpha \rightarrow T, \beta)]\text{fun}(x : \alpha).e', \beta, \phi_1 \wedge \phi_2} \\
 \\
 \text{APP} \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 \alpha, \beta, \gamma, w_1, \text{ and } w_2 \text{ are fresh} \quad \phi_3 = (T_1 = \alpha \rightarrow \beta \wedge w_1) \vee (T_1 = \alpha = \beta = \star \wedge \neg w_1) \quad \phi_4 = (T_2 = \alpha) \quad \phi_5 = (\beta = \gamma \wedge w_2) \vee (\gamma = \star \wedge \neg w_2)}{\Gamma \vdash e_1 e_2 \Rightarrow [\underline{\text{coerce}}(\beta, \gamma)]([\underline{\text{coerce}}(T_1, \alpha \rightarrow \beta)]e'_1) e'_2, \gamma, \phi_1 \wedge \phi_2 \wedge \phi_3 \wedge \phi_4 \wedge \phi_5} \\
 \\
 \text{MUL} \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 w_1, w_2, \text{ and } w_3 \text{ are fresh} \quad \phi_3 = (T_1 = \text{int} \wedge w_1) \vee (T_1 = \star \wedge \neg w_1) \quad \phi_4 = (T_2 = \text{int} \wedge w_2) \vee (T_2 = \star \wedge \neg w_2) \quad \phi_5 = (\alpha = \text{int} \wedge w_3) \vee (\alpha = \star \wedge \neg w_3)}{\Gamma \vdash e_1 \times e_2 \Rightarrow [\underline{\text{coerce}}(\text{int}, \alpha)]([\underline{\text{coerce}}(T_1, \text{int})]e'_1 \times [\underline{\text{coerce}}(T_2, \text{int})]e'_2), \alpha, \phi_1 \wedge \phi_2 \wedge \phi_3 \wedge \phi_4 \wedge \phi_5}
 \end{array}$$

Fig. 10. Constraint generation for GTLC

so there are several possible models: $\sigma = \{\alpha \mapsto \text{int}, \beta \mapsto \alpha \rightarrow \text{int}\}$ is one, as is $\sigma' = \{\alpha \mapsto \star, \dots\}$. We have $\text{SUBST}(\sigma, \beta) = \text{int} \rightarrow \text{int}$, while $\text{SUBST}(\sigma', \beta) = \star \rightarrow \text{int}$.

Finally, for succinctness, we define $\text{ground}(T)$, 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 ($\star \rightarrow \star$). ground extends the language with additional types and augments the definition of ground .

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 ??, 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 that we leave for future work (but see ?).

Constraint Generation Rules. Constraint generation is syntax directed (?). 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 ??).²

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 (??), 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 $\text{true} \times 1$). 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 ($\alpha \rightarrow T$) or the \star type. In the latter case, we also constrain the type of the function to be the ground type ($\star \rightarrow \star$). We use a weight w to prefer the former case without rejecting expressions like $1 \times (\text{fun}(x : \star).x)$.

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, $\star \rightarrow \star$ and w 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:³

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

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:

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

It is important to 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 $id \ \text{true}$, 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 ($\text{int} \rightarrow \text{int}$) is not a ground type, which APP also requires. Thus our constraint generation algorithm rules out this migration.

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):

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

²We 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 (??).

³This is a variation of the example in ??.

```

1: ▷ The only annotations in  $e$  are  $\star$ 
2: function PRECISEMIGRATE( $e$ )
3:    $e_1 \leftarrow \text{INTRODUCEMETAVARS}(e)$ 
4:    $\cdot \vdash e \Rightarrow e', T_1, \phi$ 
5:   for  $\alpha \in \phi$  do
6:     (declare-const  $\alpha$  Typ)
7:   for  $w \in \phi$  do
8:     (declare-const  $w$  Bool)
9:     (assert-soft  $w$  1)
10:  (check-sat  $\phi$ )
11:   $\sigma \leftarrow (\text{get-model})$ 
12:  return SUBST( $\sigma, e'$ )
    
```

▷ Replace every \star with a fresh α
 ▷ Generate constraints and objectives
 ▷ The set of type metavariables in ϕ
 ▷ The set of weight variables in ϕ
 ▷ Model mapping type metavariables to types
 ▷ Migrated program without explicit coercions

Fig. 11. Precise Type Migration.

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 $\star \rightarrow \star$ and leaves n and x at type \star :

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

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 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 $\text{SUBST}(\sigma, \Gamma) \vdash \text{SUBST}(\sigma, e') : \text{SUBST}(\sigma, T)$.*

PROOF. By induction on the coercion insertion judgment (see ?? for more details). □

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 ($\text{coerce}(T, T)$) 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 ?. 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 program as an input to our algorithm:

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

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

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

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 program F_2 is not substitutable for F_1 :

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

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)
$\underline{F_1} \ (\text{fun}(x : \star).0) \ (\text{fun}(k : \star).k \ \text{true})$	$\underline{F_2} \ (\text{fun}(x : \star).0) \ (\text{fun}(k : \star).k \ \text{true})$

Both programs above are well-typed. However, the static error from the previous example is now a dynamic error. As we argued in ??, 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 :

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

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)
$\underline{F_1} \ (\text{fun}(x : \star).0) \ (\text{fun}(k : \star).k \ \text{true})$	$\underline{F_3} \ (\text{fun}(x : \star).0) \ (\text{fun}(k : \star).k \ \text{true})$

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 .

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

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

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

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.

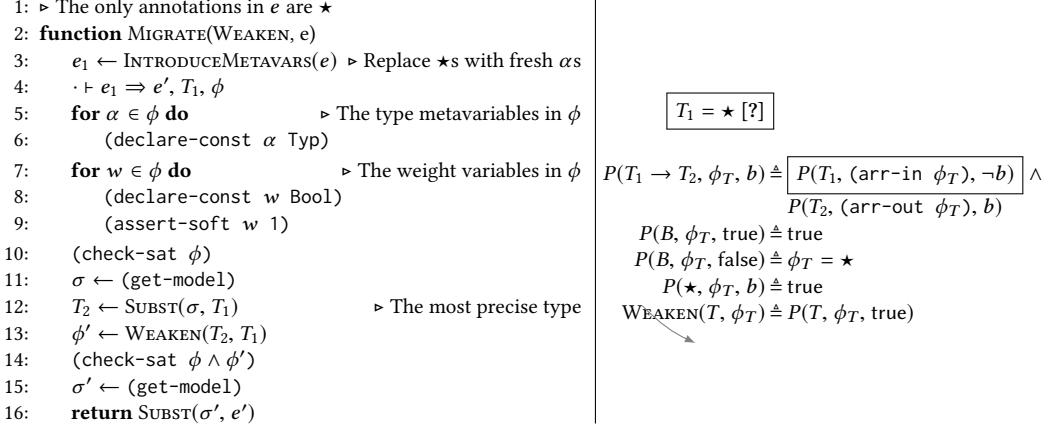


Fig. 12. The Type Migration Algorithm.

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: $[\text{coerce}(T, S)]e$. 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 .

Weakening Migrations. `TYPEWHICH` employs a two-step approach to type migration. We first generate constraints and calculate the most precise type possible, as described earlier (lines ??–?? of ??; identical to ??). We then identify all base types in negative position (following ??), and add new constraints that force them to be \star to ensure contextual compatibility. We apply a function `WEAKEN` to the output of the first-pass constraint generation (T_1 , which has metavariables) to add these additional constraints.

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 ??). This produces a new model (line ??), 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. We want to discover a type skeleton that is as precise as possible; without this stage, 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 $\text{int} \rightarrow \text{int}$. Since this type only allows the context to provide int -arguments to e , we might conclude that a better type for e is $\star \rightarrow \text{int}$. But this may be impossible: for instance, if e is the identity function, its argument

Ground types	Base Types	Constraint Metafunctions
$G := \dots \mid \text{ref}$	$B := \dots \mid \text{unit}$	$\text{ground} \in T \rightarrow \phi$
Constants	Types	$\text{ground}(T) = T \in B \vee T = \star \rightarrow \star \vee T = \text{ref } \star$
$b := \dots \mid \text{unit}$	$T := \dots \mid \text{ref } T$	Type Representation
Expressions		1 (declare-datatypes ()
$e := \dots$		2 ((Typ (star) (int) (bool)
$\text{ref } e$ Create cell		3 (ref (to Typ))
$!e$ Read cell		4 (arr (in Typ) (out Typ))))
$e_1 := e_2$ Write cell		

$$\text{IF} \frac{\Gamma \vdash e_1 \Rightarrow e'_1, T_1, \phi_1 \quad \Gamma \vdash e_2 \Rightarrow e'_2, T_2, \phi_2 \quad \Gamma \vdash e_3 \Rightarrow e'_3, T_3, \phi_3 \quad w_1, w_2, \alpha \text{ fresh}}{\Gamma \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \Rightarrow \text{if } [\text{coerce}(T_1, \text{bool})]e'_1 \text{ then } [\text{coerce}(T_2, \alpha)]e'_2 \text{ else } [\text{coerce}(T_3, \alpha)]e'_3, \alpha, \phi_1 \wedge \phi_2 \wedge \phi_3 \wedge \phi_4}$$

$$\text{ADD} \frac{\Gamma \vdash e_1 \Rightarrow T_1, e'_1, \phi_1 \quad \Gamma \vdash e_2 \Rightarrow T_2, e'_2, \phi_2 \quad w, \alpha \text{ are fresh}}{\Gamma \vdash e_1 + e_2 \Rightarrow ([\text{coerce}(T_1, \alpha)]e'_1) + ([\text{coerce}(T_2, \alpha)]e'_2), \alpha, \phi_1 \wedge \phi_2 \wedge \phi_3}$$

$$\text{REF} \frac{\Gamma \vdash e \Rightarrow T, e', \phi_1 \quad \alpha \text{ and } w \text{ are fresh}}{\Gamma \vdash \text{ref } e \Rightarrow [\text{coerce}(\text{ref } T, \text{ref } \alpha)]\text{ref } e', \alpha, \phi_1 \wedge \phi_2}$$

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

$$\text{SETREF} \frac{\Gamma \vdash e_1 \Rightarrow T_1, e'_1, \phi_1 \quad \Gamma \vdash e_2 \Rightarrow T_2, e'_2, \phi_2 \quad \alpha \text{ and } w \text{ are fresh}}{\Gamma \vdash e_1 := e_2 \Rightarrow ([\text{coerce}(T_1, \text{ref } \alpha)]e'_1) := [\text{coerce}(T_2, \alpha)]e'_2, \text{Unit}, \phi_1 \wedge \phi_2 \wedge \phi_3}$$

Fig. 13. Extensions to the GTLC.

and result types must be the same. On the other hand, if the body of e is a multiplication, then making the input type \star 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 \star , but preserves arrow types in the input. An alternative is to turn all input types to \star to maximize compatibility, similar to `?`. 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 safety in type migration. Our first-pass discovers the most precise types that we can; 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 $\vdash e \Rightarrow e', T, \phi$ such that ϕ is satisfiable in some model σ .*

PROOF. We prove that a fully dynamic model σ exists (??) and then show that such models are stable under weakening, i.e., they are still satisfiable (?? and ??). \square

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 (??). The GTLC-based languages (e.g., ?) 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 ?? 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 (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 ?? 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 [???]. 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 (??). 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 ??.

6 EVALUATION

This section presents the first comprehensive comparison of several type migration algorithms from the literature (along with `TYPEWHICH`). We compare five type migration tools on a two-part suite of 22 programs: the benchmarks from ? and a new suite of “challenge programs” that we have designed to illustrate the strengths and weaknesses of various approaches. We also evaluate `TYPEWHICH` using the Grift benchmarks from ?, to show that `TYPEWHICH` can reconstruct hand-written type annotations in Grift.

To facilitate high-level understanding of the results, we first discuss summary metrics for evaluating automated type migration tools.

6.1 How Should Type Migration Tools Be Evaluated?

As we have argued in ??, the type migration problem involves inherent trade-offs between different goals. For this reason, we avoid choosing a single evaluation metric, since this would favor one goal of type migration above the others. For instance, using the total number of type annotations improved is a good metric for type precision, but reporting only precision would obscure the fact that not all type refinements 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
$\text{fun}(f : \star).\text{fun}(x : \star).fxx$	$\text{fun}(f : \text{int} \rightarrow \text{bool} \rightarrow \star).\text{fun}(x : \star).fxx$

The migrated program has precise 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?*

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).

6.2 Type Migration Systems

We evaluate the performance of the following tools:

- (1) **TYPEWHICH**: our tool, which we run in two modes: (a) to produce the most precise migration that we can (**TYPEWHICH-P**), and (b) to produce a migration that is compatible with unmigrated code (**TYPEWHICH-C**).
- (2) **GTUBI**: *gradual typing with unification-based inference* [?] is the earliest work on gradual type migration. It does not introduce coercions that may fail.
- (3) **INSANDOUTS**: our implementation of the algorithm in ?. The algorithm is designed to not introduce coercions that may fail, and to produce a migration that is compatible with arbitrary unmigrated code.
- (4) **MAXMIGRATE**: ? 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 five as in the paper). A single program may have several maximal migrations; we take the first migration produced by the tool. We halt with no output if no migration is found.
- (5) **MGT**: our implementation of the algorithm in ? for migrating untyped or partially typed programs. We start from untyped code (all functions annotated with \star). We take the first migration produced by the tool.

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	$(\text{fun}(b : \star).b (\text{fun}(c : \star).5 \ 5) (\text{fun}(d : \star).0)) (\text{fun}(t : \star).\text{fun}(f : \star).f)$
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.

6.3 Gradual Type Migration Benchmarks

Our benchmark suite consists of two parts: a suite of benchmarks from ?, 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 ???. We describe the ten programs below. Although TYPEWHICH supports several extensions to the GTLC (???), we largely avoid their use in the challenge suite for compatibility with as many approaches as possible.

- (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 ??): 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 ??): 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 $f (f \ \text{true})$. The inner application suggests that f 's argument must be bool. However, that would rule out $\text{fun}(x : \star).0$ 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 ???) 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

Tool	1 + true rejected				
	Rejected Total Programs	New Dynamic Errors Remaining Programs	Unusable Functions Remaining Programs	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	29 / 58
MAXMIGRATE	5 / 22	3 / 17	3 / 14	3 / 11	6 / 18
TYPEWHICH-C	0 / 22	0 / 22	0 / 22	0 / 22	39 / 58
TYPEWHICH-P	0 / 22	0 / 22	0 / 22	4 / 22	25 / 58

Fig. 15. Summary of Type Migration Results. Above each column, we show an example of the kind of migrated program we count in that column.

migration must use \star rather than the more precise type. In contrast, the type of f in DOUBLE-F is $\text{bool} \rightarrow \text{bool}$. However, f can have other types in the GTLC.

6.4 Evaluation Results

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 using the evaluation scheme proposed in ??.

?? summarizes each tool’s performance on the full benchmark suite. 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 four functions. However, even in this mode, the remaining 18 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.

All other tools reject some programs. 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. MGT statically rejects some programs and restricts some functions, but it does not introduce run-time errors. INSANDOUTS rejects two programs.⁴ 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 \star serves as an “escape hatch” for programs that cannot be given more precise types. Our suite includes programs that *must* have some \star s, so every tool will have to leave some \star 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

⁴These are two programs from the ? benchmarks. From correspondence with the authors of ?, our implementation seems faithful to the presentation in the paper, and the original implementation for Adobe ActionScript is no longer accessible.

contrast, the oldest tool—GTUBI—only leaves a small fraction of annotations unimproved, but it statically rejects the majority of programs.

6.4.1 Challenge Set Results. We now examine performance on the challenge set in more detail. ?? 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 and MAXMIGRATE produce the most precise and informative result: they show that x is a boolean next to $x + 100$, which helps locate the error in the program.
- (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. TYPEWHICH produces a needless `int` annotation that restricts the program. MAXMIGRATE produces `int \rightarrow 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 $\star \rightarrow \star$; no tool produces this type.
- (8) OUTFLOWS: INSANDOUTS produces the best result. This program requires x to have two different types and thus crashes. Because the function receives an integer for x , ? gives 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 constraints 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 $([int?]x) + 1$. If we were migrating a larger program that had this function as a sub-expression, and this function were actually applied to different x arguments with different types, it would be \star .

6.4.2 ? Benchmarks. ? compare their maximal migration tool to the type migration tool in ?. We extend the comparison to include TYPEWHICH, GTUBI, and INSANDOUTS. The complete results are in ??, and we include all of these benchmarks in our summary (??).

FARG-MISMATCH	MGT MAXMIGRATE INSANDOUTS TYPEWHICH-P GTUBI	(fun f : ★.f true)(fun x : int.x + 100) identical to INSANDOUTS (fun f : bool → int.f true)(fun x : bool.x + 100) (fun f : ★ → int.f true)(fun x : ★.x + 100) constraint solving error
RANK2-POLY-ID	MGT MAXMIGRATE INSANDOUTS TYPEWHICH-P GTUBI	no improvement (fun i : ★ → ★.(fun a : int.i true)(i5))(fun x : bool.x) identical to TYPEWHICH (fun i : ★ → ★.(fun a : ★.i true)(i 5))(fun x : ★.x) constraint solving error
UNREACHABLE-ERR	MGT MAXMIGRATE INSANDOUTS TYPEWHICH-P GTUBI	less precise than TYPEWHICH-P No maximal migration (fun b : (α → α) → (★ → int) → ★ → int.b(fun c : ★.5 5)(fun d : ★.0)) (fun t : α → α.fun f : ★ → int.f) (fun b : (int → ★) → (int → int) → int → int.b(fun c : int. ([int!]5) 5)(fun d : int.0))(fun t : int → ★.fun f : int → int.f) constraint solving error
F-IN-F-OUT	MGT MAXMIGRATE INSANDOUTS TYPEWHICH-P GTUBI	identical to TYPEWHICH (fun f : int → ★.(fun y : bool.f)(f 5))(fun x : int.10 + x) (fun f : ★ → int.(fun y : int.f)(f 5))(fun x : ★.10 + x) (fun f : int → int.(fun y : int.f)(f 5))(fun x : int.10 + x) identical to TYPEWHICH
ORDER3-FUN	MGT MAXMIGRATE INSANDOUTS TYPEWHICH-P GTUBI	fun f : (★ → ★) → ★.fun x : ★ → ★.x(f x) fun f : int → int.fun x : ★.x(f x) no improvement fun f : (★ → int) → ★.fun x : ★ → int.x(f x) fun f : (α → β) → α.fun x : α → β.x(f x)
ORDER3-INTFUN	MGT MAXMIGRATE INSANDOUTS TYPEWHICH-P GTUBI	identical to TYPEWHICH fun f : int → int → int.fun g : ★.fg(g10 + 1) no improvement fun f : (int → int) → int → ★.fun g : int → int.fg(g10 + 1) fun f : (int → int) → int → α.fun g : int → int.fg(g10 + 1)
DOUBLE-F	MGT MAXMIGRATE INSANDOUTS TYPEWHICH-P GTUBI	identical to TYPEWHICH fun f : ★ → int.f(f True) no improvement fun f : bool → bool.f(f true) identical to TYPEWHICH
OUTFLOWS	MGT MAXMIGRATE INSANDOUTS TYPEWHICH-P GTUBI	no improvement (fun x : ★.x 5 + x) 5 (fun x : int.x 5 + x) 5 identical to MAXMIGRATE constraint solving error
PRECISION-RELATION	MGT MAXMIGRATE INSANDOUTS TYPEWHICH-P GTUBI	(fun f : ★.(f true) + ((fun g : int → ★.g 5)f))(fun x : ★.5) (fun f : bool → int.(f true) + ((fun g : ★ → int.g 5)f))(fun x : bool.5) (fun f : ★ → int.(f true) + ((fun g : int → int.g 5)f))(fun x : ★.5) (fun f : ★ → int.(f true) + ((fun g : ★ → int.g 5)f))(fun x : ★.5) constraint solving error
IF-TAG	MGT MAXMIGRATE INSANDOUTS TYPEWHICH-P GTUBI	Identical to TYPEWHICH-P conditionals unsupported fun tag : ★.fun x : ★.if tag then x + 1 else if x then 1 else 0 fun tag : bool.fun x : bool.if tag then ([int!]x) + 1 else if x then 1 else 0 conditionals unsupported

Fig. 16. Migrations of the challenge set with TYPEWHICH in precise mode.

6.4.3 *Summary.* Our type migration challenge suite is designed to highlight the strengths and weaknesses of different algorithms. As discussed in ??, 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.

6.5 Grift Performance Benchmarks

? 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** defines a stream library, and the hand-typed version gives streams an equirecursive type. `TYPEWHICH` improves some types, but it cannot improve the annotations on the stream library. `TYPEWHICH` infers \star rather than the `Tuple Int Dyn` migration shown below:

```
(define (stream-rest [st : (Tuple Int Dyn)])
  : (Tuple Int Dyn)
  ((tuple-proj st 1)))
```

The projection from the stream has type \star , but the function expects to return a tuple. However, our constraints will only insert a coercion from \star to a more precise type at an elimination form, so `TYPEWHICH` will not produce this migration.

6.6 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, implementation of the migration algorithm from ? and ?, 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 ???. ? substitutes metavariables that appear in type annotations with concrete types, using a variation on unification. ? builds a type inference system for ActionScript. Their system ensures that inference never fails and produces types that are compatible with all untyped contexts. ? uses variational typing to heuristically tame the exponential search space of types [?]. ? 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 that sets up gradual type migration for 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 ? and ?.

? introduces the theory of coercions that we use; ? 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 ??).

? extend ?’s work to infer principal types. Since we focus on monomorphic types, we do not directly compare against their algorithm. ? build on ?’s work, discussing the coherence issues what we point out in ??: types induce run-time checks that can affect program behavior. However, while we migrate all programs, ? use *dynamic type inference* to discover type inconsistencies and report them as run-time errors. ? 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, ? extend their previous work [?] 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.

?, ?, ?, and ? 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.

? 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, ? 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.

? 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 [?] 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. ?, 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: ? map program points to sets of abstract values, rather than types.

Alternative approaches to type migration and type inference consider sources of evidence beyond the program to be migrated. This includes work that applies supervised machine learning techniques to generate type annotations, such as [???]. Other lines of work use run-time profiling to guide type inference [???], or use programmer-supplied heuristics to guide type inference [??].

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.

Our paper also contributes to the evaluation of type migration tools. We define a multi-stage evaluation process that takes into account 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 between the many approaches to gradual type migration.

A CONSTRAINT GENERATION FOR ADDITIONAL EXPRESSIONS

$$\text{SEQUENCE} \frac{\Gamma \vdash e_1 \Rightarrow T_1, \phi_1 \quad \Gamma \vdash e_2 \Rightarrow T_2, \phi_2}{\Gamma \vdash e_1; e_2 \Rightarrow e'_1; e'_2, T_2, \phi_1 \wedge \phi_2}$$

$$\text{LET} \frac{\Gamma \vdash e_1 \Rightarrow T_1, \phi_1 \quad \Gamma, x : T_1 \vdash e_2 \Rightarrow T_2, \phi_2}{\Gamma \vdash \text{let } x = e_1 \text{ in } e_2 \Rightarrow \text{let } x = e'_1 \text{ in } e'_2, T_2, \phi_1 \wedge \phi_2}$$

$$\text{FIX} \frac{\Gamma \vdash e \Rightarrow T_1, \phi_1 \quad w \text{ and } \alpha \text{ are fresh} \quad \phi_2 = (T_1 = \alpha \wedge w) \vee (T_1 = \star \wedge \alpha = \star \rightarrow \star \wedge \neg w)}{\Gamma \vdash \text{fix } f : \alpha. e \Rightarrow \text{fix } f : \alpha. [\text{coerce}(T_1, \alpha)]e', \alpha, \phi_1 \wedge \phi_2}$$

$$\text{PAIR} \frac{\Gamma \vdash e_1 \Rightarrow T_1, \phi_1 \quad \Gamma \vdash e_2 \Rightarrow T_2, \phi_2 \quad \alpha \text{ and } w \text{ are fresh} \quad \phi_3 = (\alpha = \text{pair}(T_1, T_2) \wedge w) \vee (\alpha = \star \wedge \text{ground}(\text{pair}(T_1, T_2)) \wedge \neg w)}{\Gamma \vdash \text{pair}(e_1, e_2) \Rightarrow [\text{coerce}(\text{pair}(T_1, T_2), \alpha)]\text{pair}(e'_1, e'_2), \alpha, \phi_1 \wedge \phi_2 \wedge \phi_3}$$

$$\text{FIRST} \frac{\Gamma \vdash e \Rightarrow T, \phi_1 \quad \alpha, \beta, \text{ and } w \text{ are fresh} \quad \phi_2 = ((T = \text{pair}(\alpha, \beta) \wedge w) \vee (T = \alpha = \star \wedge \neg w))}{\Gamma \vdash \text{first}(e) \Rightarrow \text{first}([\text{coerce}(T, \text{pair}(\alpha, \beta))])e', \alpha, \phi_1 \wedge \phi_2}$$

$$\text{SECOND} \frac{\Gamma \vdash e \Rightarrow T, \phi_1 \quad \alpha, \beta, \text{ and } w \text{ are fresh} \quad \phi_2 = ((T = \text{pair}(\alpha, \beta) \wedge w) \vee (T = \beta = \star \wedge \neg w))}{\Gamma \vdash \text{second}(e) \Rightarrow \text{second}([\text{coerce}(T, \text{pair}(\alpha, \beta))])e', \beta, \phi_1 \wedge \phi_2}$$

$$\text{VECTOR} \frac{\Gamma \vdash e_1 \Rightarrow T_1, \phi_1 \quad \Gamma \vdash e_2 \Rightarrow T_2, \phi_2 \quad \alpha, w_1, \text{ and } w_2 \text{ are fresh} \quad \phi_3 = (\alpha = \text{vec}(T_1) \wedge w_1) \vee (\alpha = \star \wedge \text{ground}(\text{vector}(T_1)) \wedge \neg w_1) \quad \phi_4 = (T_2 = \text{int} \wedge w_2) \vee (T_2 = \star \wedge \neg w_2)}{\Gamma \vdash \text{vec}(e_1, e_2) \Rightarrow [\text{coerce}(\text{vec}(T_1), \alpha)]\text{vec}(e'_1, [\text{coerce}(T_2, \text{int})]e'_2), \alpha, \phi_1 \wedge \phi_2 \wedge \phi_3 \wedge \phi_4}$$

$$\text{VECGET} \frac{\Gamma \vdash e_1 \Rightarrow T_1, \phi_1 \quad \Gamma \vdash e_2 \Rightarrow T_2, \phi_2 \quad \alpha, w_1 \text{ and } w_2 \text{ are fresh} \quad \phi_3 = (T_2 = \text{int} \wedge w_1) \vee (T_2 = \star \wedge \neg w_1) \quad \phi_4 = (T_1 = \alpha = \star \wedge \neg w_2) \vee (T_1 = \text{vector}(\alpha) \wedge w_2)}{\Gamma \vdash \text{VecGet}(e_1, e_2) \Rightarrow \text{VecGet}([\text{coerce}(T_1, \text{vector}(\alpha))])e'_1, [\text{coerce}(T_2, \text{int})]e'_2, \alpha, \phi_1 \wedge \phi_2 \wedge \phi_3 \wedge \phi_4}$$

$$\text{VECSSET} \frac{\Gamma \vdash e_1 \Rightarrow T_1, \phi_1 \quad \Gamma \vdash e_2 \Rightarrow T_2, \phi_2 \quad \Gamma \vdash e_3 \Rightarrow T_3, \phi_3 \quad \alpha, w_1 \text{ and } w_2 \text{ are fresh} \quad \phi_4 = (T_2 = \text{int} \wedge w_1) \vee (T_2 = \star \wedge \neg w_1) \quad \phi_5 = (T_1 = \text{vector}(\alpha) \wedge T_2 = \alpha \wedge w_2) \vee (\alpha = \star \wedge \text{ground}(T_1) \wedge \neg w_2)}{\Gamma \vdash \text{VecSet}(e_1, e_2, e_3) \Rightarrow \text{VecSet}([\text{coerce}(T_1, \text{vector}(\alpha))])e'_1, [\text{coerce}(T_2, \alpha)]e'_2, [\text{coerce}(T_3, \text{int})]e'_3, \text{vector}(\alpha), \phi_1 \wedge \phi_2 \wedge \phi_3 \wedge \phi_4 \wedge \phi_5}$$

$$\text{LENGTH} \frac{\Gamma \vdash e_1 \Rightarrow T_1, \phi_1 \quad \Gamma \vdash e_2 \Rightarrow T_2, \phi_2 \quad \alpha, w_1 \text{ and } w_2 \text{ are fresh} \quad \phi_3 = (T_2 = \text{int} \wedge w_1) \vee (T_2 = \star \wedge \neg w_1) \quad \phi_4 = (T_1 = \text{vector}(\alpha) \wedge w_2) \vee (\alpha = \star \wedge \neg w_2)}{\Gamma \vdash \text{Length}(e_1, e_2) \Rightarrow \text{Length}([\text{coerce}(T_1, \text{vector}(\alpha))])e'_1, [\text{coerce}(T_2, \text{int})]e'_2, \text{int}, \phi_1 \wedge \phi_2 \wedge \phi_3 \wedge \phi_4}$$

B TYPING THE LANGUAGE WITH EXPLICIT COERCIONS

The rules below define type-checking for the intermediate language of GTLC, where all coercions are explicit.

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

$$\text{T-ID} \frac{\Gamma(x) = T}{\Gamma \vdash x : T}$$

$$\text{T-CONST} \frac{}{\Gamma \vdash c : \text{ty}(c)}$$

$$\text{T-FUN} \frac{\Gamma, x : S \vdash e : T}{\Gamma \vdash \text{fun}(x : S).e : S \rightarrow T}$$

$$\text{T-APP} \frac{\Gamma \vdash e_1 : T_1 \rightarrow T_2 \quad \Gamma \vdash e_2 : T_1}{\Gamma \vdash e_1 e_2 : T_2}$$

$$\text{T-MUL} \frac{\Gamma \vdash e_1 : \text{int} \quad \Gamma \vdash e_2 : \text{int}}{\Gamma \vdash e_1 \times e_2 : \text{int}}$$

$$\text{T-COERCE} \frac{\vdash k : T_1 \rightarrow T_2 \quad \Gamma \vdash e : T_1}{\Gamma \vdash [k]e : T_2}$$

$$\boxed{\vdash k : T}$$

$$\text{TC-TAG-FUN} \frac{}{\vdash \text{fun}! : (\star \rightarrow \star) \rightarrow \star}$$

$$\text{TC-CHK-FUN} \frac{}{\vdash \text{fun}? : \star \rightarrow (\star \rightarrow \star)}$$

$$\text{TC-TAG-INT} \frac{}{\vdash \text{int}? : \star \rightarrow \text{int}}$$

$$\text{TC-CHK-INT} \frac{}{\vdash \text{int}! : \text{int} \rightarrow \star}$$

$$\text{TC-SEQ} \frac{\vdash k_1 : T_1 \rightarrow T_2 \quad \vdash k_2 : T_2 \rightarrow T_3}{\vdash k_1; k_2 : T_1 \rightarrow T_3}$$

$$\text{TC-ID} \frac{}{\vdash \text{id}_T : T \rightarrow T}$$

$$\text{TC-WRAP} \frac{\vdash k_1 : T_1 \rightarrow S_1 \quad \vdash k_2 : S_2 \rightarrow T_2}{\vdash \text{wrap}(k_1, k_2) : (S_1 \rightarrow S_2) \rightarrow (T_1 \rightarrow T_2)}$$

C SOUNDNESS OF FOUND MODELS

Here we prove that if coercion insertion has a satisfiable model, it induces a well typed coercion term.

LEMMA C.1 (COERCIONS ARE WELL TYPED). $\vdash \text{coerce}(S, T) : S \rightarrow T$ using the definition of *coerce* from Figure ??.

PROOF. By induction on the sum of the sizes of the two coercions, with cases drawn from the function. Let \star have size 1 and `int` have size 2.

($S = T$) We have $\vdash \text{id}_S : S \rightarrow T$ by TC-ID.

($S = \star, T = \text{int}$) We have $\vdash \text{int}? : \star \rightarrow \text{int}$ by TC-CHK-INT.

($S = \text{int}, T = \star$) We have $\vdash \text{int}! : \text{int} \rightarrow \star$ by TC-TAG-INT.

($S = \star, T = \star \rightarrow \star$) We have $\vdash \text{fun}? : \star \rightarrow (\star \rightarrow \star)$ by TC-CHK-FUN.

($S = \star \rightarrow \star, T = \star$) We have $\vdash \text{fun}! : (\star \rightarrow \star) \rightarrow \star$ by TC-TAG-FUN.

($S = S_1 \rightarrow S_2, T = T_1 \rightarrow T_2$) By the IH on T_1 and S_1 , we have $\vdash \text{coerce}(T_1, S_1) : T_1 \rightarrow S_1$; by the IH on S_2 and T_2 , we have $\vdash \text{coerce}(S_2, T_2) : S_2 \rightarrow T_2$. By TC-WRAP on these coercions, we have $\vdash \text{wrap}(\text{coerce}(T_1, S_1), \text{coerce}(S_2, T_2)) : (S_1 \rightarrow S_2) \rightarrow (T_1 \rightarrow T_2)$.

($S = \star, T = T_1 \rightarrow T_2$) By TC-CHK-FUN, we have $\vdash \text{fun}? : \star \rightarrow (\star \rightarrow \star)$. By the IH on T_1 and \star (which are smaller in total than our original function type and \star), we know $\vdash \text{coerce}(T_1, \star) : T_1 \rightarrow \star$. Similarly, the IH on \star and T_2 , we know $\vdash \text{coerce}(\star, T_2) : \star \rightarrow T_2$. By TC-WRAP, we have $\vdash \text{wrap}(\text{coerce}(T_1, \star), \text{coerce}(\star, T_2)) : (\star \rightarrow \star) \rightarrow (T_1 \rightarrow T_2)$. Finally, by TC-SEQ, we can combine our first coercion with this to have $\vdash \text{fun}?; \text{wrap}(\text{coerce}(T_1, \star), \text{coerce}(\star, T_2)) : \star \rightarrow (T_1 \rightarrow T_2)$.

($S = T_1 \rightarrow T_2, T = \star$) By the IH on \star and T_1 (which are smaller in total than our original function type and \star), we have $\vdash \text{coerce}(\star, T_1) : \star \rightarrow T_1$. Similarly, by the IH on T_2 and \star , we have $\vdash \text{coerce}(T_2, \text{tdyn}) : T_2 \rightarrow \star$. By TC-WRAP, we have $\vdash \text{wrap}(\text{coerce}(\star, T_1), \text{coerce}(T_2, \star)) : (T_1 \rightarrow T_2) \rightarrow (\star \rightarrow \star)$. By TC-TAG-FUN, we have $\vdash \text{fun}! : (\star \rightarrow \star) \rightarrow \star$. Finally, we tie everything together with TC-SEQ: $\vdash \text{wrap}(\text{coerce}(\star, T_1), \text{coerce}(T_2, \star)); \text{fun}! : (T_1 \rightarrow T_2) \rightarrow \star$.

(otherwise) If none of the other cases apply, we generate a coercion through \star ; such a coercion is doomed to fail. It is nevertheless well typed. First, observe that neither S nor T can be \star , since one of the cases above would have adhered. So we can use the IH on S and \star or \star and T , since every other type is larger than \star .

By the IH on S and \star , we have $\vdash \text{coerce}(S, \star) : S \rightarrow \star$. Similarly, by the IH on \star and T , we have $\vdash \text{coerce}(\star, T) : \star \rightarrow T$. By TC-SEQ, we have $\vdash \text{coerce}(S, \star); \text{coerce}(\star, T) : S \rightarrow T$. \square

To keep things relatively neat notationally, we write $\sigma(X)$ to mean applying $\text{SUBST}(\sigma, -)$ to every indeterminate part of the structure X , where X might be a context Γ , expression e , or type T .

THEOREM C.2 (MODELS PRODUCE WELL TYPED TERMS (??)). If $\Gamma \vdash e \Rightarrow e', T, \phi$ and σ is a model for ϕ , then $\sigma(\Gamma) \vdash \sigma(e') : \sigma(T)$.

PROOF. By induction on the coercion insertion judgment.

(ID) By T-ID.

(CONST) By T-CONST, T-COERCE, and Lemma ??.

(FUN) Since σ is a model of $\phi_1 \wedge \phi_2$, it is also a model for ϕ_1 . So by the IH on e , we have $\sigma(\Gamma), x : \sigma(\alpha) \vdash e : \sigma(T)$. By Lemma ??, we know that $\vdash \text{coerce}(\sigma(T), \sigma(\beta)) : \sigma(T) \rightarrow \sigma(\beta)$. By T-COERCE, we have $\sigma(\Gamma), x : \sigma(\alpha) \vdash [\text{coerce}(\sigma(T), \sigma(\beta))]e : \sigma(\beta)$. Finally, by T-FUN, we have $\sigma(\Gamma) \vdash \text{fun}(x : \sigma(\alpha)).[\text{coerce}(\sigma(T), \sigma(\beta))]e : \sigma(\alpha \rightarrow \beta)$. The outer coercion is typed by T-COERCE and Lemma ??.

(APP) Since σ is a model of $\phi_1 \wedge \dots \wedge \phi_3 \wedge \phi_4 \wedge \phi_5$, it is also a model for each ϕ_i . By the IHs, we have:

$$\sigma(\Gamma) \vdash \sigma(e_1) : \sigma(T_1) \quad \text{and} \quad \sigma(\Gamma) \vdash \sigma(e_2) : \sigma(T_2).$$

By Lemma ??, we know that:

$$\vdash \text{coerce}(\sigma(T_1), \sigma(\alpha \rightarrow \beta)) : \sigma(T_1) \rightarrow \sigma(\alpha \rightarrow \beta)$$

and $\vdash \text{coerce}(\sigma(\beta), \sigma(\gamma)) : \sigma(\beta) \rightarrow \sigma(\gamma)$. We know that $\sigma(T_2) = \sigma(\alpha)$ by ϕ_4 , so by applying T-COERCE on the function and T-APP, we have:

$$\sigma(\Gamma) \vdash ([\text{coerce}(\sigma(T_1), \sigma(\alpha \rightarrow \beta))] \sigma(e_1)) \sigma(e_2) : \sigma(\beta)$$

We account for the outer coercion with T-COERCE and Lemma ??.

(MUL) Since σ is a model for $\phi_1 \wedge \phi_2$, it is also a model for ϕ_1 and ϕ_2 . By the IHs, we have:

$$\sigma(\Gamma) \vdash \sigma(e_1) : \sigma(T_1) \quad \text{and} \quad \sigma(\Gamma) \vdash \sigma(e_2) : \sigma(T_2).$$

By Lemma ??, we have:

$$\vdash \text{coerce}(\sigma(T_1), \text{int}) : \sigma(T_1) \rightarrow \text{int} \quad \text{and} \quad \vdash \text{coerce}(\sigma(T_2), \text{int}) : \sigma(T_2) \rightarrow \text{int}$$

By applying T-COERCE twice and T-MUL, we have:

$$\sigma(\Gamma) \vdash ([\text{coerce}(\sigma(T_1), \text{int})] \sigma(e_1)) \times ([\text{coerce}(\sigma(T_2), \text{int})] \sigma(e_2)) : \text{int}$$

The outer coercion is typed by T-COERCE and Lemma ??. □

D EXISTENCE OF MODELS

We show that models always exist for well scoped programs.

First, we borrow the “well scoped” relation from ?. We then show that a fully dynamic model always exists for such well scoped programs, and that it is stable under weakening. Let V be a set of variables. We say a term e is well scoped if $\emptyset \vdash e \text{ ok}$.

$$\begin{array}{c}
 \boxed{V \vdash e \text{ ok}} \\
 \\
 \text{WS-ID} \quad \frac{x \in V}{V \vdash x \text{ ok}} \qquad \qquad \qquad \text{WS-CONST} \quad \frac{}{V \vdash c \text{ ok}} \\
 \\
 \text{WS-FUN} \quad \frac{V \cup \{x\} \vdash e \text{ ok}}{V \vdash \text{fun}(x : S).e \text{ ok}} \qquad \qquad \qquad \text{WS-APP} \quad \frac{V \vdash e_1 \text{ ok} \quad V \vdash e_2 \text{ ok}}{V \vdash e_1 e_2 \text{ ok}} \\
 \\
 \text{WS-MUL} \quad \frac{V \vdash e_1 \text{ ok} \quad V \vdash e_2 \text{ ok}}{V \vdash e_1 \times e_2 \text{ ok}}
 \end{array}$$

Let $\text{dynctx}(V)$ be defined as the context that maps every variable in V to \star :

$$\begin{aligned}
 \text{dynctx}(\emptyset) &= \cdot \\
 \text{dynctx}(\Gamma, x : T) &= \text{dynctx}(\Gamma), x : \star
 \end{aligned}$$

THEOREM D.1 (WELL SCOPED TERMS HAVE DYNAMIC MODELS). *If $V \vdash e \text{ ok}$, then there exist e', T, ϕ , and ϕ such that for all dynamic models σ :*

- (1) $\text{dynctx}(V) \vdash e \Rightarrow e', T, \phi$,
- (2) ϕ is satisfiable in σ , and
- (3) $\text{SUBST}(\sigma, T) = \star$.

PROOF. By induction on the derivation of $V \vdash e \text{ ok}$. We must take the right disjunct of every constraint except for two: the outer coercion on variables and applications could safely take either disjunct.

(WS-ID) We have $x \in V$, so $x : \star \in \text{dynctx}(V)$. By ID; whether we pick the left or right disjunct, we have $\alpha = \star = \text{dynctx}(V)(x)$ (and so we will always find $\text{SUBST}(\sigma, \alpha) = \star$) and ϕ is satisfiable in all dynamic models.

(WS-CONST) We have $T = \text{ty}(c)$ and $\phi = \text{true}$ by CONST. Pick $\alpha = \star$; we have $\phi = \text{true}$. The former is just \star under SUBST.

(WS-FUN) We know that $V \cup \{x\} \vdash e \text{ ok}$; by the IH, we have $\text{dynctx}(V), x : \star \vdash e \Rightarrow e', T, \phi_1$ such that ϕ_1 is satisfiable. Since $\alpha = \star$ and $\text{SUBST}(\sigma, T) = \star$, we know that $\text{SUBST}(\sigma, \alpha \rightarrow T) = \star \rightarrow \star$, so we have $\text{ground}(\alpha \rightarrow T)$. Pick $\beta = \star$. We already know ϕ_1 is satisfiable, as is the right disjunct of ϕ_2 . We have $\text{SUBST}(\sigma, \star) = \star$ immediately.

(WS-APP) We know that $V \vdash e_1 \text{ ok}$ and $V \vdash e_2 \text{ ok}$. By the IH on e_1 , we have $\text{dynctx}(V) \vdash e_1 \Rightarrow e'_1, T_1, \phi_1$ such that ϕ_1 is satisfiable in dynamic models and $\text{SUBST}(\sigma, T_1) = \star$. Similarly, the IH on e_2 finds $\text{dynctx}(V) \vdash e_2 \Rightarrow e'_2, T_2, \phi_2$ such that ϕ_2 is satisfiable in dynamic models and $\text{SUBST}(\sigma, T_2) = \star$.

Since ϕ_1 and ϕ_2 are both satisfiable in all models where all variables map to \star , so is $\phi_1 \wedge \phi_2$. Pick $\alpha = \beta = \gamma = \star$. We satisfy the right disjunction of ϕ_3 , and we’ve already established ϕ_4 (because T_2 will substitute to \star , which is exactly equal to α). We could take either disjunction if ϕ_5 —we already know $\beta = \star$, so $\gamma = \star$ either way. We have $\text{SUBST}(\sigma, \gamma) = \star$ immediately.

(WS-MUL) We know that $V \vdash e_1$ ok and $V \vdash e_2$ ok. By the IH on e_1 , we have $\text{dynctx}(V) \vdash e_1 \Rightarrow e'_1, T_1, \phi_1$ such that ϕ_1 is satisfiable in dynamic models and $\text{SUBST}(\sigma, T_1) = \star$. Similarly, the IH on e_2 finds $\text{dynctx}(V) \vdash e_2 \Rightarrow e'_2, T_1, \phi_1$ such that ϕ_2 is satisfiable in dynamic models and $\text{SUBST}(\sigma, T_1) = \star$.

Since ϕ_1 and ϕ_2 are both satisfiable in all models where all variables map to \star , so is $\phi_1 \wedge \phi_2$. Picking $\alpha = \star$, we take the right disjuncts of ϕ_3, ϕ_4 , and ϕ_5 . We have $\text{SUBST}(\sigma, \alpha) = \star$ immediately. \square

LEMMA D.2 (DYNAMIC TERMS ARE STABLE UNDER WEAKENING).

If $\text{SUBST}(\sigma, T) = \star$, then $\text{WEAKEN}(\star, T)$ is satisfiable.

PROOF. Immediate: $\text{WEAKEN}(\star, T) = P(\star, T, \text{true}) = \text{true}$. \square

COROLLARY D.3. *If $V \vdash e$ ok, then it has a satisfiable model that is stable under weakening.*

PROOF. The term e has a dynamic model (Theorem ??) and dynamic models are stable under weakening (Lemma ??). \square

E ALL BENCHMARKS

adversarial/01-farg-mismatch.gtlc

Gtubi. : Rejected

InsAndOuts. : Compatible

```
(fun f:bool -> int. f true) (fun x:bool. x + 100)
```

MGT. : Compatible

```
(fun f : any. ((f)) ((true))) ((fun x : int. x + 100))
```

MaxMigrate. : Compatible

```
(fun f : bool -> int . f true) (fun x : any . x + 100)
```

TypeWhich2. : Compatible

```
(fun f:bool -> int. f true) (fun x:bool. (x : any) + 100)
```

TypeWhich. : Compatible

```
(fun f:any -> int. f true) (fun x:any. x + 100)
```

adversarial/02-rank2-poly-id.gtlc

Gtubi. : Rejected

InsAndOuts. : Compatible

```
(fun i:any -> any. (fun a:any. i true) (i 5)) (fun x:any. x)
```

MGT. : Compatible

```
(fun i : any. (fun a : any. ((i)) ((true))) (((i)) ((5)))) ((fun x : any. x))
```

MaxMigrate. : Runtime Error

```
(fun i : any -> any . (fun a : int . i true) (i 5)) (fun x : bool . x)
```

TypeWhich2. : Compatible

```
(fun i:any -> any. (fun a:any. i true) (i 5)) (fun x:any. x)
```

TypeWhich. : Compatible

```
(fun i:any -> any. (fun a:any. i true) (i 5)) (fun x:any. x)
```

adversarial/03-unreachable-error.gtlc

Gtubi. : Rejected

InsAndOuts. : Compatible

```
(fun b:(t4418 -> t4418) -> (any -> int) -> any -> int. b
  (fun c . (fun x:int. x x) 5 5) (fun d:any. 0))
(fun t:t4418 -> t4418. fun f:any -> int. f)
```

Solver-based Gradual Type Migration

MGT. : Compatible

```
(fun b : (any -> any) -> (any -> int) -> any -> int.  
b (fun c : any. (((fun x : any. ((x) x) ((5)))) ((5))) (fun d : any. 0))  
(fun t : any -> any. fun f : any -> int. f)
```

MaxMigrate. : Rejected

TypeWhich2. : Compatible

```
(fun b:(any -> any) -> (any -> int) -> any -> int. b (fun c:any.  
  (fun x:any. x x) 5 5) (fun d:any. 0)) (fun t:any -> any.  
fun f:any -> int. f)
```

TypeWhich. : Compatible

```
(fun b:(any -> any) -> (any -> int) -> any -> int. b (fun c:any.  
  (fun x:any. x x) 5 5) (fun d:any. 0)) (fun t:any -> any.  
fun f:any -> int. f)
```

adversarial/04-f-in-f-out.gtlc

Gtubi. : Compatible

```
((fun f : (int -> int). ((fun y : int. f) (f 5))) (fun x : int. (10 + x)))
```

InsAndOuts. : Compatible

```
(fun f:any -> int. (fun y:int. f) (f 5)) (fun x:any. 10 + x)
```

MGT. : Compatible

```
(fun f : int -> int. (fun y : int. f) (f 5)) (fun x : int. 10 + x)
```

MaxMigrate. : Runtime Error

```
(fun f : int -> any . (fun y : bool . f) (f 5)) (fun x : int . 10 + x)
```

TypeWhich2. : Compatible

```
(fun f:any -> int. (fun y:int. f) (f 5)) (fun x:any. 10 + x)
```

TypeWhich. : Compatible

```
(fun f:int -> int. (fun y:int. f) (f 5)) (fun x:int. 10 + x)
```

adversarial/05-order3-fun.gtlc

Gtubi. : Compatible

```
(fun f : ((beta@1 -> beta@2) -> beta@1). (fun x : (beta@1 -> beta@2). (x (f x))))
```

InsAndOuts. : Compatible

```
fun f:any. fun x:any. x (f x)
```

MGT. : Compatible

```
fun f : (any -> any) -> any. fun x : any -> any. x (f x)
```

MaxMigrate. : Unusable

```
fun f : int -> int . fun x : any . x (f x)
```

TypeWhich2. : Compatible

```
fun f:any. fun x:any. x (f x)
```

TypeWhich. : Compatible

```
fun f:(any -> any) -> any. fun x:any -> any. x (f x)
```

adversarial/06-order3-intfun.gtlc

Gtubi. : Compatible

```
(fun f : ((int -> int) -> (int -> beta@5)). (fun g : (int -> int). ((f g) ((g 10) + 1))))
```

InsAndOuts. : Compatible

```
fun f:any. fun g:any. f g (g 10 + 1)
```

MGT. : Compatible

```
fun f : (int -> int) -> int -> any. fun g : int -> int. f g ((g 10) + 1)
```

MaxMigrate. : Restricted

```
fun f : int -> int -> int . fun g : any . f g (g 10 + 1)
```

TypeWhich2. : Compatible

```
fun f:any. fun g:any. f g (g 10 + 1)
```

TypeWhich. : Restricted

```
fun f:(int -> int) -> int -> int. fun g:int -> int. f g (g 10 + 1)
```

adversarial/07-double-f.gtlc

Gtubi. : Restricted

```
(fun f : (bool -> bool). (f (f true)))
```

InsAndOuts. : Compatible

```
fun f:any. f (f true)
```

MGT. : Restricted

```
fun f : bool -> bool. f (f true)
```

MaxMigrate. : Restricted

```
fun f : any -> int . f (f true)
```

TypeWhich2. : Compatible

```
fun f:any. f (f true)
```


Solver-based Gradual Type Migration

TypeWhich. : Restricted

```
fun f:bool -> bool. f (f true)
```

adversarial/08-outflows.gtlc

Gtubi. : Rejected

InsAndOuts. : Compatible

```
(fun x:int. x 5 + x) 5
```

MGT. : Compatible

```
(fun x : any. (((x)) ((5))) + x) ((5))
```

MaxMigrate. : Compatible

```
(fun x : any . x 5 + x) 5
```

TypeWhich2. : Compatible

```
(fun x:int. (x : any) 5 + x) 5
```

TypeWhich. : Compatible

```
(fun x:int. (x : any) 5 + x) 5
```

adversarial/09-precision-relation.gtlc

Gtubi. : Rejected

InsAndOuts. : Compatible

```
(fun f:any -> int. f true + (fun g:int -> int. g 5) f) (fun x:any. 5)
```

MGT. : Compatible

```
(fun f : any. (((f)) ((true))) + ((fun g : int -> any. g 5) ((f))))  
((fun x : any. 5))
```

MaxMigrate. : Runtime Error

```
(fun f : bool -> int . f true + (fun g : any -> int . g 5) f) (fun x : bool . 5)
```

TypeWhich2. : Compatible

```
(fun f:any -> int. f true + (fun g:any -> int. g 5) f) (fun x:any. 5)
```

TypeWhich. : Compatible

```
(fun f:any -> int. f true + (fun g:any -> int. g 5) f) (fun x:any. 5)
```

adversarial/10-if-tag.gtlc

Gtubi. : Rejected

InsAndOuts. : Compatible

```
fun tag:any. fun x:any. if tag then x + 1 else if x then 1 else 0
```

MGT. : Restricted

```
fun tag : bool.  
fun x : bool. if tag then (((x)) + ((1))) else if x then 1 else 0
```

MaxMigrate. : Rejected

TypeWhich2. : Compatible

```
fun tag:any. fun x:any. if tag then x + 1 else if x then 1 else 0
```

TypeWhich. : Restricted

```
fun tag:bool. fun x:bool. if tag  
  then (x : any) + 1  
  else if x then 1 else 0
```

migeed/01-apply-add.gtlc

Gtubi. : Rejected

InsAndOuts. : Compatible

```
fun x:any. x (x + 1)
```

MGT. : Compatible

```
fun x : any -> any. x (((x)) + ((1)))
```

MaxMigrate. : Compatible

```
fun x : any . x (x + 1)
```

TypeWhich2. : Compatible

```
fun x:any. x (x + 1)
```

TypeWhich. : Compatible

```
fun x:any -> any. x ((x : any) + 1)
```

migeed/02-add-applied.gtlc

Gtubi. : Rejected

InsAndOuts. : Compatible

```
fun x:any. x (x true + 1)
```

MGT. : Restricted

```
fun x : bool -> any. x (((x true) + ((1))))
```

Solver-based Gradual Type Migration

MaxMigrate. : Compatible

```
fun x : any -> int . x (x true + 1)
```

TypeWhich2. : Compatible

```
fun x:any. x (x true + 1)
```

TypeWhich. : Compatible

```
fun x:any -> int. x (x true + 1)
```

migeed/03-add-two-applies.gtlc

Gtubi. : Rejected

InsAndOuts. : Compatible

```
fun x:any. x 4 + x true
```

MGT. : Compatible

```
fun x : any. (((x)) ((4))) + (((x)) ((true)))
```

MaxMigrate. : Compatible

```
fun x : any -> int . x 4 + x true
```

TypeWhich2. : Compatible

```
fun x:any. x 4 + x true
```

TypeWhich. : Compatible

```
fun x:any -> int. x 4 + x true
```

migeed/04-identity-four.gtlc

Gtubi. : Compatible

```
((fun x : int. x) 4)
```

InsAndOuts. : Compatible

```
(fun x:int. x) 4
```

MGT. : Compatible

```
(fun x : int. x) 4
```

MaxMigrate. : Compatible

```
(fun x : int . x) 4
```

TypeWhich2. : Compatible

```
(fun x:int. x) 4
```

TypeWhich. : Compatible

(fun x:int. x) 4

migeed/05-succ-id-id.gtlc

Gtubi. : Rejected

InsAndOuts. : Compatible

1 + (fun y:bool. y) ((fun x:bool. x) true)

MGT. : Compatible

1 + ((fun y : int. y) (((fun x : any. x) ((true))))))

MaxMigrate. : Compatible

1 + (fun y : int . y) ((fun x : any . x) true)

TypeWhich2. : Compatible

1 + (fun y:bool. y) ((fun x:bool. x) true)

TypeWhich. : Compatible

1 + (fun y:bool. (y : any)) ((fun x:bool. x) true)

migeed/06-identity.gtlc

Gtubi. : Compatible

(fun x : _t0. x)

InsAndOuts. : Compatible

fun x:any. x

MGT. : Compatible

fun x : any. x

MaxMigrate. : Restricted

fun x : int . x

TypeWhich2. : Compatible

fun x:any. x

TypeWhich. : Restricted

fun x:int. x

migeed/07-apply2.gtlc

Gtubi. : Compatible

(fun x : _t1. (fun y : (_t1 -> (_t1 -> beta@2)). ((y x) x)))

Solver-based Gradual Type Migration

InsAndOuts. : Compatible

```
fun x:any. fun y:any. y x x
```

MGT. : Compatible

```
fun x : any. fun y : any -> any -> any. y x x
```

MaxMigrate. : Unusable

```
fun x : any . fun y : int -> bool -> int . y x x
```

TypeWhich2. : Compatible

```
fun x:any. fun y:any. y x x
```

TypeWhich. : Compatible

```
fun x:any. fun y:any -> any -> any. y x x
```

migeed/08-indirect-apply-self.gtlc

Gtubi. : Rejected

InsAndOuts. : Compatible

```
fun x:any. (fun y:any. x) x x
```

MGT. : Compatible

```
fun x : any. (((fun y : any. x) x)) x
```

MaxMigrate. : Unusable

```
fun x : any . (fun y : int . x) x x
```

TypeWhich2. : Compatible

```
fun x:any. (fun y:any. x) x x
```

TypeWhich. : Compatible

```
fun x:any -> any. (fun y:any -> any. x) x x
```

migeed/09-the-long-one.gtlc

Gtubi. : Compatible

```
(fun x : _t3. (((fun f : (_t3 -> int). (((fun xx : (_t3 -> int). (fun y : int. xx)) f) (f x))) (fun z : _t3. 1
```

InsAndOuts. : Compatible

```
fun x:any. (fun f:any -> int. (fun xx:any -> int. fun y:int. xx) f (f  
  x)) (fun z:any. 1)
```

MGT. : Compatible

```
fun x : any.
(fun f : any -> int. (fun xx : any -> int. fun y : int. xx) f (f x))
(fun z : any. 1)
```

MaxMigrate. : Compatible

```
fun x : int . (fun f : any . (fun xx : int . fun y : int . xx) f (f x)) (fun z : int . 1)
```

TypeWhich2. : Compatible

```
fun x:any. (fun f:any -> int. (fun xx:any -> int. fun y:int. xx) f (f
x)) (fun z:any. 1)
```

TypeWhich. : Compatible

```
fun x:any. (fun f:any -> int. (fun xx:any -> int. fun y:int. xx) f (f
x)) (fun z:any. 1)
```

migeed/10-apply-self.gtlc

Gtubi. : Rejected

InsAndOuts. : Compatible

```
fun x:any. x x
```

MGT. : Compatible

```
fun x : any. ((x)) x
```

MaxMigrate. : Rejected

TypeWhich2. : Compatible

```
fun x:any. x x
```

TypeWhich. : Compatible

```
fun x:any -> any. x x
```

migeed/11-untypable-in-sys-f.gtlc

Gtubi. : Rejected

InsAndOuts. : Rejected

MGT. : Compatible

```
(fun x : (any -> any -> any) -> any.
fun y : any -> any -> any.
y (x ((fun x : any. x))) (x (fun b : any. fun c : any. b)))
((fun d : any. ((d)) d))
```

MaxMigrate. : Rejected

TypeWhich2. : Compatible

```
(fun x:(any -> any) -> any. fun y:any. y (x (fun x:any. x)) (x
  (fun b:any. fun c:any. b))) (fun d:any -> any. d d)
```

TypeWhich. : Compatible

```
(fun x:(any -> any) -> any. fun y:any -> any -> any. y (x (fun x:any.
  x)) (x (fun b:any. fun c:any. b))) (fun d:any -> any. d d)
```

migeed/12-self-interpreter.gtlc

Gtubi. : Rejected

InsAndOuts. : Rejected

MGT. : Compatible

```
(fun h : ((any -> any) ->
  (any -> any -> any) -> ((any -> any) -> any -> any) -> any) ->
  ((any -> any) -> (any -> any -> any) -> ((any -> any) -> any -> any) -> any) ->
  any.
  (fun x : any. h (((x) x)))
    ((fun x : (any -> any) ->
      (any -> any -> any) -> ((any -> any) -> any -> any) -> any.
        h x x)))
  ((fun e : any.
    fun m : (any -> any) ->
      (any -> any -> any) -> ((any -> any) -> any -> any) -> any.
        m (fun x : any. x) (fun m : any. fun n : any. (((e) m)) ((e) n))
          (fun m : any -> any. fun v : any. ((e) (m v))))))
```

MaxMigrate. : Rejected

TypeWhich2. : Compatible

```
(fun h:any -> any -> any. (fun x:any -> any. h (x x)) (fun x:any. h x
  x)) (fun e:any. fun m:any. m (fun x:any. x) (fun m:any. fun n:any. e
  m (e n)) (fun m:any. fun v:any. e (m v)))
```

TypeWhich. : Compatible

```
(fun h:any -> any -> any. (fun x:any -> any. h (x x)) (fun x:any. h x
  x)) (fun e:any. fun m:any. m (fun x:any. x) (fun m:any. fun n:any. e
  m (e n)) (fun m:any. fun v:any. e (m v)))
```

F GRIFT BENCHMARKS

Name	Description
array	Adapted from ?
sieve	Adapted from ?
n-body	Adapted from ?
tak	Adapted from ?
ray	Adapted from ?
blackscholes	Adapted from ?
matmult	400 × 400 matrix multiplication
quicksort	Quicksort on worst-case input
quicksort-pairs	Quicksort implemented with pairs
fft	Adapted from ?
cps	Mutually recursive even-odd in continuation-passing style