



Gradual Liquid Type Inference

NIKI VAZOU, IMDEA, Spain

ÉRIC TANTER, University of Chile, Chile

DAVID VAN HORN, University of Maryland, USA

Refinement types allow for lightweight program verification by enriching types with logical predicates. Liquid typing provides a decidable refinement inference mechanism that is convenient but subject to two major issues: (1) inference is global and requires top-level annotations, making it unsuitable for inference of modular code components and prohibiting its applicability to library code, and (2) inference failure results in obscure error messages. These difficulties seriously hamper the migration of existing code to use refinements.

This paper shows that *gradual liquid type inference*—a novel combination of liquid inference and gradual refinement types—addresses both issues. Gradual refinement types, which support imprecise predicates that are optimistically interpreted, can be used in argument positions to constrain liquid inference so that the global inference process effectively infers modular specifications usable for library components. Dually, when gradual refinements appear as the result of inference, they signal an inconsistency in the use of static refinements. Because liquid refinements are drawn from a finite set of predicates, in gradual liquid type inference we can enumerate the *safe concretizations* of each imprecise refinement, *i.e.*, the static refinements that justify why a program is gradually well-typed. This enumeration is useful for static liquid type error explanation, since the safe concretizations exhibit all the potential inconsistencies that lead to static type errors.

We develop the theory of gradual liquid type inference and explore its pragmatics in the setting of Liquid Haskell. To demonstrate the utility of our approach, we develop an interactive tool, GuiLT, for gradual liquid type inference in Liquid Haskell that both infers modular types and explores safe concretizations of gradual refinements. We report on the use of GuiLT for error reporting and discuss a case study on the migration of three commonly-used Haskell list manipulation libraries into Liquid Haskell.

CCS Concepts: • **Software and its engineering** → **Software verification**;

Additional Key Words and Phrases: liquid types, refinement types, gradual typing, error explanation

ACM Reference Format:

Niki Vazou, Éric Tanter, and David Van Horn. 2018. Gradual Liquid Type Inference. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 132 (November 2018), 25 pages. <https://doi.org/10.1145/3276502>

1 INTRODUCTION

Refinement types [Freeman and Pfenning 1991] allow for lightweight program verification by decorating existing program types with logical predicates. For instance, the type $\{x:\text{Int} \mid 0 < x\}$ denotes strictly positive integer values and can be used to validate *at compile time* the absence of division-by-zero errors by refining the type of division:

$$(/) :: \text{Int} \rightarrow x:\{ \text{Int} \mid 0 < x \} \rightarrow \text{Int}$$

*This work is partially funded by CONICYT, FONDECYT Regular Project 1150017, Chile.

Authors' addresses: Niki Vazou, IMDEA, Software Institute, 28223 Pozuelo de Alarcón, Madrid, 28906, Spain; Éric Tanter, University of Chile, Computer Science Department (DCC), Beauchef 851, Santiago, Chile; David Van Horn, University of Maryland, Dept. of Computer Science, 8223 Paint Branch Drive, College Park, Maryland, 20742, USA.



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

© 2018 Copyright held by the owner/author(s).

2475-1421/2018/11-ART132

<https://doi.org/10.1145/3276502>

A major challenge with refinement types is to support decidable automatic checking and *inference*. To this end, *liquid types* restrict refinement predicates to decidable theories [Rondon et al. 2008]. The type $\{ x:\text{Int} \mid k \}$ describes integer values refined with some predicate k , that is automatically solved based on unifying the constraints generated at each use of x , resulting in a concrete refinement drawn from a *finite* domain of template refinements. The attractiveness of liquid types for programmers is usability. Verification only requires specification of top-level functions, while all intermediate types can be automatically inferred and checked.

It has long been recognized that automatic type inference makes error reporting very challenging [Wand 1986], resulting in severe usability problems. In particular, at an ill-typed function application, the system should decide whether to blame the function definition or the client; the particular choice will unfortunately and inevitably expose some of the internals of the inference procedure to the user. The situation is even worse with liquid type inference, because refinement verification is more involved than standard Hindley-Milner style unification. The difficulty of understanding error messages from liquid inference in turn makes it really hard to progressively migrate non-refined code to refined code, *e.g.*, from Haskell to Liquid Haskell [Vazou et al. 2014b].

For instance, consider the following function `divIf` that either inverts its argument x if it is positive, or inverts $1-x$ otherwise:

```
divIf x = if isPos x then 1/x else 1/(1-x)
```

The function relies on an imported function `isPos` of type $\text{Int} \rightarrow \text{Bool}$. If we want to give `divIf` a refined type, liquid inference starts from the signature:

```
divIf :: x:{ Int | kx } → { o:Int | ko }
```

and solves the predicate variables k_x and k_o based on the use sites. However, without any uses of `divIf` in the considered source code, and for reasons we will clarify in due course, liquid inference infers the useless refinements:

```
divIf :: x:{ Int | false } → { o:Int | false }
```

The `false` refinement on the argument means that `divIf` is dead code, since liquid type inference relies on a *closed world assumption*; and `divIf` has, at inference time, no clients.

Conversely, if the code base at inference time includes a “positive” client call `divIf 1`, the inferred precondition would be $0 < x$, triggering a type error in the *definition* of `divIf` due to the lack of information about `isPos`. This hard-to-predict and moving blame is not unique to liquid type inference; it frequently appears in type inference engines, yielding hard-to-debug error messages.

A key contribution of this paper is to recognize that, in such situations, treating the inferred refinement as an *unknown* refinement — in the sense of gradual typing [Siek and Taha 2006] — allows us to eliminate the closed world assumption, as well as globally explain liquid type errors. Specifically, we adapt the gradual refinement types of Lehmann and Tanter [2017] to the setting of liquid type inference, yielding gradual liquid type inference. Programmers can introduce gradual refinements such as $\{ x:\text{Int} \mid ? \}$ and inference exhaustively searches for *safe concretizations* (SCs for short), *i.e.*, the concrete refinements that can replace each occurrence of a gradually-refined variable to make the program well-typed. These SCs can then be used to understand liquid type errors and assist in migrating programs to adopt liquid types.

Contributions. This paper makes the following contributions:

- We give a semantics and inference algorithm for gradual liquid types (§ 4), by exploiting the abstract interpretation approach to gradual language design [Garcia et al. 2016]. We prove that inference is correct and satisfies the static criteria for gradual languages [Siek et al. 2015]. In

particular, the static gradual guarantee ensures that only making types more precise can introduce new type errors. This result is, to the best of our knowledge, novel for a gradual inference system.

- We implement gradual liquid type inference in GuiLT, as an extension of Liquid Haskell (§ 5). The implementation integrates a number of optimizations and heuristics in order to be applicable to existing Haskell libraries. GuiLT implements the static aspects of gradual liquid typing, leaving insertion of runtime checks to future work.
- GuiLT takes as input a Haskell program annotated with gradual refinements and generates an interactive program that presents all valid static choices to interpret each separate occurrence of an unknown refinement, if any. The user can explore suggested predicates and decide which ones to use so that their program is well-typed (§ 6).
- We use GuiLT for user-guided migration of three existing Haskell libraries (1260 LoC) to Liquid Haskell (§ 7), demonstrating that gradual liquid type inference can feasibly be supported and used interactively for program migration.

The practical experience we report on with GuiLT in this article has to be understood as a proof of *feasibility* of the approach; more work remains to be done in order to allow the technique to scale adequately.

The essence of the novelty of our approach—that gradual inference based on abstract interpretation can be fruitfully used for error reporting and program migration—is not restricted to refinement types. We conjecture that the principles of gradual types can be used for type error explanation in further typing systems.

This paper proceeds as follows. We first provide an informal overview of liquid type inference, gradual refinement types, and gradual liquid type inference (§ 2). Next, we formally establish the necessary background on liquid types and gradual refinements (§ 3), before presenting the main results (§ 4–7). Finally, we discuss related work (§ 8) and conclude (§ 9). Auxiliary definitions as well as proofs of theorems can be found in the extended version Vazou et al. [2018a]. Source code for GuiLT is available at:

<https://github.com/ucsd-progsys/liquidhaskell/tree/gradual>

2 BACKGROUND AND OVERVIEW

We start by recalling the basic notions of refinement types (§ 2.1), their decidable fragment known as liquid types (§ 2.2), and gradual refinements (§ 2.3). Then (§ 2.4), we combine gradual and liquid types to build a usable interactive inference procedure that can be used for both error explanation and program migration.

2.1 Refinement Types

To explain refinement type checking, assume the below type signatures for the example of § 1:

```
isPos :: x:Int → { b:Bool | b ⇔ 0 < x }
divIf :: Int → Int
```

With these specifications, `divIf` is well-typed because the positive restriction in the precondition of `(/)` is provably satisfied. Refinement type checking proceeds in three steps, described below.

Step 1: Constraint Generation. Based on the code and the specifications, refinement subtyping constraints are generated; for our example, two subtyping constraints are generated, one for each call to `(/)`. They stipulate that, in the environment with the argument x and the boolean branching guard b , the second argument to `(/)` (i.e., $v = x$ and $v = 1-x$, resp.) is *safe*, that is, it respects the precondition $0 < v$.

$$b:\{ (b \Leftrightarrow 0 < x) \wedge b \} \quad \vdash \{ v \mid v = x \} \leq \{ v \mid 0 < v \}$$

$$b:\{ (b \Leftrightarrow 0 < x) \wedge \text{not } b \} \vdash \{ v \mid v = 1 - x \} \leq \{ v \mid 0 < v \}$$

For space, we write $\{ v \mid p \}$ to denote $\{ v:t \mid p \}$ when the type t is clear; we omit the refinement variables from the environment, simplifying $x:\{ x \mid p \}$ to $x:\{ p \}$, and we skip uninformative refinements such as $x:\{ \text{true} \}$.

In both constraints the branching guard b is refined with the result refinement of `isPos`: $b \Leftrightarrow 0 < x$. Also, the environment is strengthened with the value of the condition in each branch: b in the `then` branch, `not b` in the `else` branch.

Step 2: Verification Conditions. Each subtyping constraint is reduced to a logical verification condition (VC), that validates if, assuming all the refinements in the environment, the refinement on the left-hand side implies the one on the right-hand side. For instance, the two constraints above reduce to the following VCs.

$$\begin{aligned} (b \Leftrightarrow 0 < x) \wedge b &\Rightarrow v = x &\Rightarrow 0 < v \\ (b \Leftrightarrow 0 < x) \wedge \text{not } b &\Rightarrow v = 1 - x &\Rightarrow 0 < v \end{aligned}$$

Step 3: Implication Checking. Finally, an SMT solver is used to check the validity of the generated VCs and thus determine if the program is well-typed. Here, the SMT solver determines that both VCs are valid, thus `divIf` is well-typed.

2.2 Liquid Types

When not all refinement types are spelled out explicitly, one can use *inference*. For instance, the well-typedness of `divIf` crucially relies on the guard predicate, as propagated by the refinement type of `isPos`. We now explain how liquid typing [Rondon et al. 2008] infers a type for `divIf` in case `isPos` is an imported, *unrefined* function. Liquid types introduce refinement type variables, known as *liquid variables*, for unspecified refinements. As in § 1, the type of `divIf` is assigned the liquid variables k_x and k_o for the input and the output refinements, *resp.*:

$$\text{divIf} :: x:\{ \text{Int} \mid k_x \} \rightarrow \{ o:\text{Int} \mid k_o \}$$

After generating subtyping constraints as in § 2.1, the inference procedure attempts to find a solution for the liquid variables such that all the constraints are satisfied. If no solution can be found, the program is deemed ill-typed.

Step 1: Constraint Generation. After introduction of the liquid variables, the following subtyping constraints are generated for `divIf`.

$$\begin{aligned} x:\{ k_x \}, b:\{ b \} \vdash \{ v \mid v = x \} &\leq \{ v \mid 0 < v \} \\ x:\{ k_x \}, b:\{\text{not } b\} \vdash \{ v \mid v = 1 - x \} &\leq \{ v \mid 0 < v \} \end{aligned}$$

Step 2: Constraint Solving. Liquid inference then solves the liquid variables k so that the subtyping constraints are satisfied. The solving procedure takes as input a *finite* set of refinement *templates* \mathbb{Q}^\star abstracted over program variables. For example, the template set \mathbb{Q}^\star below describes ordering predicates, with \star ranging over program variables.

$$\mathbb{Q}^\star = \{ 0 < \star, 0 \leq \star, \star < 0, \star \leq 0, \star < \star, \star \leq \star \}$$

Next, for each liquid variable, the set \mathbb{Q}^\star is instantiated with all program variables in scope, to generate well-sorted predicates. Instantiation of \mathbb{Q}^\star for the liquid variables k_x and k_o leads to the following concrete predicate candidates, where trivially true (*i.e.*, $x \leq x$) and trivially false (*i.e.*, $x < x$) predicates are ignored for simplicity.

$$\begin{aligned} \mathbb{Q}^x &= \{ 0 < x, 0 \leq x, x < 0, x \leq 0 \} \\ \mathbb{Q}^o &= \{ 0 < o, 0 \leq o, o < 0, o \leq 0, o < x, x < o, \dots \} \end{aligned}$$

Finally, inference iteratively computes the strongest solution for each liquid variable that satisfies the constraints. It starts from an initial solution that maps each variable to the logical conjunction of all the instantiated templates $A = \{k_x \mapsto \bigwedge Q^x, k_o \mapsto \bigwedge Q^o\}$. It then repeatedly filters out predicates of the solution until all constraints are satisfied.

In our example, the initial solution includes both the predicates $0 < x$ and $x < 0$, which are *contradictory*, so the conjunction of solutions implies false: consequently, both liquid variables are solved to false. As discussed in § 1, this solution is inferred as false based on a closed-world assumption that is useless in practice¹. With client code that imposes additional constraints, the inferred solution can be more useful, though the reported errors can be hard to interpret.

2.3 Gradual Refinement Types

We observe that we can exploit gradual typing in order to assist inference and provide better support for error explanation and program migration. Instead of interpreting an unspecified refinement as a liquid variable, let us use the unknown gradual refinement $\{ \text{Int} \mid ? \}$ for the argument type of `divIf` [Lehmann and Tanter 2017].

$$\text{divIf} :: x:\{ \text{Int} \mid ? \} \rightarrow \text{Int}$$

This gradual precondition specifies that for each *usage occurrence* of the argument x , there must *exist* a concrete refinement (which we call a *safe concretization*, SC) for which the (non-gradual) program type checks. Key to this definition is that the refinement that exists need not be unique to all occurrences of the identifier. Gradual refinement type checking proceeds as follows.

Step 1: Constraint Generation. First, we generate the subtyping constraints derived from the definition of `divIf` that now contain gradual refinements.

$$\begin{aligned} x:\{ ? \}, b:\{ b \} &\vdash \{ v \mid v = x \} \leq \{ v \mid 0 < v \} \\ x:\{ ? \}, b:\{ \text{not } b \} &\vdash \{ v \mid v = 1 - x \} \leq \{ v \mid 0 < v \} \end{aligned}$$

Step 2: Gradual Verification Conditions. Each subtyping reduces to a VC, where each gradual refinement such as $x:\{ ? \}$ translates intuitively to an existential refinement ($\exists p. p \ x$). The solution of these existentials are the safe concretizations (SCs) of the program. Here, we informally use $\exists^? p$ to denote such existentials over predicates and call the corresponding verification conditions *gradual VCs* (GVCs). For example, the GVCs for `divIf` are the following.

$$\begin{aligned} (\exists^? p_{\text{then}} \cdot p_{\text{then}} \ x) \wedge b &\Rightarrow v = x \Rightarrow 0 < v \\ (\exists^? p_{\text{else}} \cdot p_{\text{else}} \ x) \wedge \text{not } b &\Rightarrow v = 1 - x \Rightarrow 0 < v \end{aligned}$$

Step 3: Gradual Implication Checking. Checking the validity of the generated GVCs is an open problem. In the `divIf` example, we can, by observation, find the SCs that render the GVCs valid: $p_{\text{then}} \ x \mapsto 0 < x$ and $p_{\text{else}} \ x \mapsto x \leq 0$.

More importantly, we can present these solutions to the user as the conditions under which `divIf` is well-typed. We use the SCs to explain to the user that for `divIf` to type check under a static type, the gradual precondition $?$ should be replaced with a refinement that implies $0 < x$ in the `then` branch and $x \leq 0$ in the `else` branch. Since such a refinement cannot exist, we use SCs to explain to the user that the $?$ cannot be replaced with a refinement in the type of `divIf`, unless the user modifies the code (here, strengthen the postcondition of `isPos`, as in § 2.1).

Our goal is to find an algorithmic procedure that solves GVCs. Lehmann and Tanter [2017] describe how GVCs over linear arithmetic can be checked, while Courcelle and Engelfriet [2012]

¹The actual implementation of Liquid Haskell departs from this behavior, to comply with Haskell's module system that does not assume closed-world. Instead, Liquid Haskell assumes a trivially refined type for each unused, exported function.

describe a more general logical fragment (monadic second order logic) with an algorithmic decision procedure. Yet, in both cases we lose the justifications, *i.e.*, the SCs, and thus the opportunity to use such SCs for error explanation and migration assistance.

2.4 Gradual Liquid Type Inference

To algorithmically solve GVCs we exhaustively search for SCs in the *finite* predicate domain of liquid types. In between constraint generation and constraint solving, gradual liquid type inference concretizes the constraints by instantiating gradual refinements with each possible liquid template.

Step 1: Constraint Generation. Constraint generation is performed exactly like gradual refinement types, leading to the constraints of § 2.3 for the `divIf` example.

Step 2: Constraint Concretization. We exhaustively generate all the possible concretizations of the constraints. For example, $x:\{?\}$ can be concretized to any predicate from the \mathbb{Q}^x set of § 2.2, yielding $|\mathbb{Q}^x|^2 = 16$ concrete constraint sets, among which the following two:

(1) **Concretization for** $p_{\text{then}} x \mapsto \emptyset < x, p_{\text{else}} x \mapsto \emptyset < x$:

$$\begin{aligned} x:\{\emptyset < x\}, b:\{b\} &\vdash \{v \mid v = x\} \leq \{v \mid \emptyset < v\} \\ x:\{\emptyset < x\}, b:\{\text{not } b\} &\vdash \{v \mid v = 1 - x\} \leq \{v \mid \emptyset < v\} \end{aligned}$$

(2) **Concretization for** $p_{\text{then}} x \mapsto \emptyset < x, p_{\text{else}} x \mapsto x \leq 0$:

$$\begin{aligned} x:\{\emptyset < x\}, b:\{b\} &\vdash \{v \mid v = x\} \leq \{v \mid \emptyset < v\} \\ x:\{x \leq 0\}, b:\{\text{not } b\} &\vdash \{v \mid v = 1 - x\} \leq \{v \mid \emptyset < v\} \end{aligned}$$

Step 3: Constraint Solving. After concretization, liquid constraint solving finds out the valid ones. In our example, the constraint 1 above is invalid while 2 is valid. Out of the 16 concrete constraints, only two are valid, with $p_{\text{then}} x \mapsto \emptyset < x$ and $p_{\text{else}} x \mapsto x \leq 0$ or $p_{\text{else}} x \mapsto x < 0$. Thus, `divIf` (gradually) type checks and also the inference provides to the user the SCs of p_{then} and p_{else} as an explanation of type checking.

Application to Error Explanation. The contradictory solutions of the gradual refinement of the argument indicates to the user that the code cannot statically type check. The user needs to edit either the code or the refinement types. For example, the code can be fixed by providing a precise type for `isPos` (following § 2.1) or by restricting `divIf`'s domain to positive numbers (rendering the `else` branch as dead code).

Unlike current liquid type inference, our algorithm does not assume a closed world, since the SCs it provides do not depend on call sites of the functions. Thus, the output of the algorithm is modular: it explains the code contradictions that lead to type errors but these contradictions are generated in the function definition and do not rely on the arguments of function calls.

Importantly, gradual liquid types are used for error explanation using the language of contradictions in refinement predicates that the user understands. In this example, the output of our algorithm informs the user that the program cannot type check because the same refinement should be solved to $\emptyset < x$ in the `then` branch and to $x < 0$ in the `else` branch, which is impossible. This explanation is much more informative than the current liquid type algorithm, which would generate a type error either in the definition of `divIf` or at its call sites, following the closed world assumption, as discussed in § 1.

In § 6 we use gradual liquid types to explain to the user a typical off-by-one error. Since the algorithm is exhaustively searching all potential solutions, it is exponentially slow on the number of potential solutions. Yet, in § 7 we show that we can reasonably apply our technique on real Haskell code due to the three following reasons:

Constants	$c ::= \wedge \mid \neg \mid = \mid \dots$ $\mid \text{true} \mid \text{false}$ $\mid 0, 1, -1, \dots$	Basic Types	$b ::= \text{int} \mid \text{bool}$
Values	$v ::= c \mid \lambda x. e$	Types	$t ::= \{x:b \mid p\} \mid x:t \rightarrow t$
Expressions	$e ::= v \mid x \mid e x$ $\mid \text{if } x \text{ then } e \text{ else } e$ $\mid \text{let } x = e \text{ in } e$ $\mid \text{let } x:t = e \text{ in } e$	Environment	$\Gamma ::= \cdot \mid \Gamma, x:t$
Predicates	$p ::= e$	Substitution	$\sigma ::= \cdot \mid \sigma, (x, e)$
		Constraint	$C ::= \Gamma \vdash \{v:b \mid p\}$ $\mid \Gamma \vdash \{v:b \mid p\} \leq \{v:b \mid p\}$

Fig. 1. Syntax of λ_R^e .

- Our algorithm is a modular, per-function analysis. Thus inference time depends on the size of the analyzed function, and not on the size of the whole codebase to be type checked.
- Our implementation is user interactive, thus the exponential complexity is not necessarily a problem in practice, since our algorithm runs in the background while expecting the user input. Once a SC is found, the system presents it to the user, while looking for the next SCs in the background.
- Finally, in § 5 we discuss further technical optimizations that make our theoretically-exponential algorithm tractable and quite usable in practice, at least to justify the feasibility of the approach. Admittedly, though, scalability remains a challenge to be addressed in future work.

In § 4 we formalize the inference steps and prove the correctness and the gradual criteria of our algorithm. Various implementation considerations necessary for the algorithm to be usable in practice are described in § 5. We report on its use for error explanation and program migration in § 6 and § 7.

3 LIQUID TYPES AND GRADUAL REFINEMENTS

We briefly provide the technical background required to describe gradual liquid types. We start with the semantics and rules of a generic refinement type system (§ 3.1) which we then adjust to describe both liquid types (§ 3.2) and gradual refinement types (§ 3.3).

3.1 Refinement Types

Syntax. Figure 1 presents the syntax of a standard functional language with refinement types, λ_R^e . The expressions of the language include constants, lambda terms, variables, function applications, conditionals, and let bindings. Note that the argument of a function application needs to be syntactically a variable, as must the condition of a conditional; this normalization is standard in refinement types as it simplifies the formalization [Rondon et al. 2008]. There are two let binding forms, one where the type of the bound variable is inferred and one where it is explicitly declared.

λ_R^e types include *base refinements* $\{x:b \mid p\}$ where b is a base type (int or bool) refined with the logical predicate p . A predicate can be any expression e , which can refer to x . Types also include dependent function types $x:t_x \rightarrow t$, where x is bound to the function argument and can appear in the result type t . As usual, we write b as a shortcut for $\{x:b \mid \text{true}\}$ and $t_x \rightarrow t$ as a shortcut for $x:t_x \rightarrow t$ when x does not appear in t .

Denotations. Following Knowles and Flanagan [2010], each type of λ_R^e denotes a set of expressions. The denotation of a base refinement includes all expressions that either diverge or evaluate to base values that satisfy the associated predicate. We write $\models e$ to represent that e is (operationally) valid and $e \Downarrow$ to represent that e terminates:

$$\models e \doteq e \hookrightarrow^* \text{true} \quad e \Downarrow \doteq \exists v. e \hookrightarrow^* v$$

Typing $\Gamma \vdash e:t$

$$\begin{array}{c}
\frac{\Gamma(x) = \{v:b \mid _ \}}{\Gamma \vdash x:\{v:b \mid v = x\}} \text{ T-VAR-BASE} \quad \frac{\Gamma(x) \text{ is a function type}}{\Gamma \vdash x:\Gamma(x)} \text{ T-VAR} \quad \frac{}{\Gamma \vdash c:ty(c)} \text{ T-CONST} \\
\\
\frac{\Gamma \vdash e:t_e \quad \Gamma \vdash t \quad \Gamma \vdash t_e \leq t}{\Gamma \vdash e:t} \text{ T-SUB} \quad \frac{\Gamma, x:t_x \vdash e:t \quad \Gamma \vdash x:t_x \rightarrow t}{\Gamma \vdash \lambda x.e:(x:t_x \rightarrow t)} \text{ T-FUN} \\
\\
\frac{\Gamma \vdash y:t_x}{\Gamma \vdash e:(x:t_x \rightarrow t)} \text{ T-APP} \quad \frac{\text{fresh } x' \quad \Gamma_1 \doteq \Gamma, x':\{v:\text{bool} \mid x\} \quad \Gamma_2 \doteq \Gamma, x':\{v:\text{bool} \mid \neg x\}}{\Gamma \vdash x:\{v:\text{bool} \mid _ \} \quad \Gamma_1 \vdash e_1:t \quad \Gamma_2 \vdash e_2:t \quad \Gamma \vdash t} \text{ T-IF} \\
\Gamma \vdash e \ y:t[y/x] \quad \Gamma \vdash \text{if } x \text{ then } e_1 \text{ else } e_2:t \\
\\
\frac{\Gamma \vdash e_x:t_x \quad \Gamma, x:t_x \vdash e:t \quad \Gamma \vdash t}{\Gamma \vdash \text{let } x = e_x \text{ in } e:t} \text{ T-LET} \quad \frac{\Gamma \vdash e_x:t_x \quad \Gamma, x:t_x \vdash e:t \quad \Gamma \vdash t \quad \Gamma \vdash t_x}{\Gamma \vdash \text{let } x:t_x = e_x \text{ in } e:t} \text{ T-SPEC}
\end{array}$$

Sub-Typing $\Gamma \vdash t \leq t$

$$\frac{\text{isValid}(\Gamma \vdash \{v:b \mid p_1\} \leq \{v:b \mid p_2\})}{\Gamma \vdash \{v:b \mid p_1\} \leq \{v:b \mid p_2\}} \text{ S-BASE} \quad \frac{\Gamma \vdash t_{x_2} \leq t_{x_1} \quad \Gamma, x:t_{x_2} \vdash t_1 \leq t_2}{\Gamma \vdash x:t_{x_1} \rightarrow t_1 \leq x:t_{x_2} \rightarrow t_2} \text{ S-FUN}$$

Well-Formedness $\Gamma \vdash t$

$$\frac{\text{isValid}(\Gamma \vdash \{v:b \mid p\})}{\Gamma \vdash \{v:b \mid p\}} \text{ W-BASE} \quad \frac{\Gamma \vdash t_x \quad \Gamma, x:t_x \vdash t}{\Gamma \vdash x:t_x \rightarrow t} \text{ W-FUN}$$

Fig. 2. Static Semantics of λ_R^e . (Types colored in blue need to be inferred.)

where $\cdot \hookrightarrow^* \cdot$ is the reflexive, transitive closure of the small-step reduction relation. Denotations are naturally extended to function types and environments (as sets of substitutions).

$$\begin{aligned}
\llbracket \{x:b \mid p\} \rrbracket &\doteq \{e \mid e \vdash b, \text{ if } e \Downarrow \text{ then } \models p[e/x]\} \\
\llbracket x:t_x \rightarrow t \rrbracket &\doteq \{e \mid \forall e_x \in \llbracket t_x \rrbracket. e \ e_x \in \llbracket t[e_x/x] \rrbracket\} \\
\llbracket \Gamma \rrbracket &\doteq \{\sigma \mid \forall x:t \in \Gamma. (x, e) \in \sigma \wedge e \in \llbracket \sigma \cdot t \rrbracket\}
\end{aligned}$$

Static semantics. Figure 2 summarizes the standard typing rules that characterize whether an expression belongs to the denotation of a type [Knowles and Flanagan 2010; Rondon et al. 2008]. Namely, $e \in \llbracket t \rrbracket$ iff $\vdash e:t$. We define three kinds of relations 1. typing, 2. subtyping, and 3. well-formedness.

(1) *Typing:* $\Gamma \vdash e:t$ iff $\forall \sigma \in \llbracket \Gamma \rrbracket. \sigma \cdot e \in \llbracket \sigma \cdot t \rrbracket$.

Rule T-VAR-BASE refines the type of a variable with its exact value. Rule T-CONST types a constant c using the function $ty(c)$ that is assumed to be sound, *i.e.*, we assume that for each constant c , $c \in \llbracket ty(c) \rrbracket$. Rule T-SUB allows to weaken the type of a given expression by subtyping, discussed below. Rule T-IF achieves path sensitivity by typing each branch under an environment strengthened with the value of the condition. Finally, the two let binding rules T-LET and T-SPEC only differ in whether the type of the bound variable is inferred or taken from the syntax. Note that the last premise, a well-formedness condition, ensures that the bound variable does not escape (at the type level) the scope of the let form.

(2) *Subtyping:* $\Gamma \vdash t_1 \leq t_2$ iff $\forall \sigma \in \llbracket \Gamma \rrbracket, e \in \llbracket \sigma \cdot t_1 \rrbracket. e \in \llbracket \sigma \cdot t_2 \rrbracket$.

Rule S-BASE uses the relation $\text{isValid}(\cdot)$ to check subtyping on basic types; we leave this relation abstract for now since we will refine it in the course of this section. Knowles and Flanagan

[2010] define subtyping between base refinements as:

$$\text{isValid}(\Gamma \vdash \{x:b \mid p_1\} \leq \{x:b \mid p_2\}) \quad \text{iff} \quad \forall \sigma \in \llbracket \Gamma, x:b \rrbracket. \text{if } \models \sigma \cdot p_1 \text{ then } \models \sigma \cdot p_2$$

This definition makes checking undecidable, as it quantifies over all substitutions. We come back to decidability below.

- (3) *Well-Formedness*: Rule W-BASE overloads $\text{isValid}(\cdot)$ to refer to well-formedness on base refinements. A base refinement $\{x:b \mid p\}$ is well-formed only when p is typed as a boolean:

$$\text{isValid}(\Gamma \vdash \{x:b \mid p\}) \quad \text{iff} \quad \Gamma, x:b \vdash p:\text{bool}$$

Inference. In addition to being undecidable, the typing rules in Figure 2 are not syntax directed: several types do not come from the syntax of the program, but have to be guessed—they are colored in blue in Figure 2. These are: the argument type of a function (Rule T-FUN), the common (least upper bound) type of the branches of a conditional (Rule T-IF), and the resulting type of let expressions (Rules T-LET and T-SPEC), which needs to be weakened to not refer to variable x in order to be well-formed. Thus, to turn the typing relation into a type checking algorithm, one needs to address both decidability of subtyping judgments and inference of the aforementioned types.

3.2 Liquid Types

Liquid types [Rondon et al. 2008] provide a decidable and efficient inference algorithm for the typing relation of Figure 2. For decidability, the key idea of liquid types is to restrict refinement predicates to be drawn from a *finite* set of *predefined*, SMT-decidable predicates $q \in \mathbb{Q}$.

Syntax. The syntax of *liquid predicates*, written \hat{p} , is:

$$\begin{array}{ll} \hat{p} ::= & \text{true} \quad \text{True} \\ & | \quad q \quad \text{Predicate, with } q \in \mathbb{Q} \\ & | \quad \hat{p} \wedge \hat{p} \quad \text{Conjunction} \\ & | \quad \kappa \quad \text{Liquid Variable} \\ A ::= & \cdot \mid A, \kappa \mapsto \bar{q} \quad \text{Solution} \end{array}$$

A liquid predicate can be true (true), an element from the predefined set of predicates (q), a conjunction of predicates ($\hat{p} \wedge \hat{p}$), or a predicate variable (κ), called a *liquid variable*. A *solution* A is a mapping from liquid variables to a set of elements of \mathbb{Q} . The set \bar{q} represents a variable-free liquid predicate using true for the empty set and conjunction to combine the elements otherwise.

Checking. When all the predicates in \mathbb{Q} belong to SMT-decidable theories, validity checking of λ_R^e : $\text{isValid}(\Gamma \vdash \{x:b \mid p_1\} \leq \{x:b \mid p_2\})$ which quantifies over all embeddings of the typing environment, can be SMT automated in a sound and complete way. Concretely, a subtyping judgment $\Gamma \vdash \{x:b \mid \hat{p}_1\} \leq \{x:b \mid \hat{p}_2\}$ is valid *iff* under all the assumptions of Γ , the predicate \hat{p}_1 implies the predicate \hat{p}_2 .

$$\begin{array}{c} \text{isValid}(\Gamma \vdash \{x:b \mid \hat{p}_1\} \leq \{x:b \mid \hat{p}_2\}) \\ \text{iff} \\ \text{isSMTValid}(\bigwedge \{\hat{p} \mid x:\{x:b \mid \hat{p}\} \in \Gamma\} \Rightarrow \hat{p}_1 \Rightarrow \hat{p}_2) \end{array}$$

Inference. The liquid inference algorithm, defined in Figure 3, first applies the rules of Figure 2 using liquid variables as the refinements of the types that need to be inferred and then uses an iterative algorithm to solve the liquid variable as a subset of \mathbb{Q} [Rondon et al. 2008] (steps 1 and 2 of § 2.2).

More precisely, given a typing environment $\hat{\Gamma}$, an expression \hat{e} , and the fixed set of predicates \mathbb{Q} , the function $\text{infer } \hat{\Gamma} \hat{e} \mathbb{Q}$ returns just the type of the expression \hat{e} under the environment $\hat{\Gamma}$, if it exists, or nothing otherwise. It first generates a template type \hat{t} and a set of constraints C that

```

infer :: Env → Expr → Quals → Maybe Type
infer  $\hat{\Gamma} \hat{e} \mathbb{Q} = A \langle * \rangle \hat{t}$ 
  where
    A = solve C A0
    ( $\hat{t}$ , C) = cons  $\hat{\Gamma} \hat{e}$ 

cons :: Env → Expr → (Maybe Type, [Cons])
solve :: [Cons] → Sol → Maybe Sol

```

Fig. 3. **Liquid Inference Algorithm** (cons and solve are defined in [Vazou et al. 2018a]).

contain liquid variables in the types to be inferred. Then it generates a solution A that satisfies all the constraints in C . Finally, it returns the type \hat{t} in which all the liquid variables have been substituted by concrete refinements from the mapping in A .

The function $\text{cons } \hat{\Gamma} \hat{e}$ uses the typing rules in Figure 2 to generate the template type $\text{Just } \hat{t}$ of the expression \hat{e} , *i.e.*, a type that potentially contains liquid variables, and the basic constraints that appear in the leaves of the derivation tree of the judgment $\hat{\Gamma} \vdash \hat{e}:\hat{t}$. If the derivation rules fail, then $\text{cons } \hat{\Gamma} \hat{e}$ returns Nothing and an empty constraint list. The function $\text{solve } C A$ uses the decidable validity checking to iteratively pick a constraint in $c \in C$ that is not satisfied, while such a constraint exists, and weakens the solution A so that c is satisfied. The function $A \langle * \rangle \hat{t}$ applies the solution A to the type \hat{t} , if both contain Just values, otherwise returns Nothing . Here and in the following, we pose: $A_0 = \lambda\kappa.\mathbb{Q}$.

The algorithm $\text{infer } \hat{\Gamma} \hat{e} \mathbb{Q}$ is terminating and sound and complete with respect to the typing relation $\hat{\Gamma} \vdash \hat{e}:\hat{t}$ as long as all the predicates are conjunctions of predicates drawn from \mathbb{Q} .

3.3 Gradual Refinement Types

Gradual refinement types [Lehmann and Tanter 2017] extend the refinements of λ_R^e to include imprecise refinements like $x > 0 \wedge ?$. While they describe the static and dynamic semantics of gradual refinements, inference is left as an open challenge. Our work extends liquid inference to gradual refinements, therefore we hereby summarize their basics.

Syntax. The syntax of gradual predicates in $\lambda_G^{\tilde{p}}$ is

$$\begin{array}{ll} \tilde{p} ::= p & \text{Precise Predicate} \\ | p \wedge ? & \text{Imprecise Predicates, where } p \text{ is local} \end{array}$$

A predicate is either *precise* or *imprecise*. The syntax of an imprecise predicate $p \wedge ?$ allows for a *static part* p . Intuitively, with the predicate $x > 0 \wedge ?$, x is statically (and definitely) positive, but the type system can optimistically assume stronger, non-contradictory requirements about x . To make this intuition precise and derive the complete static and dynamic semantics of gradual refinements, Lehmann and Tanter [2017] follow the Abstracting Gradual Typing methodology (AGT) [Garcia et al. 2016]. AGT starts from a specification of the meaning of gradual types as denoting *sets of static types*. For instance, in simple typing $?$ denotes any static type, while $? \rightarrow \text{int}$ denotes all functions that return integers. This denotation is captured by a *concretization* function γ . For instance, $\gamma(?) = \text{TYPE}$ and $\gamma(? \rightarrow \text{int}) = \{t \rightarrow \text{int} \mid t \in \text{TYPE}\}$.

Following AGT, a gradual refinement type (resp. predicate) is given meaning by concretization to the set of static types (resp. predicates) it represents. Defining this concretization requires introducing two important notions.

Specificity. First, we say that p_1 is *more specific* than p_2 , written $p_1 \leq p_2$, iff p_2 is true when p_1 is:

$$p_1 \leq p_2 \doteq \forall \sigma. \text{if } \models \sigma \cdot p_1 \text{ then } \models \sigma \cdot p_2$$

Locality. Additionally, in order to prevent imprecise formulas from introducing contradictions—which would defeat the purpose of refinement checking—Lehmann and Tanter [2017] identify the need for the static part of an imprecise refinement to be *local*. Intuitively, a refinement is local if it only restricts the refinement variable and not other variables in scope. Using an explicit syntax $p(x)$ to explicitly declare the variable x refined by the predicate p , a refinement is local if there exists a value v for which $p[v/x]$ is true; and this, for any (well-typed) substitution that closes the predicate:

$$\text{isLocal}(p(x)) \doteq \forall \sigma, \exists v. \models \sigma \cdot p[v/x]$$

Concretization. Using specificity and locality, the concretization function $\gamma(\cdot)$ maps gradual predicates to the set of the static predicates they represent.

$$\begin{aligned} \gamma(p(x)) &\doteq \{p\} \\ \gamma((p \wedge ?)(x)) &\doteq \{p' \mid p' \leq p, \text{isLocal}(p'(x))\} \end{aligned}$$

A precise predicate concretizes to itself (singleton), while an imprecise predicate denotes all the local predicates more specific than its static part. This definition extends naturally to types and environments.

$$\begin{aligned} \gamma(\{x:b \mid \tilde{p}\}) &\doteq \{\{x:b \mid p\} \mid p \in \gamma(\tilde{p}(x))\} \\ \gamma(x:\tilde{t}_x \rightarrow \tilde{t}) &\doteq \{x:t_x \rightarrow t \mid t_x \in \gamma(\tilde{t}_x), t \in \gamma(\tilde{t})\} \\ \gamma(\tilde{\Gamma}) &\doteq \{\Gamma \mid x:t \in \Gamma \text{ iff } x:\tilde{t} \in \tilde{\Gamma}, t \in \gamma(\tilde{t})\} \end{aligned}$$

The denotations of gradual refinement types are similar to those from § 3.1. The denotation of a base *imprecise* gradual refinement $\{x:b \mid p \wedge ?\}$ includes all (gradually-typed) expressions that satisfy at least p .

Type Checking. Figure 2 is used “as is” to type gradual expressions $\tilde{\Gamma} \vdash \tilde{e}:\tilde{t}$, save for the fact that the validity predicate must be lifted to operate on gradual types. $\text{isValid}(\cdot)$ holds if there exists a justification, by concretization, that the static judgment holds. Precisely:

$$\begin{aligned} \text{isValid}(\tilde{\Gamma} \vdash \tilde{t}_1 \leq \tilde{t}_2) &\doteq \exists \Gamma \in \gamma(\tilde{\Gamma}), t_1 \in \gamma(\tilde{t}_1), t_2 \in \gamma(\tilde{t}_2). \text{isValid}(\Gamma \vdash t_1 \leq t_2) \\ \text{isValid}(\tilde{\Gamma} \vdash \tilde{t}) &\doteq \exists \Gamma \in \gamma(\tilde{\Gamma}), t \in \gamma(\tilde{t}). \text{isValid}(\Gamma \vdash t) \end{aligned}$$

4 GRADUAL LIQUID TYPES

We now formalize the combination of liquid type inference and gradual refinements to later use gradual liquid types for both error explanation and program migration. We extend the work of Lehmann and Tanter [2017] by adapting the liquid type inference algorithm to the gradual setting. To do so, we apply the abstract interpretation approach of AGT [Garcia et al. 2016] to lift the infer function (defined in § 3.2) so that it operates on gradual liquid types.

Below is the syntax of predicates in $\lambda_{GL}^{\tilde{p}}$, a gradual liquid core language whose predicates are gradual predicates where the static part of an imprecise predicate is a liquid predicate, with the additional requirement that it is *local* (def. in § 3.3).

$$\begin{aligned} \tilde{p} &::= \hat{p} && \text{Precise Liquid Predicate} \\ &| \hat{p} \wedge ? && \text{Imprecise Liquid Predicate, where } \hat{p} \text{ is local} \end{aligned}$$

The elements of $\lambda_{GL}^{\tilde{p}}$ are both gradual and liquid; *i.e.*, expressions \tilde{e} could also be written as \tilde{e} . Also, we write $?$ as a shortcut for the imprecise predicate $\text{true} \wedge ?$.

Our goal is to define $\text{infer} \checkmark \checkmark \mathbb{Q}$ so that it returns a type \checkmark such that $\checkmark \vdash \checkmark : \checkmark$. After deriving infer using AGT (§ 4.1), we provide an algorithmic characterization of infer (§ 4.3), which serves as the basis for our implementation. We present the properties that infer satisfies in § 4.4.

4.1 Lifting Liquid Inference

We define the function infer using the abstracting gradual typing (AGT) methodology [Garcia et al. 2016]. In AGT, once concretization γ is fixed, one needs to define the sound and optimal abstraction function α that, together with γ , forms a Galois connection. Then, AGT explains how to lift the static semantics of the language to operate on gradual types through this Galois connection. In particular, AGT defines the consistent lifting of a function f as: $\checkmark f \checkmark = \alpha(\{f t \mid t \in \gamma(\checkmark)\})$.

The question is how to apply this general approach to the liquid type inference algorithm. We answer this question via trial-and-error.

Try 1. Lifting infer. Assume we lift infer in a similar manner, *i.e.*, we pose

$$\text{infer} \checkmark \checkmark \mathbb{Q} = \alpha(\{\text{infer} \Gamma e \mathbb{Q} \mid \Gamma \in \gamma(\checkmark), e \in \gamma(\checkmark)\})$$

This definition of infer is too strict: it rejects expressions that should be accepted. Consider for instance the following expression \checkmark that defines a function f with an imprecisely-refined argument:

```
// onlyPos :: { v: Int | 0 < v } → Int
// check :: Int → Bool
let f :: x: { Int | ? } → Int
    f x = if check x then onlyPos x else onlyPos (-x)
in f 42
```

There is no single static expression $e \in \gamma(\checkmark)$ such that the definition of f above type checks. For any \mathbb{Q} we will get $\text{infer} \{ \} e \mathbb{Q} = \text{Nothing}$ which denotes a type inference failure. This behavior occurs because by the definition of infer the gradual argument of f needs to be concretized before calling infer , which breaks the flexibility programmers expect from gradual refinements (this example is based on the motivation example of Lehmann and Tanter [2017]). One expects that $\text{infer} \{ \} \checkmark \mathbb{Q}$ should simply return Int . A similar scenario appears in Garcia et al. [2016], where lifting the typing relation as a whole would be too imprecise and instead the lifted typing relation is defined by lifting the type functions and predicates used to define typing. Here, as described in § 3.2, infer calls the functions cons and solve , which in turn calls the function isValid . Which of these functions should we lift? Since cons is merely an algorithmic definition of the typing rules, it is not affected by the gradualization of the system, thus does not require lifting. On the contrary, solve is calling isValid that operates on gradual refinements. To get a precise inference system we first attempted to lift isValid .

Try 2. Lifting isValid. To our surprise, using gradual validity checking (the lifting of isValid , § 3.3) leads to an unsound inference algorithm! This is because, soundness of static inference implicitly relies on the property of validity checking that if two refinements p_1 and p_2 are right-hand-side valid, then so is their conjunction, *i.e.*,

$$\begin{aligned} \text{isValid}(\Gamma \vdash \{v:b \mid p\} \leq \{v:b \mid p_1\}) \quad \text{and} \quad \text{isValid}(\Gamma \vdash \{v:b \mid p\} \leq \{v:b \mid p_2\}) \\ \Rightarrow \\ \text{isValid}(\Gamma \vdash \{v:b \mid p\} \leq \{v:b \mid p_1 \wedge p_2\}) \end{aligned}$$

But this property does not hold for gradual validity checking, because for any logical predicate q , it is true that $(q \Rightarrow p_1 \wedge q \Rightarrow p_1)$ implies $(q \Rightarrow p_1 \wedge p_2)$, but $(\exists q.(q \Rightarrow p_1)) \wedge (\exists q.(q \Rightarrow p_1))$ does not imply that $\exists q.(q \Rightarrow p_1 \wedge p_2)$.

Try 3. *Lifting solve*. Let us try to lift solve:

$$\text{so}\check{\text{lve}} A \check{C} = \{\text{solve } A C \mid C \in \gamma(\check{C})\}$$

where \check{C} denotes a gradual constraint (from Figure 1). This approach is successful and leads to a provably sound and complete inference algorithm (§ 4.4).

Note that in the definition of $\text{so}\check{\text{lve}}$, we do not appeal to abstraction. This is because we can directly define $\text{in}\check{\text{fer}}$ to consider all produced solutions instead.

$$\begin{aligned} \text{in}\check{\text{fer}} &:: \text{Env} \rightarrow \text{Expr} \rightarrow \text{Quals} \rightarrow \text{Set Type} \\ \text{in}\check{\text{fer}} \check{\Gamma} \check{\epsilon} \mathbb{Q} &= \{\check{i}' \mid \text{Just } \check{i}' \leftarrow A \langle * \rangle \check{i}, A \in \text{so}\check{\text{lve}} A_0 \check{C}\} \\ \text{where } (\check{i}, \check{C}) &= \text{cons } \check{\Gamma} \check{\epsilon} \end{aligned}$$

First, function cons derives the typing constraints \check{C} , and if successful, the template type \check{i} (step 1 of § 2.4).² Next, we use the lifted $\text{so}\check{\text{lve}}$ to concretize and solve all the derived constraints (steps 2 and 3 of § 2.4, *resp.*). By keeping track of the concretizations that return non-Nothing solutions, we derive the safe concretizations of § 2. Finally, we apply each solution to the template gradual type, yielding a set of inferred types. We do not explicitly abstract the set of inferred types back to a single gradual type. Adding abstraction by exploiting the abstraction function defined by Lehmann and Tanter [2017] is left for future work.

The core of the gradual inference algorithm is a combination of the liquid and gradual refinement type systems. Yet, as we exposed by the failing attempt to lift isValid , this combination was not trivial, since blind application of the AGT [Garcia et al. 2016] methodology could lead to unsound inference. In § 4.4, we prove that our algorithm is sound and in § 5 we discuss an optimized implementation and its applications. The applications discussed in § 6 and § 7 make explicit use of the inferred set in order to assist users in understanding errors and migrating programs.³

4.2 When Does Gradual Liquid Inference Fail?

Given the flexibility induced by gradual refinements, at this point the reader might wonder when an actual type inference failure can occur. First, as we will formally prove later in this section, gradual liquid inference is a conservative extension of liquid inference, and therefore, in the absence of imprecise refinements, gradual liquid inference fails exactly when standard liquid inference fails.

More crucially, in presence of imprecise refinements, gradual inference only succeeds when there exist possible justifications for each individual occurrences of gradually-refined variables. Said otherwise, gradual liquid inference fails when there exists at least one occurrence of a gradually-refined variable for which there does not exist any valid concretization.

Consider the example below:

```
// onlyPos :: { v:Int | 0 < v } → Int
f :: x:{ Int | x ≤ 0 ∧ ? } → Int
f x = onlyPos x
```

Each valid concretization of the imprecise refinement $x \leq 0 \wedge ?$ should imply the precondition of onlyPos , namely $0 < v$. This however contradicts the static part of the imprecise refinement, $x \leq 0$. Therefore, gradual liquid inference fails for this program.

²Note that Cons is unchanged from the static system, because it only depends on the structure of the types, and not on the refinements themselves.

³In standard gradual typing, the set of static types denoted by a gradual type can be infinite, hence abstraction is definitely required. In contrast, here the structure of types is fixed, and the set of possible liquid refinements, even if potentially large, is finite. We can therefore do without abstraction. We discuss implementation considerations in § 6.

4.3 Algorithmic Concretization

To make infer algorithmic, we need to define an algorithmic concretization function of a set of constraints, $\gamma(\check{C})$. We do so, by using the the finite domain of predicates \mathbb{Q} .

Recall the concretization of gradual predicates: $\gamma((p \wedge ?)(x)) \doteq \{p' \mid p' \leq p, \text{isLocal}(p'(x))\}$. In general, this function cannot be algorithmically computed, since it ranges over the infinite domain of predicates. In gradual liquid refinements, the domain of predicates is restricted to the powerset of the *finite* domain \mathbb{Q} . We define the algorithmic concretization function $\gamma_{\mathbb{Q}}(\check{p}(x))$ as the intersection of the powerset of the finite domain \mathbb{Q} with the gradual concretization function.

$$\gamma_{\mathbb{Q}}(\check{p}(x)) \doteq 2^{\mathbb{Q}} \cap \gamma(\check{p}(x))$$

Concretization of gradual predicates reduces to (decidable) locality and specificity checking on the elements of \mathbb{Q} .

$$\gamma_{\mathbb{Q}}((\hat{p} \wedge ?)(x)) \doteq \{\hat{p}' \mid \hat{p}' \in 2^{\mathbb{Q}}, \hat{p}' \leq \hat{p}, \text{isLocal}(\hat{p}'(x))\}$$

We naturally extend the algorithmic concretization function to typing environments, constraints, and list of constraints.

$$\begin{aligned} \gamma_{\mathbb{Q}}(\check{\Gamma}) &\doteq \{\hat{\Gamma} \mid x:\hat{t} \in \hat{\Gamma} \text{ iff } x:\check{t} \in \check{\Gamma}, \hat{t} \in \gamma_{\mathbb{Q}}(\check{t})\} \\ \gamma_{\mathbb{Q}}(\check{\Gamma} \vdash \check{t}_1 \leq \check{t}_2) &\doteq \{\hat{\Gamma} \vdash \hat{t}_1 \leq \hat{t}_2 \mid \hat{\Gamma} \in \gamma_{\mathbb{Q}}(\check{\Gamma}), \hat{t}_i \in \gamma_{\mathbb{Q}}(\check{t}_i)\} \\ \gamma_{\mathbb{Q}}(\check{\Gamma} \vdash \check{t}) &\doteq \{\hat{\Gamma} \vdash \hat{t} \mid \hat{\Gamma} \in \gamma_{\mathbb{Q}}(\check{\Gamma}), \hat{t} \in \gamma_{\mathbb{Q}}(\check{t}), \} \\ \gamma_{\mathbb{Q}}(\check{C}) &\doteq \{\hat{C} \mid c \in \hat{C} \text{ iff } \check{c} \in \check{C}, \hat{c} \in \gamma_{\mathbb{Q}}(\check{c})\} \end{aligned}$$

We use $\gamma_{\mathbb{Q}}(\cdot)$ to define an algorithmic version of solve :

$$\text{solve } A \check{C} = \{\text{solve } A \hat{C} \mid \hat{C} \in \gamma_{\mathbb{Q}}(\check{C})\}$$

which in turn yields an algorithmic version of infer .

4.4 Properties of Gradual Liquid Inference

We prove that the inference algorithm infer satisfies the correctness criteria of [Rondon et al. \[2008\]](#), as well as the static criteria for gradually-typed languages [[Siek et al. 2015](#)].⁴ The corresponding proofs can be found in supplementary material [[Vazou et al. 2018a](#)].

4.4.1 Correctness of Inference. The algorithm infer is sound, complete, and terminates.

THEOREM 4.1 (CORRECTNESS). *Let \mathbb{Q} be a finite set of predicates from an SMT-decidable logic, $\check{\Gamma}$ a gradual liquid environment, and \check{e} a gradual liquid expression. Then*

- **Termination** $\text{infer } \check{\Gamma} \check{e} \mathbb{Q}$ terminates.
- **Soundness** If $\check{t} \in \text{infer } \check{\Gamma} \check{e} \mathbb{Q}$, then $\check{\Gamma} \vdash \check{e}:\check{t}$.
- **Completeness** If $\text{infer } \check{\Gamma} \check{e} \mathbb{Q} = \emptyset$, then $\nexists \check{t}. \check{\Gamma} \vdash \check{e}:\check{t}$.

The proof of termination is straightforward from the careful definition of the infer algorithm. Soundness and completeness rely on the property that a constraint is gradually valid *iff* there exists a concrete solution that renders it valid (*i.e.*, $\text{isValid}(A \cdot \check{c}) \text{ iff } \exists \hat{c} \in \gamma_{\mathbb{Q}}(\check{c}). \text{isValid}(A \cdot \hat{c})$). We use this property to prove soundness by the definition of the algorithm and completeness by contradiction. We note that, unlike the infer algorithm that provably returns the strongest possible solution, it is not clear how to relate the set of solutions returned by $\text{infer } \check{\Gamma} \check{e} \mathbb{Q}$ with the rest of the types that satisfy $\check{\Gamma} \vdash \check{e}:\check{t}$.

⁴Because this work focuses on the static semantics, *i.e.*, the inference algorithm, we do not discuss the dynamic part of the gradual guarantee, which has been proven for gradual refinement types [[Lehmann and Tanter 2017](#)].

4.4.2 *Gradual Typing Criteria.* Siek et al. [2015] list three criteria for the static semantics of a gradual language, which the infer algorithm satisfies. These criteria require the gradual type system (i) is a conservative extension of the static type system, (ii) is flexible enough to accommodate the dynamic end of the typing spectrum (in our case, unrefined types), and (iii) supports a smooth connection between both ends of the spectrum.

(i) *Conservative Extension.* The gradual inference algorithm infer coincides with the static algorithm infer on terms that only rely on precise predicates. More specifically, if infer infers a static type for a term, then infer returns only that type, for the same term. Conversely, if a term is not typeable with infer , it is also not typeable with infer .

THEOREM 4.2 (CONSERVATIVE EXTENSION). *If $\text{infer } \hat{\Gamma} \hat{e} \mathbb{Q} = \text{Just } \hat{t}$, then $\text{infer } \hat{\Gamma} \hat{e} \mathbb{Q} = \{\hat{t}\}$. Otherwise, $\text{infer } \hat{\Gamma} \hat{e} \mathbb{Q} = \emptyset$.*

The proof follows by the definitions of infer and the concretization function.

(ii) *Embedding of imprecise terms.* We then prove that given a well-typed unrefined term (i.e., simply-typed), refining all base types with the unknown predicate $?$ yields a well-typed gradual term. This property captures the fact that it is possible to “import” a simply-typed term into the gradual liquid setting. (In contrast, this is not possible without gradual refinements: just putting true refinements to all base types does not yield a well-typed program.)

To state this theorem, we use t_s to denote simple types (b and $t_s \rightarrow t_s$) and similarly e_s and Γ_s for terms and environments. The simply-typed judgment is the standard one.

The $[\cdot]$ function turns simple types into gradual liquid types by introducing the unknown predicate on every base type (and naturally extended to environments and terms):

$$[b] = \{v:b \mid ?\} \qquad [t_1 \rightarrow t_2] = x:[t_1] \rightarrow [t_2]$$

THEOREM 4.3 (EMBEDDING OF UNREFINED TERMS). *If $\Gamma_s \vdash e_s:t_s$, then $\text{infer } [\Gamma_s] [e_s] \mathbb{Q} \neq \emptyset$.*

Let t_s denote the unrefined version of the gradually refined type \check{t} . We first prove that for any gradual type \check{t} with only local refinement both $\Gamma \vdash \check{t} \leq [t_s]$ and $\Gamma \vdash [t_s] \leq \check{t}$ hold. Then, we prove that if $\Gamma_s \vdash e_s:t_s$, then $[\Gamma_s] \vdash [e_s]:[t_s]$. Finally, we use completeness of the inference algorithm to prove theorem 4.3.

For the expression $[e_s]$, the inference algorithm will generate a set of refinement types, all of which have the structure of t_s . The liquid algorithm infers types with strongest postconditions and the gradual inference algorithm is not generating any fresh $?$, so, the inference returns all valid types with shape t_s with the dynamic refinements instantiated to all valid local predicates and for each instantiation, the unknown refinements are solved following the principle of strongest postconditions.

(iii) *Static Gradual Guarantee.* Finally, the solid foundations for gradualization based on abstract interpretation allow us to effectively satisfy the *gradual guarantee*, which we believe was never proved for any gradual inference work. The gradual guarantee stipulates that typeability is monotonic in the precision of type information. In other words, making type annotations less precise cannot introduce new type errors.

We first define the notion of precision in terms of algorithmic concretization:

Definition 4.4 (Precision of Gradual Types). \check{t}_1 is less precise than \check{t}_2 , written as $\check{t}_1 \sqsubseteq \check{t}_2$, iff $\gamma_{\mathbb{Q}}(\check{t}_1) \subseteq \gamma_{\mathbb{Q}}(\check{t}_2)$.

Precision naturally extends to type environments and terms.

THEOREM 4.5 (STATIC GRADUAL GUARANTEE). *If $\check{\Gamma}_1 \sqsubseteq \check{\Gamma}_2$ and $\check{e}_1 \sqsubseteq \check{e}_2$, then for every $\check{t}_{1i} \in \text{infer } \check{\Gamma}_1 \check{e}_1 \mathbb{Q}$ there exists $\check{t}_{2i} \sqsubseteq \check{t}_{1i}$ so that $\check{t}_{2i} \in \text{infer } \check{\Gamma}_2 \check{e}_2 \mathbb{Q}$.*

The intuition of the proof is that since $\check{t}_{1i} \in \text{infer } \check{\Gamma}_1 \check{e}_1 \mathbb{Q}$, then the algorithm inferred a solution of the unknown refinements. We prove that this solution is also inferred for \check{e}_1 .

For a given term and type environment, the theorem ensures that, for every inferred type, the algorithm infers a less precise type when run on a less precise term and environment.

5 IMPLEMENTATION

We implemented `infer` as `GuiLT`, an extension to Liquid Haskell [Vazou et al. 2014a] that takes a Haskell program annotated with gradual refinement specifications and returns an `.html` interactive file that lets the user explore all safe concretizations.

Concretely, `GuiLT` uses the existing API of Liquid Haskell to implement the three steps of gradual liquid type checking steps described in § 2.4 (and formalized in § 4): 1) First, `GuiLT` calls the Liquid Haskell API to generate subtyping constraints that contain both liquid variables and imprecise predicates. 2) Next, it calls the liquid API to collect all the refinement templates. The templates are used to map each occurrence of imprecise predicates in the constraints to a set of concretizations. These concretizations are combined to generate the possible concretizations of the constraints. 3) Finally, using Liquid Haskell’s constraint solving it decides the validity of each concretized constraint, while all the safe concretizations SCs are interactively presented to the user.

The implementation of `GuiLT` closely follows the theory of § 4, apart from a syntactic detail (we use `??` instead of `?` to denote the unknown part of a refinement) and several practical adjustments that we discuss here and are crucial for usability.

Templates. To generate the refinement templates we use Liquid Haskell’s existing API. The generated templates consist of a predefined set of predicates for linear arithmetic ($v [< | \leq | > | \geq | = | \neq] x$), comparison with zero ($v [< | \leq | > | \geq | = | \neq] 0$), and length operations ($\text{len } v [\geq | >] 0$, $\text{len } v = \text{len } x$, $v = \text{len } x$, $v = \text{len } x + 1$), where v and x respectively range over the refinement variable and any program variable. Application-specific templates are automatically abstracted from user-provided specifications and the user can explicitly define custom templates. For instance, assume a user defined data type `Tree a` and an uninterpreted function `size :: Tree a → Int`. If the user writes a specification $x: \text{Tree } a \rightarrow \{v: \text{Int} \mid v < \text{size } x\}$, then the template $v < \text{size } x$ is generated where v and x respectively range over the refinement variable and any program variable.

Depth. For completeness in the theory, we check the validity of any solution, including all possible *conjunctions* of elements of \mathbb{Q} , which is not tractable in practice. The implementation uses an instantiation depth parameter that is 1 by default, meaning each `??` ranges over single templates. At depth 2, `??` ranges over conjunctions of (single) templates, *etc.*

Sensibility Checking. Each `??` can be instantiated with any templates that are local and specific. The implementation uses the SMT solver to check both. As an *optimization*, we perform a syntactic locality check to reject templates that are “non-sensible”, for instance, syntactic contradictions of the form $x < v \wedge v < x$. As a *heuristic*, we further filter out as non-sensible type-directed instantiations of the templates that, based on our experience, a user would not write, such as arithmetic operations on lists and booleans (e.g., $x < \text{False}$)—although potentially correct in Haskell through overloading.

Locality Checking. To encode Haskell functions (e.g., `len`) in the refinements, Liquid Haskell is using uninterpreted SMT functions [Vazou et al. 2018b]. As Lehmann and Tanter [2017] note, locality checking breaks under the presence of uninterpreted functions. For instance, the predicate

Fig. 4. The .html output of GuiLT for error explanation of incompatible indexing.

$0 < \text{len } i$ is not local on i , because $\exists i. 0 < \text{len } i$ is not SMT valid due to a model in which len is always negative. To check locality under uninterpreted functions we define a fresh variable (e.g., $\text{len}i$) for each function application (e.g., $\text{len } i$). For instance, $0 < \text{len } i$ is local on i because under the new encoding $\exists \text{len}i. 0 < \text{len}i$ is SMT valid.

Partitions. A critical optimization for efficiency is that after generation and before solving, the set of constraints is partitioned based on the constraint dependencies so that each partition is solved independently. Two constraints depend on each other when they contain the same liquid variable or when they contain different variables (e.g., k_1 and k_2) that depend on each other (e.g., $x:\{k_1\} \vdash \{v \mid \text{true}\} \leq \{v \mid k_2\}$). That way, we reduce the number of ?? that appear in each set of independent constraints and thus the number of concretization combinations that need to be checked (which increases exponentially with the number of ?? accounting for all combinations of all concretizations).

6 APPLICATION I: ERROR EXPLANATION

We illustrate how GuiLT can be used interactively for error explanation. Consider the list indexing function (!!) in Haskell:

```
(x:_) !! 0 = x
( _:xs) !! i = xs!!(i - 1)
_      !! _ = error "Out of Bounds!"
```

Indexing is 0-based and signals a runtime error for length $\text{xs} \leq i$. Now consider a client indexing a list by its length:

```
client = [1, 2, 3] !! 3
```

Refinement types can be used to impose a false precondition on the error function and detect the out-of-bound crash:

```
error :: { i:String | false } → a
```

But when the two above incompatible definitions coexist is the error due to the *definition* of indexing or due to the *client* of indexing?

This question of who is to blame has no definitive answer in the general case. Gradual liquid type inference is helpful to explore possible resolutions of the error and decide for a suitable solution. To do so, we transform the statically ill-typed program into a gradually well-typed program by giving an imprecise refinement as a precondition to the indexing function:

```
(!!) :: xs:[a] → { i:Int | ?? } → a
```

Fig. 5. The .html output of GuiLT for liquid type exploration.

Using GuiLT, we can explore all potential predicates that can be substituted for `??`. Among the candidates are both $0 \leq i < \text{len } xs$ and $0 < i \leq \text{len } xs$. While the former induces a type error in the client, the latter induces a type error in the definition of `(!!)`.

Figure 4 was generated after running GuiLT on the above specification and code. The result of GuiLT is an .html file where each user specified `??` is turned into a colored button. Pressing a `??` button once highlights all of its usage occurrence, using different colors for each `??` (Figure 4(left)). Pressing it again presents all safe concretizations to scroll through (Figure 4(right)). In the indexing example, three gradual constraints are generated: 1) for the recursive call of indexing, 2) for the unreachable (due to the false precondition) error call, and 3) for the client. Unsurprisingly, the SCs for the recursive call and the client include the incompatible $0 \leq i < \text{len } xs$ and $0 < i \leq \text{len } xs$, respectively. Interestingly, the unreachable constraint enjoys many SCs including $0 \leq i \leq \text{len } xs$ and one given in the right of Figure 4, *i.e.*, $i < \text{len } xs \wedge i == 0$, since the case of indexing 0 from a non-empty list is covered in the first case of the indexing function.

The user explores all SCs with the `<<` and `>>` buttons. In the example, the three SCs are independent, but in many cases SCs can depend on each other (*e.g.*, dependencies in function preconditions). In such cases, pressing a navigation button changes the values of all the dependent occurrences.

When the goal is to replace the `??` with a concrete refinement, going through all SCs can be overwhelming. In the example, there are 22, 10, and 13 SCs for the unreachable, client, and recursive calls, respectively. To accommodate the user, GuiLT generates an alternative interactive .out.html file by which the user can replace the `??` with any SC and observe the generated refinement errors. In Figure 5-left the `??` is replaced with the concrete refinement $0 \leq i < \text{len } xs$, generating an error at the client site. Pressing `>>`, the user explores the next concrete refinement, where $0 < i \leq \text{len } xs$ generates an error in the recursive case of indexing (Figure 5-right). This exposes all the concretizations of `??` that render at least one constraint safe (here 22).

Gradual Types for Error Explanation. This interactive replacement of `??` allows the user to try all the possible concretizations inferred by the algorithm. For each selection the associated type errors are generated and highlighted. This process is effective for two major reasons:

- (1) *Provide all error locations:* The inference algorithm does not arbitrarily blame one program location based on one inferred predicate, instead the user navigates via *all* the potential error locations. Thus, during navigation the user can observe all potential sources of program errors.
- (2) *Predicate Synthesis:* Moreover, the user does not have to come up with the refinements, since the inference procedure synthesizes them. It is only up to the user to pick the adequate refinement that she chooses as the desired specification—which could be fully static, or imprecise.⁵

Quantitative Evaluation. Table 1 summarizes a quantitative evaluation of the indexing example. We run GuiLT with instantiation depth (Depth) 1 and 2, *i.e.*, the size of the conjunctions of the templates we consider (§ 5). The `# ?` column gives the number of imprecise refinements (as added by the user) and `Occs` gives the number each `??` appearing in the generated constraints. Column `Cands` denotes the candidate (*i.e.*, well-typed) templates for each occurrence (§ 5). In the example,

⁵For sound imprecise refinements, runtime support for gradual refinements is needed, as described by Lehmann and Tanter [2017]. Implementing such support in Liquid Haskell is an interesting perspective, outside the scope of this work.

Table 1. Quantitative Evaluation of the Indexing Example

Depth	# ?	Occs	Cands	Sens	Local	Spec	Parts	# γ	SCs	Sols	Time (s)
1	1	[4]	[12*]	[11*]	[11*]	[11*]	4/13	11*	[8,6,6,0]	0	0.51
2	1	[4]	[68*]	[38*]	[34*]	[34*]	4/13	34*	[22,13,13,10]	0	1.61

for depth 2, 68 templates are generated for each occurrence (*i.e.*, [68, 68, 68, 68] simplified as 68* for space). Then, GuiLT decides how many of the templates are sensible (Sens), local (Local), and specific (Spec); here, 38, 34, and 34, respectively. We note that all the local templates are specific, when the static part of the gradual refinement is true (*i.e.*, true \wedge ??, simplified as ??). Moreover, note that most non-local solutions were filtered out by the sensibility check. The constraints were split in 13 partitions (Parts), out of which 4 contained gradual refinements and had to be concretized. Each partition had 34 concretizations (# γ) out of which 22, 13, 13, and 10 were safe concretizations (SCs), respectively. None of these concretizations was common for all the occurrences of the ??, thus the GuiLT reports 0 static solutions (Sols), as expected. Finally, the running time of GuiLT was 1.61 sec.

In short, we observe how the implementation optimizations, as discussed in § 5, allow for possible error reporting using gradual liquid types. With depth of 2 and 4 occurrences of the ?? each with 68 candidates, the exponential inference algorithm would need to run 68⁴ times. In practice, we filter the candidates to 34 and since each occurrence appears in a different partition, we only run the algorithm 34 * 4 times. Next, we use these metrics to evaluate GuiLT as a possible tool for migrating three standard Haskell list libraries to Liquid Haskell.

7 APPLICATION II: MIGRATION ASSISTANCE

As a second application, we use GuiLT's error reporting GUI from § 6 to assist the migration of commonly used Haskell libraries to Liquid Haskell. Inter-language migration to strengthen type safety guarantees is one of the main motivations for gradual type systems [Tobin-Hochstadt and Felleisen 2006]. Our study confirms that gradual liquid type inference provides an effective bridge to migrate programs written in a traditional polymorphic functional language like Haskell to a stronger discipline with refinement types as provided by Liquid Haskell.

Benchmarks. We used GuiLT to migrate three interdependent Haskell list libraries:

- `GHC.List`: provides commonly used list functions available at Haskell's Prelude,
- `Data.List`: defines more sophisticated list functions, *e.g.*, list transposing, and
- `Data.List.NonEmpty`: lifts list functions to a non-empty list data type.

Migration Process. A library consists of a set of function imports and definitions. Each function comes with its Haskell type and may be assigned to a gradual refinement type during migration. Migration is complete when, if possible, all functions are given fully-static refinement types and type check under Liquid Haskell. The process proceeds in four steps:

- Step 1:** run Liquid Haskell to generate a set of type errors,
- Step 2:** fix the errors by *manually* inserting gradual refinements (??),
- Step 3:** use GuiLT to replace ?? with an *automatically* generated SC, and
- Step 4:** go back to **Step 1** until no type errors are reported.

This process is iterative and interactive since refinement errors propagate between imported and client libraries and it is up to the user to decide how to resolve these errors.

Step 1: At the beginning of the migration process the source files given to Liquid Haskell have no refinement type specifications. Still, there are two sources⁶ of refinement type errors: 1) failure to satisfy imported functions preconditions, *e.g.*, in § 6 the error function assumes the false precondition and 2) incomplete patterns, *i.e.*, a pattern-match that might fail at runtime, *e.g.*, `scanr`'s result is matched to a non-empty list.

```
scanr _ q [] = [q]
scanr f q (x:xs) = f x q : qs where qs@(q:_) = scanr f q xs
```

Each time **Step 1** is reiterated, new type errors can occur as a consequence of new function preconditions added in the following steps.

Step 2: The insertion of gradual refinements (??) is left to the user; they can add ?? either in the preconditions of defined functions or postconditions of imported functions, thereby resolving the type errors of **Step 1**. If the user places redundant ??, then they will be solved to the True static refinement. If the user misses some ??, type errors will remain.

Step 3: Once the program gradually typechecks, the user explores the generated SCs and chooses which one to replace ?? with. In our benchmarks, often, the decision is trivial since only one SC coincides for all occurrences of the gradually-refined variable.

Step 4: Depending on the refinement inserted at **Step 3**, new errors might appear in the current or imported libraries; there are three possible scenarios:

- If the user refined the precondition of a function (*e.g.*, `head` requires non-empty lists), then a type error might be generated at clients of the function both inside and outside the library;
- If the user refined the post-condition of a function (*e.g.*, `scanr` always returns non-empty lists), then no error can be generated;
- If the user refined the postcondition of an imported function upon which the function to be verified relies, then the imported function's specification may not be satisfied by its implementation. In this case the user can either *assume* the imported type (thus relying on gradual checking), or update and re-check the imported library.

Evaluation. Table 2 summarizes the migration case study; there are three subtables, one for each library: `GHC.List`, `Data.List`, and `Data.List.NonEmpty`. Within each of these tables, there is a row for every function which requires a refinement type to type check. Rows are collected into groupings of rounds (*i.e.*, **Steps 1 - 4**), where round i lead to refinement type preconditions that trigger type errors in the functions of round $i + 1$. The columns of the table have the same meaning as that of Table 1; all results are for depth 1.

`GHC.List`: Liquid Haskell reported three static errors on the original version of `GHC.List`, *i.e.*, with no user refinement type specifications. The function `errorEmp` is rejected as it is merely a wrapper around the error function; `scanr` and `scanr1` each have incomplete patterns assuming a non-empty list postcondition. `GuiLT` performed bad at all these three initial cases. It was unable to generate any SC for `errorEmp`, since the required false precondition is non local. It required more than one minutes to generate SCs of `scanr` and it timed-out in the case of `scanr1`. The reason for this is that ?? in post-conditions generated dependent set of constraints that made the partitioning optimization (§ 5) useless, despite its importance in the other scenarios. Yet, `GuiLT` was very efficient in the next round. The specification of `errorEmp` introduced errors in ten functions that were fixed using `GuiLT`: in nine cases the generated SCs correspond to exactly one predicate

⁶Termination checking, by default activate in Liquid Haskell, is a third source of type errors in programs without any refinement type specifications. To simplify our case study we deactivated termination checking.

Table 2. Evaluation of Migrations Assistance. Rnd: number of iterations to verify the function. Function: name of the function. # ? : number of ? inserted. Occs: times each ? is used. For each occurrence, we give the number of template candidates (Cands) and how many are specific (Spec). Parts: the number of partitions. For each partition, we give the number of concretizations (# γ) and safe concretizations (SCs). Sols: number of static solutions found. Time: time in sec.

Rnd	Function	# ?	Occs	Cands	Spec	Parts	# γ	SCs	Sols	Time (s)
GHC.List (56 functions defined and verified)										
1 st	errorEmp	1	[1]	[[5]]	[[4]]	1/4	[4]	[0]	0	1
	scanr	1	[4]	[6*]	[5*]	1/4	[625]	[125]	1	98
	scanr1	2	[6, 4]	[5*,12*]	[4*,5*]	1/5	[2.5M]	N/A	N/A	N/A
2 nd	head	1	[1]	[[5]]	[[4]]	1/3	[4]	[1]	1	1
	tail	1	[1]	[[5]]	[[4]]	1/5	[4]	[1]	1	2
	last	1	[2]	[5*]	[4*]	2/4	4*	[4,1]	1	2
	init	1	[3]	[5*]	[4*]	2/8	[16,4]	[16,1]	1	2
	fold1	1	[3]	[5*]	[4*]	2/5	[16,4]	[16,1]	1	2
	foldr1	1	[1]	[[5]]	[[4]]	1/2	[4]	[1]	1	2
	(!!)	2	[4,4]	[47*,12*]	[8*,6*]	4/9	48*	[28,48,24,48]	8	16
	cycle	1	[2]	[5*]	[4*]	2/6	4*	[4,1]	1	1
	maximum	1	[2]	[5*]	[4*]	2/4	4*	[4,1]	1	1
	minimum	1	[2]	[5*]	[4*]	2/4	4*	[4,1]	1	1
Data.List (115 functions defined and verified)										
1 st	maximumBy	1	[4]	[5*]	[4*]	2/5	[64,4]	[64,1]	1	4
	minimumBy	1	[4]	[5*]	[4*]	2/5	[64,4]	[64,1]	1	4
	transpose	1	[5]	[5*]	[4*]	2/9	[256,4]	[28,3]	1	92
	genIndex	2	[6,6]	[2*,5*]	[1*,4*]	6/12	4*	[4,4,1,3,1,4]	0	2
Data.List.NonEmpty (57 functions defined and verified)										
1 st	fromList	1	[2]	[5*]	[4*]	2/10	4*	[4,1]	1	3
	cycle	3	[1,1,1]	[[12],[2],[5]]	[[5],[1],[1]]	1/2	[5]	[0]	0	3
	-toList	1	[1]	[[5]]	[[4]]	1/2	[4]	[3]	3	3
	(!!)	1	[4]	[10*]	[9*]	4/11	9*	[9,9,3,5]	1	4
	transpose	2	[1,1]	[[6],[7]]	[[1],[1]]	1/6	[1]	[0]	0	3
2 nd	cycle	2	[1,1]	[[12],[5]]	[[5],[1]]	1/2	[5]	[2]	2	3
	groupAll	2	[1,0]	[[12],[0]]	[[5],[0]]	1/7	[5]	[2]	2	3
	groupBy	2	[1,0]	[[12],[0]]	[[5],[0]]	1/3	[5]	[1]	1	4
	groupBy1	2	[1,0]	[[12],[0]]	[[5],[0]]	1/3	[5]	[1]	1	3
	lift	1	[1]	[[6]]	[[5]]	1/6	[5]	[2]	2	3
	inits	1	[1]	[[6]]	[[5]]	1/5	[5]	[2]	2	5
	tails	2	[1,1]	[[5],[6]]	[[4],[5]]	1/5	[20]	[6]	6	5
	scan1	1	[1]	[[6]]	[[5]]	1/5	[5]	[2]	2	3
	scanl1	2	[1,1]	[[6],[5]]	[[5],[4]]	1/5	[20]	[6]	6	6
	insert	1	[1]	[[12]]	[[5]]	1/5	[5]	[2]	2	3
	transpose	1	[1]	[[12]]	[[5]]	1/5	[5]	[2]	2	3
	3 rd	reverse	2	[1,1]	[[5],[12]]	[[4],[5]]	2/3	[5,4]	[1,2]	2
sort		2	[1,1]	[[12],[5]]	[[5],[4]]	2/3	[5,4]	[1,2]	2	3
sortBy		1	[1]	[[12]]	[[5]]	1/2	[5]	[1]	1	2

that was used to replace the ??. For the case of (!!), GuiLT generated 8 correct predicates, out of which we manually picked the most general one.

In short, out of the 56 functions defined, 13 required refinement type specifications, since the rest were typed with the true-default types generated by Liquid Haskell. GuiLT was unusable in 2 cases, but synthesized exactly one, the correct, predicate in 9 of them. The two cases where GuiLT

was unusable are due to many dependencies within the same partition; as future work, we should explore whether it is possible to devise a more advanced partitioning scheme.

`Data.List`: Migration of `Data.List` only required one round. Four functions errored due to incomplete patterns or violation of preconditions of functions imported from previously verified `GHC.List`. Unsurprisingly, `GuiLT` was unable to find any SCs for `genIndex`, a generic variant of (`!!`) that indexes lists using any integral (instead of integer) index, because it lacks arithmetic templates for integrals. To complete migration, the user needs to either manually provide the refinement type of `genIndex` or appropriately extend the set of templates.

`Data.List.NonEmpty`: Migration was more interesting for the `Data.List.NonEmpty` library that manipulates the data type `NonEmpty a` of non-empty lists. The first round exposed that `fromList` requires the non-empty precondition. The `GHC.List` function `cycle` has a non-empty precondition, thus lifted to non-empty lists does not type check.

```
cycle :: NonEmpty a → NonEmpty a
cycle = fromList . List.cycle . toList
```

To migrate `cycle`, we first gradually refined the result type of `toList :: NonEmpty a → { xs:[a] | ?? }` for which `GuiLT` suggested the single static refinement of $0 < \text{len } xs$.

Verification of non-empty list indexing calls requires invariants that relate the lengths of the empty and non-empty lists. Similarly, `GuiLT` finds SCs only after predicate templates that express such invariants are added. In general, to aid migration on user-defined structures, `GuiLT` requires the definition of domain-specific templates.

On the second round, the non-empty precondition of `fromList` triggers errors to eleven clients (step 4a), all of which call functions that return (non-provably) non-empty lists. For example, `inits` lifts `List.inits` to non-empty lists.

```
inits = fromList . List.inits . toList
```

To migrate such functions, we first assume an unknown specification, for example:

```
assume List.inits :: [a] → { o:[a] | ?? }
```

We then use `GuiLT` to solve the unknown specification to $0 < \text{len } o$. At this point, we can either update the imported function with the discovered static specification and recheck the library (step 4c), or stick to gradual checking and stay with the unknown specification for `inits`.

The function `cycle` reappears in the second round, due to the new precondition of `fromList`. Since the imported `List.cycle` already has a precondition $\{i:[a] \mid 0 < \text{len } i\}$, at this round we further strengthened the existing precondition with the imprecise refinement:

```
assume List.cycle :: i:{ [a] | 0 < len i ∧ ?? } → { o:[a] | ?? }
```

This is the only case in our experiments where we used a gradual refinement with a static part and thus the only case in which some local templates are rejected as non-specific (here 3 out of 4 local templates are not specific).

Finally, we use `GuiLT` to derive higher-order specifications. The `lift` function lifts a list transformation to non-empty lists, `lift f = fromList . f . toList`, and comes with a comment that “If the provided function returns an empty list, this will raise an error.”. Alerted by this comment, we use a `??` in higher-order position:

```
lift :: (i:[a] → { o:[b] | ?? }) → NonEmpty a → NonEmpty b
```

`GuiLT` produces two static solutions $0 < \text{len } o$ and $\text{len } i == \text{len } o$. We choose the second, which leads to type errors in three clients, which are resolved in later rounds.

To sum up, GuiLT indeed is aiding Haskell to Liquid Haskell migration of real libraries, since the user can place the initial $??$ s and choose from the suggested SCs, instead of writing specifications from scratch. Often times, there is only one possible SC coinciding to all concretizations, thus the choice is trivial. When no suggestions are generated, e.g., `errorEmp` or `genIndex`, the user can fall back to the standard verification process.

Our case studies indicate that GuiLT can indeed be used to aid program migration and opens up future work for practical program migration. Our suggested search algorithm is optimized, but still exponential, leading to scalability problems, as indicated by the inability to generate suggestions in two cases. It seems plausible that in practice, users would provide partial refinements that could effectively prune the search space. This should be further evaluated empirically. As another future direction, automatic insertion of $??$ s would further reduce the required user input. Theorem 4.3 suggests a complete way to automatically insert $??$ s so that the migrating libraries gradually type check. The types suggested by this theorem are very imprecise, yet they provide a sound starting point, upon which one can increase precision so long as the libraries type check.

8 RELATED WORK

Liquid Types. Dependent types allow arbitrary expressions at the type level to express theorems on programs, while theorem proving is simplified by various automations ranging from tactics (e.g., Coq [Bertot and Castéran 2004], Isabelle [Wenzel 2016]) to external SMT solvers (e.g., F* [Swamy et al. 2016]). Liquid types [Rondon et al. 2008] restrict the expressiveness of the type specifications to decidable fragments of logic to achieve decidable type checking and inference.

Gradual Refinement Types. Several refinement type systems mix static verification with runtime checking. Hybrid types [Knowles and Flanagan 2010] use an external prover to statically verify subtyping when possible, otherwise a cast is automatically inserted to defer checking at runtime. Soft contract verification [Nguyen et al. 2014, 2018; Tobin-Hochstadt and Horn 2012] works in the other direction, statically verifying contracts wherever possible, and otherwise leaving unverified contracts for checking at runtime. Ou et al. [2004] allow the programmer to explicitly annotate whether an assertion is verified at compile- or runtime. Manifest contracts [Greenberg et al. 2010] formalize the metatheory of refinement typing in the presence of dynamic contract checking. Lehmann and Tanter [2017] developed the first gradual refinement type system, which adheres to the refined criteria of Siek et al. [2015]. None of these systems support inference; on the contrary, because refinements can be arbitrary, inference is impossible in these systems. Here we restrict gradual refinements to a finite set of predicates, in order to achieve inference by adaptation of the liquid inference procedure.

Gradual Type Inference. Several approaches have been developed to combine type inference and gradual types. Siek and Vachharajani [2008] infer gradual types using unification, while Rastogi et al. [2012] exploit type inference to improve the performance of gradually-typed programs. Garcia and Cimini [2015] lift an inference algorithm from a core system to its gradual counterpart, by ignoring the unification constraints imposed by gradual types. In contrast, in gradual liquid inference the constraints imposed by gradual refinements cannot be ignored, since unlike Hindley-Milner inference, the liquid algorithm starts from the strongest solution for the liquid variables and uses constraints to iteratively weaken the solution. Our work is the first gradual inference work to be systematically derived from the static algorithm based on the Abstracting Gradual Typing approach (AGT) [Garcia et al. 2016], and proven to satisfy the static gradual guarantee [Siek et al. 2015].

Error Reporting. Properly localizing, diagnosing and reporting type errors is a long-standing challenge, especially for inference algorithms. Since the seminal work of Wand [1986], which keeps

track of all unification steps in order to help debugging, many algorithms have been proposed to better assist programmers. A large variety of techniques has been explored, including slicing [Haack and Wells 2003; Tip and Dinesh 2001], heuristics [Zhang et al. 2015], SMT constraint solving [Pavlinovic et al. 2014], counter-example generation [Nguyen and Horn 2015; Seidel et al. 2016], machine learning [Seidel et al. 2017], and other search-based approaches, such as Seminal [Lerner et al. 2007] and counterfactual change inference [Chen and Erwig 2018].

Apart from Seidel et al. [2016], which focuses on identifying dynamic witnesses for static refinement type errors, none of the above approaches target refinement types. In this work, we uncover a novel application of gradual typing for error explanation, by observing that the notion of concretizations of gradual types that stem from AGT, embed useful justifications (or lack thereof) that can be exploited to identify refinement errors and guide migration.

Gradual Program Migration. Recently, Campora et al. [2018] exploit variational typing [Chen et al. 2014] to assist programmers in making their program more static, by analyzing the impact of replacing unknown types with static types. Variational typing greatly reduces the complexity of exploring all possible combinations. Though the considered type system is simple, it seems likely that the approach could be extended to refinements and combined with our technique. This might be an effective way to address the scalability issues we have encountered with our implementation approach in certain scenarios.

9 CONCLUSION

This paper makes the novel observation that gradual inference based on abstract interpretation can be fruitfully exploited to assist in explaining type errors and migrating programs to a stronger typing discipline. We develop this intuition in the context of refinement types, yielding a novel integration of liquid type inference and gradual refinements. Gradual liquid type inference computes possible concretizations of unknown refinements in order for a program to be well-typed. In addition to laying down the theoretical foundations of gradual liquid type inference, we provide an implementation integrated with Liquid Haskell. Thanks to a number of heuristics and optimizations, the current implementation of the interactive tool GuiLT proves useful for migrating existing Haskell libraries to the stronger discipline of Liquid Haskell. Our experience however also shows that enhancing the scalability of our prototype further is necessary; we believe that variational typing could be particularly helpful towards that goal. Similarly, integrating prior work on ranking suggested type error sources would help our inference algorithm suggest the most relevant concretizations first. Finally, we conjecture that the idea of using gradual typing for error reporting and migration generalizes to other typing disciplines.

ACKNOWLEDGMENTS

We thank Ron Garcia for his very helpful feedback on earlier versions of this paper. We thank the anonymous reviewers for their helpful constructive feedback, which allowed us to enhance the presentation of this work.

REFERENCES

- Yves Bertot and Pierre Castéran. 2004. *Coq'Art: The Calculus of Inductive Constructions*. Springer Verlag.
- John Peter Campora, Sheng Chen, Martin Erwig, and Eric Walkingshaw. 2018. Migrating Gradual Types. *PACMPL (POPL)*.
- Sheng Chen and Martin Erwig. 2018. Systematic identification and communication of type errors. *JFP* 28.
- Sheng Chen, Martin Erwig, and Eric Walkingshaw. 2014. Extending Type Inference to Variational Programs. In *TOPLAS*.
- Bruno Courcelle and Joost Engelfriet. 2012. *Graph Structure and Monadic Second-Order Logic: A Language-Theoretic Approach* (1st ed.). Cambridge University Press.
- Tim Freeman and Frank Pfenning. 1991. Refinement Types for ML. In *PLDI*.

- Ronald Garcia and Matteo Cimini. 2015. Principal Type Schemes for Gradual Programs. In *POPL*.
- Ronald Garcia, Alison M. Clark, and Éric Tanter. 2016. Abstracting Gradual Typing. In *POPL*.
- Michael Greenberg, Benjamin C. Pierce, and Stephanie Weirich. 2010. Contracts Made Manifest. In *POPL*.
- Christian Haack and Joe Wells. 2003. Type Error Slicing in Implicitly-Typed Higher-Order Languages. In *ESOP*.
- K.W. Knowles and C. Flanagan. 2010. Hybrid type checking. In *TOPLAS*.
- Nico Lehmann and Éric Tanter. 2017. Gradual Refinement Types. In *POPL*.
- Benjamin Lerner, Matthew Flower, Dan Grossman, and Craig Chambers. 2007. Searching for Type-Error Messages. In *PLDI*.
- Phuc C. Nguyen and David Van Horn. 2015. Relatively complete counterexamples for higher-order programs. In *PLDI*.
- Phuc C. Nguyen, Sam Tobin-Hochstadt, and David Van Horn. 2014. Soft Contract Verification. In *ICFP*.
- Phuc C. Nguyen, Sam Tobin-Hochstadt, and David Van Horn. 2018. Soft Contract Verification for Higher-order Stateful Programs. In *POPL*.
- Xinming Ou, Gang Tan, Yitzhak Mandelbaum, and David Walker. 2004. Dynamic Typing with Dependent Types, Jean-Jacques Levy, Ernst W. Mayr, and John C. Mitchell (Eds.). *IFIP*.
- Zvonimir Pavlinovic, Tim King, and Thomas Wies. 2014. Finding minimum type error sources. In *OOPSLA*.
- Aseem Rastogi, Avik Chaudhuri, and Basil Hosmer. 2012. The ins and outs of gradual type inference. In *POPL*.
- P. Rondon, M. Kawaguchi, and R. Jhala. 2008. Liquid Types. In *PLDI*.
- Eric Seidel, Ranjit Jhala, and Westley Weimer. 2016. Dynamic witnesses for static type errors (or, ill-typed programs usually go wrong). In *ICFP*.
- Eric L. Seidel, Huma Sibghat, Kamalika Chaudhuri, Westley Weimer, and Ranjit Jhala. 2017. Learning to blame: localizing novice type errors with data-driven diagnosis. *OOPSLA (2017)*.
- Jeremy Siek and Walid Taha. 2006. Gradual Typing for Functional Languages. In *Scheme and Functional Programming Workshop*.
- Jeremy G. Siek and Manish Vachharajani. 2008. Gradual Typing with Unification-based Inference. In *Dynamic Languages Symposium*.
- Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, and John Tang Boyland. 2015. Refined Criteria for Gradual Typing. In *SNAPL*.
- N. Swamy, C. Hrițcu, C. Keller, A. Rastogi, A. Delignat-Lavaud, S. Forest, K. Bhargavan, C. Fournet, P. Y. Strub, M. Kohlweiss, J. K. Zinzindohoue, and S. Zanella-Béguelin. 2016. Dependent Types and Multi-Monadic Effects in F*. In *POPL*.
- Frank Tip and T. B. Dinesh. 2001. A slicing-based approach for locating type error. In *TOSEM*.
- Sam Tobin-Hochstadt and Matthias Felleisen. 2006. Interlanguage migration: From scripts to programs. In *DLS*.
- Sam Tobin-Hochstadt and David Van Horn. 2012. Higher-order Symbolic Execution via Contracts. In *OOPSLA*.
- Niki Vazou, Eric L. Seidel, and Ranjit Jhala. 2014a. LiquidHaskell: Experience with Refinement Types in the Real World. In *Haskell*.
- Niki Vazou, Eric L. Seidel, Ranjit Jhala, Dimitrios Vytiniotis, and Simon Peyton-Jones. 2014b. Refinement Types for Haskell. In *ICFP*.
- Niki Vazou, Éric Tanter, and David Van Horn. 2018a. Gradual Liquid Types. *ArXiv e-prints* (2018). arXiv:1807.02132
- Niki Vazou, Anish Tondwalkar, Vikraman Choudhury, Ryan Scott, Ryan Newton, Philip Wadler, and Ranjit Jhala. 2018b. Refinement Reflection: Complete Verification with SMT. In *POPL*.
- Mitchell Wand. 1986. Finding the Source of Type Errors. In *POPL*.
- Makarius Wenzel. 2016. The Isabelle System Manual. (2016). <https://www.cl.cam.ac.uk/research/hvg/Isabelle/dist/Isabelle2016-1/doc/system.pdf>
- Danfeng Zhang, Andrew C. Myers, Dimitrios Vytiniotis, and Simon Peyton-Jones. 2015. Diagnosing Type Errors with Class. In *PLDI*.