

Toward Hole-Driven Development in Liquid Haskell

PATRICK REDMOND, University of California, Santa Cruz, USA

GAN SHEN, University of California, Santa Cruz, USA

LINDSEY KUPER, University of California, Santa Cruz, USA

Liquid Haskell is an extension to the Haskell programming language that adds support for *refinement types*: data types augmented with SMT-decidable logical predicates that refine the set of values that can inhabit a type. Furthermore, Liquid Haskell’s support for *refinement reflection* enables the use of Haskell for general-purpose mechanized theorem proving. A growing list of large-scale mechanized proof developments in Liquid Haskell take advantage of this capability. Adding theorem-proving capabilities to a “legacy” language like Haskell lets programmers directly verify properties of real-world Haskell programs (taking advantage of the existing highly tuned compiler, run-time system, and libraries), just by writing Haskell. However, more established proof assistants like Agda and Coq offer far better support for interactive proof development and insight into the proof state (for instance, what subgoals still need to be proved to finish a partially-complete proof). In contrast, Liquid Haskell provides only coarse-grained feedback to the user — either it reports a type error, or not — unfortunately hindering its usability as a theorem prover.

In this paper, we propose improving the usability of Liquid Haskell by extending it with support for Agda-style *typed holes* and interactive editing commands that take advantage of them. In Agda, typed holes allow programmers to indicate unfinished parts of a proof, and incrementally complete the proof in a dialogue with the compiler. While GHC Haskell already has its own Agda-inspired support for typed holes, we posit that typed holes would be especially powerful and useful if combined with Liquid Haskell’s refinement types and SMT automation. We discuss how typed holes might work in Liquid Haskell, and we consider possible implementation approaches and next steps.

1 INTRODUCTION

Refinement types [Rushby et al. 1998; Xi and Pfenning 1998] are data types augmented with logical predicates, called *refinement predicates*, that restrict, or *refine*, the set of values that can inhabit a type. Depending on the expressivity of the language used for the refinement predicates, programmers can specify rich program properties using refinement types, sometimes at the expense of the decidability of type checking. *Liquid types* [Rondon et al. 2008; Vazou et al. 2014] are refinement types that avoid such undecidability by restricting refinement predicates to a decidable logic that can be checked by an off-the-shelf SMT solver. Liquid types are beginning to make their way into general-purpose, industrial-strength programming languages through tools such as *Liquid Haskell* [Vazou et al. 2014], an extension to the Haskell programming language that adds support for liquid types.

Liquid Haskell’s support for *refinement reflection* [Vazou et al. 2018, 2017b] lets programmers use refinement types to specify arbitrary “extrinsic” properties that can relate multiple functions, and then prove those properties by writing Haskell programs to inhabit the specified refinement types. Thanks to refinement reflection and Liquid Haskell’s *proof combinators* [Vazou et al. 2018, 2017b], a Liquid Haskell programmer can leverage the Curry-Howard correspondence and prove arbitrary (undecidable) properties by *inhabiting types with programs*, just as she might do with a dependent-types-based proof assistant like, say, Agda [Norell 2008], thus “turning Haskell into a theorem prover” [Vazou et al. 2017b] and enabling general-purpose mechanized theorem proving in Liquid Haskell. A growing number of large-scale mechanized proof developments in Liquid Haskell [Liu et al. 2020; Parker et al. 2019; Vazou et al. 2017a] make use of this capability.

As previous work [Parker et al. 2019; Vazou et al. 2017a] has discussed, deciding to use Liquid Haskell for mechanized proofs involves a trade-off. On one hand, with the addition of refinement reflection, Liquid Haskell checks many of the boxes on a modern proof engineer’s wish list: Curry-Howard-style proving by programming, SMT automation, and above all, seamless integration with an existing, general-purpose language (with no additional code extraction step required). Attaching theorem-proving capabilities to a “legacy” language like Haskell lets programmers directly verify properties of real-world Haskell programs (taking advantage of the existing highly tuned compiler, run-time system, and libraries), just by writing Haskell [Vazou 2016]. On the other hand, traditional Curry-Howard-based proof assistants like Agda [Norell 2008] or Coq [Bertot and Castan 2010] not only rest on a firm theoretical foundation that Liquid Haskell arguably lacks, but they also currently offer far better support for *interactive* proof development and insight into the proof state (for instance, what subgoals still need to be proved to finish a partially-complete proof). Liquid Haskell provides only coarse-grained feedback to the programmer: either it reports a type error, which means there is still more work to do to complete a proof, or it does not, which means the proof is done. While such coarse-grained feedback might be fine for intrinsic refinement type specifications or short proofs, the increasing use of Liquid Haskell for large-scale proof development motivates the need for better interactive proof development tools.

In Agda, the interactive proof development experience is enabled in large part by a feature called *typed holes*. Typed holes allow Agda programmers to indicate parts of a proof that they need help with filling in. In an interactive environment (such as an IDE), the programmer can use holes to interact with the proof checker by, for example, asking for the type expected by the hole (that is, the proposition that needs to be proven to fill in that hole). Additional IDE commands can assist the programmer in filling in holes [The Agda Team 2021]. For example, if the programmer has partially filled a hole with an expression whose *return* type matches the type expected by the hole, then an IDE command can generate new holes indicating the types of the arguments of the expression. For a hole on the right-hand side of a definition, another IDE command can case-split on a pattern variable and generate new holes for each case. The result is that, even though Agda favors a style in which programmers explicitly write proof terms (compared to, for instance, proof assistants like Coq or Isabelle [Wenzel et al. 2008], in which programmers more typically use a tactic language or proof automation to indirectly construct proof objects [Ringer et al. 2019]), much of the effort of this explicit proof-writing in Agda is carried out with helpful machine assistance.

In this paper, we propose improving the interactivity of Liquid Haskell by extending it with support for Agda-style typed holes and interactive editing commands that take advantage of them. In fact, vanilla GHC Haskell already has support for typed holes [GHC Team 2014], a feature that was itself inspired by Agda. However, we hypothesize that typed holes could be especially powerful in combination with refinement types and SMT automation. In the rest of this paper, we give an overview of refinement types and Liquid Haskell (§2), and describe how typed holes might work in Liquid Haskell (§3). We conclude with a discussion of possible implementation approaches and next steps (§4).

2 BACKGROUND: REFINEMENT TYPES AND LIQUID HASKELL

Refinement types [Rushby et al. 1998; Xi and Pfenning 1998] let programmers specify data types augmented with logical predicates, called *refinement predicates*, that restrict the set of values that can inhabit the type. Depending on the expressivity of the language of refinement predicates, programmers can specify rich program properties using refinement types, sometimes at the expense of the decidability of type checking. Liquid Haskell avoids that problem by restricting refinement predicates to an SMT-decidable logic [Rondon et al. 2008; Vazou et al. 2014]. For example, in Liquid Haskell we could define the type of even integers by refining the Haskell type `Int` using the

refinement type $\{ v:\text{Int} \mid v \bmod 2 == 0 \}$, where $v \bmod 2 == 0$ is the refinement predicate and $v:\text{Int}$ binds the name v for values of type Int that appear in the refinement predicate. One could define an analogous refinement type for odd integers, and then write a Liquid Haskell function for adding them:

```
type EvenInt = { v:Int | v mod 2 == 0 }
type OddInt  = { v:Int | v mod 2 == 1 }

oddAdd :: OddInt -> OddInt -> EvenInt
oddAdd x y = x + y
```

The type `OddInt` of the arguments to `oddAdd` expresses the *precondition* that x and y will be odd, and the return type `EvenInt` expresses the *postcondition* that $x + y$ will evaluate to an even number. Liquid Haskell automatically proves that such postconditions hold by generating verification conditions that are checked at compile time by the underlying SMT solver, Z3 [de Moura and Bjørner 2008]. If the solver finds a verification condition to be invalid, typechecking fails. If the return type of `oddAdd` had been `OddInt`, for instance, the above code would fail to typecheck.

Aside from preconditions and postconditions of individual functions, Liquid Haskell makes it possible to verify *extrinsic properties* that are not specific to any particular function’s definition. For example, the type of `sumOdd` below expresses the extrinsic property that the sum of an odd and an even number is an odd number:

```
sumOdd :: x : OddInt -> y : EvenInt -> { _:Proof | (x + y) mod 2 == 1 }
sumOdd _ _ = ()
```

Here, `sumOdd` is a Haskell function that returns a *proof* that the sum of x and y is odd. (In Liquid Haskell, `Proof` is a type alias for Haskell’s `()` (unit) type, but refined with constraints that express the property we wish to prove.) Because the proof of this particular property is easy for the SMT solver to carry out automatically, the body of the `sumOdd` function need not say anything but `()`, the sole inhabitant of the unit type. In general, however, programmers can specify arbitrary extrinsic properties in refinement types, including properties that refer to arbitrary Haskell functions via *refinement reflection* [Vazou et al. 2017b]. The programmer can then prove those extrinsic properties by writing Haskell programs that inhabit those refinement types, using Liquid Haskell’s provided *proof combinators* – with the help of the underlying SMT solver to simplify the construction of these proofs-as-programs [Vazou et al. 2018, 2017b].

Liquid Haskell thus occupies a unique position at the intersection of SMT-based program verifiers such as Dafny [Leino 2010], and proof assistants that leverage the Curry-Howard correspondence such as Coq and Agda. A Liquid Haskell program can consist of both application code like `oddAdd` (which runs at execution time, as usual) and verification code like `sumOdd` (which only “runs” at compile time), but, pleasantly, both are just Haskell programs, albeit annotated with refinement types. Being based on Haskell enables programmers to gradually port code from vanilla Haskell to Liquid Haskell, adding richer specifications to code as they go. Furthermore, verified Liquid Haskell libraries can be used directly in arbitrary Haskell programs, letting programmers take advantage of formally-verified components from unverified code written in an industrial-strength, general-purpose language. Finally, unlike Coq or Agda, which require an executable implementation to be *extracted* from the code written in the proof assistant’s vernacular language, Liquid Haskell enables proving properties both *in* and *about* the code to be executed, resulting in immediately executable verified code with no need for a further extraction step.

2.1 Why do Liquid Haskell programmers need typed holes?

Because the proof of `sumOdd` above can be carried out automatically by the SMT solver, there is no need to write anything but `()` as the body of `sumOdd`. If SMT automation makes proofs so easy to write, why would Liquid Haskell programmers need typed holes, anyway? Our answer is that, thanks to refinement reflection, large-scale proof development [Liu et al. 2020; Parker et al. 2019; Vazou et al. 2017a] is now possible in Liquid Haskell, which in turn motivates the need for better machine support in proof construction.

For example, Liu et al. [2020] used Liquid Haskell to develop a framework for programming distributed applications based on *conflict-free replicated data types* (CRDTs) [Shapiro et al. 2011]. CRDTs are data structures whose operations must satisfy certain mathematical properties that can be leveraged to ensure *strong convergence* [Shapiro et al. 2011], meaning that replicas are guaranteed to have equivalent states given that they have received and applied the same *unordered* set of update operations. For example, consider a data structure representing the contents of a shopping cart, replicated across data centers in Seattle, Frankfurt, Mumbai, and Tokyo for fault tolerance and data locality. Due to network partitions and message latency, updates to the cart’s contents may arrive in an arbitrary order at each data center, but strong convergence ensures that under reliable message delivery assumptions, the replicas will eventually agree. Liu et al. [2020]’s framework required proofs that CRDT operations commute. While these proofs were trivial for some simple CRDTs, more sophisticated ones required on the order of thousands of lines of Liquid Haskell (and hours of solver time) [Liu et al. 2020, Table 3], constituting one of the largest Liquid Haskell proof developments to date. Liu et al. describe the effort as “strenuous”, highlighting the need for tools that enable visibility into the in-progress proof state.

3 WHAT COULD WE DO WITH TYPED HOLES IN LIQUID HASKELL?

As a simple mock-up of how typed holes might work in Liquid Haskell, consider a `listLength` function that we want to prove has the same behavior as Liquid Haskell’s built-in `len`. One way to state this property is as an intrinsic property on the type of `listLength` itself, as follows:

```
listLength :: xs:_ -> { v : Nat | v == len xs }
listLength []     = _0
listLength (y:ys) = 1 + _1
```

Here, `_0` and `_1` are holes, indicating that the programmer needs help filling in these parts of the code. Given the refinement type annotation on `listLength`’s return type, together with the context that `xs == []` in this case, we could expect that the type checker would be able to tell us that the type of the first hole is `{ v : Nat | v == len [] }`. The second hole could have type `{ v : Nat | 1 + v == len (y:ys) }`, but the type checker should be able to take advantage of the underlying SMT solver’s automated reasoning about linear arithmetic to conclude that the type of the second hole is `{ v : Nat | v == len (y:ys) - 1 }`. So we have:

```
_0 : { v : Nat | v == len [] }
_1 : { v : Nat | v == len (y:ys) - 1 }
```

which can be further simplified to

```
_0 : { v : Nat | v == 0 }
_1 : { v : Nat | v == len ys }
```

at which point we know we need `0` and `listLength ys` for `_0` and `_1` respectively.

The preceding example shows how a programmer might use typed holes for writing code with intrinsic refinement type specifications, but how about for writing extrinsic proofs? Instead of giving a refinement type to `listLength`, we can express the property that relates `listLength` and `len` extrinsically, as follows:

```
listLength :: [a] -> Nat
listLength [] = 0
listLength (y:ys) = 1 + listLength ys

listLengthProof :: xs:_ -> { _:Proof | listLength xs == len xs }
listLengthProof = _0
```

The body of `listLengthProof` is a hole, written as `_0`. When we attempt to compile this code, since the property we are trying to show relates the return values of `listLength` and `len`, Liquid Haskell might start by suggesting a case split on one of the functions:

```
Found hole `_0' of type `xs:[a] -> { _:Proof | listLength xs == len xs }'.
  Consider a case split as in the body of `listLength'.
```

The case split could be automated with a keystroke, or manually completed. In either case, the result is a partially completed proof with two holes:

```
listLengthProof :: xs:_ -> { _:Proof | listLength xs == len xs }
listLengthProof [] = _0
listLengthProof (y:ys) = _1
```

The first hole, `_0`, has expected type `{ _:Proof | listLength xs == len xs }`. In `_0`'s context, Liquid Haskell could substitute `[]` in place of `xs` in `listLength xs` and then evaluate the call to `listLength` to arrive at `0`, with similar steps to evaluate `len xs` to `0`, to complete the proof. This process could be automated after Liquid Haskell encounters the hole by checking whether `()` completes the branch's proof in the background and then issuing a message:

```
Found hole `_0' of type `{ _:Proof | xs == [] && listLength xs == len xs }'.
  This can be completed with `()'.
```

The replacement of `_0` with `()` could again be automated with a keystroke, or manually completed. Even better, completing branches with `()` could be folded into automated case-splitting such that each branch is generated with either `()` or a new hole, according to whether more work is needed.

The second hole, `_1`, has expected type `{ _:Proof | listLength xs == len xs }`. In `_1`'s context, Liquid Haskell could substitute `(y:ys)` in place of `xs` in `listLength xs` and then evaluate the call to `listLength xs` to arrive at `1 + listLength ys`, with similar steps for `len`, to generate a new goal `{ _:Proof | 1 + listLength ys == 1 + len ys }`. This could be automated after Liquid Haskell encounters the hole and finds that replacing it with `()` does not complete the proof. Liquid Haskell could then issue a message:

```
Found hole `_1' of type `{ _:Proof | xs == (y:ys) && listLength xs == len xs }'.
  Conclusion expands to `1 + listLength ys == 1 + len ys',
  which is simplified to `listLength ys == len ys'.
```

This message does not suggest any action directly, but it does show us an intermediate goal which if proved would be sufficient to show the overall conclusion `listLength xs == len xs`. Users may recognize this goal as the *inductive assumption*. Replacing `_1` with `listLengthProof ys` could potentially be automated with a keystroke in this simple example, although it is unlikely to be quite this obvious in nontrivial proofs.

4 NEXT STEPS

In this section, we consider ways to implement the behavior sketched in §3 by extending the existing Liquid Haskell implementation, and discuss how the proposed behavior differs from what GHC Haskell's typed holes already offer.

Existing Liquid Haskell implementation. In 2020, Liquid Haskell transitioned from an external tool called `liquid` to one implemented as a GHC plugin [Jhala 2020]. Following the GHC build sequence, the `LiquidHaskell` plugin checks refinement types after GHC completes its typechecking. The plugin obtains a refinement type for every top-level binding and component expression through a combination of annotations and inference.

Existing (unpublished) work on synthesizing Haskell expressions using Liquid Haskell has already seen the plugin internals set up to capture holes of the form `_goal`. When a hole is captured, Liquid Haskell tracks the logical environment at the location of the hole. This environment has the potential to be used as a Coq-style proof-writing environment so that the user can learn what information is missing. Leveraging this existing work, a path forward might be:

- Collect examples of hole placements in nontrivial proofs.
- Observe the information tracked in the logical environment at those holes.
- Experiment with ways to convey that information concisely to the user in compiler or IDE output.

With this foundation, we'll be able to decide how to proceed. Options include:

- Transforming the logical environment with SMT-aided simplification, as in the `listLength` and `listLengthProof` examples above.
- Completing the interactivity loop from compiler output (including on-disk suggestion metadata) to editor commands which trigger automated changes to the user's code.

Comparison with GHC's typed holes. Typed holes in GHC currently communicate the inferred principal type of the expression in an error message to the user. This is useful, and has inspired some Haskell programmers to advocate for a "hole-driven development" style of programming. However, in the case of the incomplete `listLength` function and the incomplete `listLengthProof` from §3, the principal types are obvious, and the real utility would come from analysis of the refinement type context to generate the constrained refinement type and synthesize expressions that fit.

Recent work [Gissurarson 2018] has extended GHC's typed hole support to combine the principal type information available from GHC's inference with the types of in-scope binders, such as those imported from library modules. This work suggests both binders with exact type matches and also ones with more general types that fit. However, for a program like `listLength` above, without taking refinement type information into account, there may be many completions that are well-typed but behaviorally incorrect. In the case of `listLengthProof`, there is simply not sufficient information at the level of Haskell types to guide suggestions or synthesis meaningfully. In fact, Gissurarson [2018] suggests integration with Liquid Haskell as a possible solution to the problem of principal types in GHC being too vague to suggest useful completions. Our hypothesis is that integrating typed holes into Liquid Haskell will allow more knowledge of the programmer's intent, expressed in refinement types, to drive suggestions and code synthesis.

REFERENCES

- Yves Bertot and Pierre Castran. 2010. *Interactive Theorem Proving and Program Development: Coq'Art The Calculus of Inductive Constructions* (1st ed.). Springer Publishing Company, Incorporated.
- Leonardo de Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems*, C. R. Ramakrishnan and Jakob Rehof (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 337–340.
- GHC Team April 2014. Typed Holes. https://downloads.haskell.org/~ghc/7.8.1/docs/html/users_guide/typed-holes.html.
- Matthias Páll Gissurarson. 2018. Suggesting Valid Hole Fits for Typed-Holes (Experience Report). In *Proceedings of the 11th ACM SIGPLAN International Symposium on Haskell (Haskell 2018)*. Association for Computing Machinery, New York, NY, USA, 179–185. <https://doi.org/10.1145/3242744.3242760>

- Ranjit Jhala. 2020. Liquid Haskell is a GHC Plugin. <https://ucsd-progsys.github.io/liquidhaskell-blog/2020/08/20/lh-as-a-ghc-plugin.lhs/>.
- K. Rustan M. Leino. 2010. Dafny: An Automatic Program Verifier for Functional Correctness. In *Proceedings of the 16th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR'10)*. Springer-Verlag, Berlin, Heidelberg, 348–370.
- Yiyun Liu, James Parker, Patrick Redmond, Lindsey Kuper, Michael Hicks, and Niki Vazou. 2020. Verifying Replicated Data Types with Typeclass Refinements in Liquid Haskell. *Proc. ACM Program. Lang.* 4, OOPSLA, Article 216 (Nov. 2020), 30 pages. <https://doi.org/10.1145/3428284>
- Ulf Norell. 2008. Dependently Typed Programming in Agda. In *Proceedings of the 6th International Conference on Advanced Functional Programming (AFP'08)*. Springer-Verlag, Berlin, Heidelberg, 230–266.
- James Parker, Niki Vazou, and Michael Hicks. 2019. LWeb: Information Flow Security for Multi-Tier Web Applications. *Proc. ACM Program. Lang.* 3, POPL, Article 75 (Jan. 2019), 30 pages. <https://doi.org/10.1145/3290388>
- Talia Ringer, Karl Palmskog, Ilya Sergey, Milos Gligoric, and Zachary Tatlock. 2019. QED at Large: A Survey of Engineering of Formally Verified Software. *Foundations and Trends® in Programming Languages* 5, 2-3 (2019), 102–281. <https://doi.org/10.1561/25000000045>
- Patrick M. Rondon, Ming Kawaguci, and Ranjit Jhala. 2008. Liquid Types. In *Proceedings of the 29th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '08)*. Association for Computing Machinery, New York, NY, USA, 159–169. <https://doi.org/10.1145/1375581.1375602>
- J. Rushby, S. Owre, and N. Shankar. 1998. Subtypes for specifications: predicate subtyping in PVS. *IEEE Transactions on Software Engineering* 24, 9 (1998), 709–720. <https://doi.org/10.1109/32.713327>
- Marc Shapiro, Nuno Preguiça, Carlos Baquero, and Marek Zawirski. 2011. Conflict-Free Replicated Data Types. In *Stabilization, Safety, and Security of Distributed Systems*, Xavier Défago, Franck Petit, and Vincent Villain (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 386–400.
- The Agda Team 2021. Holes and case splitting. <https://agda.readthedocs.io/en/v2.6.2/getting-started/a-taste-of-agda.html#holes-and-case-splitting>.
- Niki Vazou. 2016. *Liquid Haskell: Haskell as a theorem prover*. Ph.D. Dissertation.
- Niki Vazou, Joachim Breitner, Rose Kunkel, David Van Horn, and Graham Hutton. 2018. Theorem Proving for All: Equational Reasoning in Liquid Haskell (Functional Pearl). In *Proceedings of the 11th ACM SIGPLAN International Symposium on Haskell (Haskell 2018)*. Association for Computing Machinery, New York, NY, USA, 132–144. <https://doi.org/10.1145/3242744.3242756>
- Niki Vazou, Leonidas Lampropoulos, and Jeff Polakow. 2017a. A Tale of Two Provers: Verifying Monoidal String Matching in Liquid Haskell and Coq. In *Proceedings of the 10th ACM SIGPLAN International Symposium on Haskell (Haskell 2017)*. Association for Computing Machinery, New York, NY, USA, 63–74. <https://doi.org/10.1145/3122955.3122963>
- Niki Vazou, Eric L. Seidel, Ranjit Jhala, Dimitrios Vytiniotis, and Simon Peyton-Jones. 2014. Refinement Types for Haskell. In *Proceedings of the 19th ACM SIGPLAN International Conference on Functional Programming (ICFP '14)*. Association for Computing Machinery, New York, NY, USA, 269–282. <https://doi.org/10.1145/2628136.2628161>
- Niki Vazou, Anish Tondwalkar, Vikraman Choudhury, Ryan G. Scott, Ryan R. Newton, Philip Wadler, and Ranjit Jhala. 2017b. Refinement Reflection: Complete Verification with SMT. *Proc. ACM Program. Lang.* 2, POPL, Article 53 (Dec. 2017), 31 pages. <https://doi.org/10.1145/3158141>
- Makarius Wenzel, Lawrence C. Paulson, and Tobias Nipkow. 2008. The Isabelle Framework. In *Theorem Proving in Higher Order Logics*, Otmane Ait Mohamed, César Muñoz, and Sofiène Tahar (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 33–38.
- Hongwei Xi and Frank Pfenning. 1998. Eliminating Array Bound Checking through Dependent Types. In *Proceedings of the ACM SIGPLAN 1998 Conference on Programming Language Design and Implementation (PLDI '98)*. Association for Computing Machinery, New York, NY, USA, 249–257. <https://doi.org/10.1145/277650.277732>