

Lower Your Guards

A Compositional Pattern-Match Coverage Checker

SEBASTIAN GRAF, Karlsruhe Institute of Technology, Germany

SIMON PEYTON JONES, Microsoft Research, UK

RYAN G. SCOTT, Indiana University, USA

A compiler should warn if a function defined by pattern matching does not cover its inputs—that is, if there are missing or redundant patterns. Generating such warnings accurately is difficult for modern languages due to the myriad of language features that interact with pattern matching. This is especially true in Haskell, a language with a complicated pattern language that is made even more complex by extensions offered by the Glasgow Haskell Compiler (GHC). Although GHC has spent a significant amount of effort towards improving its pattern-match coverage warnings, there are still several cases where it reports inaccurate warnings.

We introduce a coverage checking algorithm called Lower Your Guards, which boils down the complexities of pattern matching into *guard trees*. While the source language may have many exotic forms of patterns, guard trees only have three different constructs, which vastly simplifies the coverage checking process. Our algorithm is modular, allowing for new forms of source-language patterns to be handled with little changes to the overall structure of the algorithm. We have implemented the algorithm in GHC and demonstrate places where it performs better than GHC’s current coverage checker, both in accuracy and performance.

CCS Concepts: • **Software and its engineering** → **Compilers**; *Software maintenance tools*; Procedures, functions and subroutines; *Constraints*; *Functional languages*; *Multiparadigm languages*.

Additional Key Words and Phrases: Haskell, pattern matching, guards, strictness

ACM Reference Format:

Sebastian Graf, Simon Peyton Jones, and Ryan G. Scott. 2020. Lower Your Guards: A Compositional Pattern-Match Coverage Checker. *Proc. ACM Program. Lang.* 4, ICFP, Article 107 (August 2020), 39 pages. <https://doi.org/10.1145/3408989>

1 INTRODUCTION

Pattern matching is a tremendously useful feature in Haskell and many other programming languages, but it must be used with care. Consider this example of pattern matching gone wrong:

```
f :: Int → Bool
f 0 = True
f 0 = False
```

The function f has two serious flaws. One obvious problem is that there are two clauses that match on 0, and due to the top-to-bottom semantics of pattern matching, this makes the $f\ 0 = \text{False}$ clause completely unreachable. Even worse is that f never matches on any patterns besides 0, making it not fully defined. Attempting to invoke $f\ 1$, for instance, will fail.

Authors’ addresses: Sebastian Graf, Karlsruhe Institute of Technology, Karlsruhe, Germany, sebastian.graf@kit.edu; Simon Peyton Jones, Microsoft Research, Cambridge, UK, simonpj@microsoft.com; Ryan G. Scott, Indiana University, Bloomington, Indiana, USA, rgscott@indiana.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2020 Copyright held by the owner/author(s).

2475-1421/2020/8-ART107

<https://doi.org/10.1145/3408989>

To avoid these mishaps, compilers for languages with pattern matching often emit warnings (or errors) if a function is missing clauses (i.e., if it is *non-exhaustive*), if one of its right-hand sides will never be entered (i.e., if it is *inaccessible*), or if one of its equations can be deleted altogether (i.e., if it is *redundant*). We refer to the combination of checking for exhaustivity, redundancy, and accessibility as *pattern-match coverage checking*. Coverage checking is the first line of defence in catching programmer mistakes when defining code that uses pattern matching.

Coverage checking for a set of equations matching on algebraic data types is a well studied (although still surprisingly tricky) problem—see [Section 8](#) for this related work. But the coverage-checking problem becomes *much* harder when one includes the raft of innovations that have become part of a modern programming language like Haskell, including: view patterns, pattern guards, pattern synonyms, overloaded literals, bang patterns, lazy patterns, as-patterns, strict data constructors, empty case expressions, and long-distance effects ([Section 4](#)). Particularly tricky are: *GADTs* where the *type* of a match can determine what *values* can possibly appear [[Xi et al. 2003](#)]; and *local type-equality constraints* brought into scope by pattern matching [[Vytiniotis et al. 2011](#)].

The current state of the art for coverage checking in a richer language of this sort is *GADTs Meet Their Match* [[Karachalias et al. 2015](#)], or GMTM for short. It presents an algorithm that handles the intricacies of checking *GADTs*, lazy patterns, and pattern guards. However GMTM is monolithic and does not account for a number of important language features; it gives incorrect results in certain cases; its formulation in terms of structural pattern matching makes it hard to avoid some serious performance problems; and its implementation in GHC, while a big step forward over its predecessors, has proved complex and hard to maintain.

In this paper we propose a new, compositional coverage-checking algorithm, called Lower Your Guards (LYG), that is simpler, more modular, *and* more powerful than GMTM (see [Section 8.1](#)). Moreover, it avoids GMTM’s performance pitfalls. We make the following contributions:

- We characterise some nuances of coverage checking that not even GMTM handles ([Section 2](#)). We also identify issues in GHC’s implementation of GMTM.
- We describe a new, compositional coverage checking algorithm, LYG, in [Section 3](#). The key insight is to abandon the notion of structural pattern matching altogether, and instead desugar all the complexities of pattern matching into a very simple language of *guard trees*, with just three constructs ([Section 3.1](#)). Coverage checking on these guard trees becomes remarkably simple, returning an *annotated tree* ([Section 3.2](#)) decorated with *refinement types*. Finally, provided we have access to a suitable way to find inhabitants of a refinement type, we can report accurate coverage errors ([Section 3.3](#)).
- We demonstrate the compositionality of LYG by augmenting it with several language extensions ([Section 4](#)). Although these extensions can change the source language in significant ways, the effort needed to incorporate them into the algorithm is comparatively small.
- We discuss how to optimize the performance of LYG ([Section 5](#)) and implement a proof of concept in GHC ([Section 6](#)).
- The evaluation against a large number of Haskell packages ([Section 6](#)) provides evidence that LYG is sound. In order to discuss soundness formally in [Section 7](#), we turn the informal semantics of guard trees and refinement types in [Section 3](#) into formal semantics. We also list mechanisms that render LYG incomplete in order to guarantee good performance.
- We discuss the wealth of related work in [Section 8](#).

2 THE PROBLEM WE WANT TO SOLVE

What makes coverage checking so difficult in a language like Haskell? At first glance, implementing a coverage checking algorithm might appear simple: just check that every function matches on

every possible combination of data constructors exactly once. A function must match on every possible combination of constructors in order to be exhaustive, and it must match on them exactly once to avoid redundant matches.

This algorithm, while concise, leaves out many nuances. What constitutes a “match”? Haskell has multiple matching constructs, including function definitions, `case` expressions, and guards. How does one count the number of possible combinations of data constructors? This is not a simple exercise since term and type constraints can make some combinations of constructors unreachable if matched on, and some combinations of data constructors can overlap others. Moreover, what constitutes a “data constructor”? In addition to traditional data constructors, GHC features *pattern synonyms* [Pickering et al. 2016], which provide an abstract way to embed arbitrary computation into patterns. Matching on a pattern synonym is syntactically identical to matching on a data constructor, which makes coverage checking in the presence of pattern synonyms challenging.

Prior work on coverage checking (discussed in Section 8) accounts for some of these nuances, but not all of them. In this section we identify some key language features that make coverage checking difficult. While these features may seem disparate at first, we will later show in Section 3 that these ideas can all fit into a unified framework.

2.1 Guards

Guards are a flexible form of control flow in Haskell. Here is a function that demonstrates various capabilities of guards:

```
guardDemo :: Char → Char → Int
guardDemo c1 c2 | c1 == 'a'           = 0
                 | 'b' ← c1           = 1
                 | let c1' = c1, 'c' ← c1', c2 == 'd' = 2
                 | otherwise          = 3
```

This function has four *guarded right-hand sides* or GRHSs for short. The first GRHS has a *boolean guard*, (`c1 == 'a'`), that succeeds if the expression in the guard returns `True`. The second GRHS has a *pattern guard*, (`'b' ← c1`), that succeeds if the pattern in the guard successfully matches. The next line illustrates that each GRHS may have multiple guards, and that guards include `let` bindings, such as `let c1' = c2`. The fourth GRHS uses *otherwise*, which is simply defined as `True`.

Guards can be thought of as a generalization of patterns, and we would like to include them as part of coverage checking. Checking guards is significantly more complicated than checking ordinary structural pattern matches, however, since guards can contain arbitrary expressions. Consider this implementation of the *signum* function:

```
signum :: Int → Int
signum x | x > 0 = 1
         | x == 0 = 0
         | x < 0 = -1
```

Intuitively, *signum* is exhaustive since the combination of (`>`), (`==`), and (`<`) covers all possible `Int`s. This is hard for a machine to check, because doing so requires knowledge about the properties of `Int` inequalities. Clearly, coverage checking for guards is undecidable in general. However, while we cannot accurately check *all* uses of guards, we can at least give decent warnings for some common cases. For instance, take the following functions:

$not :: Bool \rightarrow Bool$	$not2 :: Bool \rightarrow Bool$	$not3 :: Bool \rightarrow Bool$
$not\ b \mid False \leftarrow b = True$	$not2\ False = True$	$not3\ x \mid False \leftarrow x = True$
$\mid True \leftarrow b = False$	$not2\ True = False$	$not3\ True = False$

Clearly all are equivalent. Our coverage checking algorithm should find that all three are exhaustive, and indeed, LYG does so.

2.2 Programmable Patterns

Expressions in guards are not the only source of undecidability that the coverage checker must cope with. GHC extends the pattern language in other ways that are also impossible to check in the general case. We consider two such extensions here: view patterns and pattern synonyms.

2.2.1 View Patterns. View patterns allow arbitrary computation to be performed while pattern matching. When a value v is matched against a view pattern ($f \rightarrow p$), the match is successful when $f\ v$ successfully matches against the pattern p . For example, one can use view patterns to succinctly define a function that computes the length of Haskell’s opaque Text data type:

```
Text.null :: Text → Bool    -- Checks if a Text is empty
Text.uncons :: Text → Maybe (Char, Text)  -- If a Text is non-empty, return Just (x, xs),
                                         -- where x is the first character and xs is the rest
```

```
length :: Text → Int
length (Text.null → True)      = 0
length (Text.uncons → Just (_, xs)) = 1 + length xs
```

Again, it would be unreasonable to expect a coverage checking algorithm to prove that *length* is exhaustive, but one might hope for a coverage checking algorithm that handles some common usage patterns. For example, LYG indeed *is* able to prove that the *safeLast* function is exhaustive:

```
safeLast :: [a] → Maybe a
safeLast (reverse → [ ])      = Nothing
safeLast (reverse → (x : _)) = Just x
```

2.2.2 Pattern Synonyms. Pattern synonyms [Pickering et al. 2016] allow abstraction over patterns themselves. Pattern synonyms and view patterns can be useful in tandem, as the pattern synonym can present an abstract interface to a view pattern that does complicated things under the hood. For example, one can define *length* with pattern synonyms like so:

pattern Nil :: Text	$length :: Text \rightarrow Int$
pattern Nil \leftarrow (Text.null \rightarrow True)	$length\ Nil = 0$
pattern Cons :: Char \rightarrow Text \rightarrow Text	$length\ (Cons\ _ xs) = 1 + length\ xs$
pattern Cons $x\ xs \leftarrow$ (Text.uncons \rightarrow Just (x, xs))	

The pattern synonym Nil matches precisely when the view pattern Text.null \rightarrow True would match, and similarly for Cons.

How should a coverage checker handle pattern synonyms? One idea is to simply “look through” the definitions of each pattern synonym and verify whether the underlying patterns are exhaustive. This would be undesirable, however, because (1) we would like to avoid leaking the implementation details of abstract pattern synonyms, and (2) even if we *did* look at the underlying implementation, it would be challenging to automatically check that the combination of Text.null and Text.uncons is exhaustive.

Nevertheless, `Text.null` and `Text.uncons` together are in fact exhaustive, and GHC allows programmers to communicate this fact to the coverage checker using a `COMPLETE` pragma [GHC team 2020]. A `COMPLETE` set is a combination of data constructors and pattern synonyms that should be regarded as exhaustive when a function matches on all of them. For example, declaring `{-# COMPLETE Nil, Cons #-}` is sufficient to make the definition of `length` above compile without any exhaustivity warnings. Since GHC does not (and cannot, in general) check that all of the members of a `COMPLETE` set actually comprise a complete set of patterns, the burden is on the programmer to ensure that this invariant is upheld.

2.3 Strictness

The evaluation order of pattern matching can impact whether a pattern is reachable or not. Consider:

```
f :: Bool → Bool → Int
f _   False = 1
f True False = 2
f _   _     = 3
```

Is the second clause redundant? In a strict language such as OCaml or Lean the answer is “Yes”, but in lazy Haskell this the answer is “No”. To see that, consider the call `f (error "boom") True`, an expression that in a strict language would immediately evaluate the error in the argument (`error "boom"`) by-value before making the call.

In Haskell, after falling through the first clause that does not match in the second argument, the second clause will evaluate the first argument in order to match against `True`. Doing so forces the error, to much the same effect as in a strict language, and we see that `f (error "boom") True` evaluates to `error "boom"`. However, if we *remove* the second clause, `f (error "boom") True` would evaluate to 3, because the first argument is never needed during pattern-matching. Since removing the clause changes the semantics of the function, it cannot be redundant, but its right-hand side is *inaccessible* still (Section 2.3.1).

We could have used a different error such as `undefined` or a non-terminating expression such as `loop = loop` in the example above. The Haskell Language Report [Marlow et al. 2010] does not distinguish between these different kinds of divergence, referring to them as \perp , and we do the same here. There are a number of language features which interact with \perp in the context of pattern-matching.

2.3.1 Redundancy versus Inaccessibility. The example function `f` above demonstrates that when reporting unreachable equations, we must distinguish between *redundant* and *inaccessible* cases. A redundant equation can be removed from a function without changing its semantics, whereas an inaccessible equation cannot, even though its right-hand side is unreachable. The examples below illustrate the challenges for LYG in more detail:

<pre>u :: () → Int u () False = 1 True = 2 u _ = 3</pre>	<pre>u' :: () → Int u' () False = 1 False = 2 u' _ = 3</pre>
---	--

Within `u`, the equations that return 1 and 3 could be deleted without changing the semantics of `u`, so they are classified as *redundant*. Within `u'`, the right-hand sides of the equations that return 1 and 2 are inaccessible, but they cannot both be redundant because their clause evaluates the parameter; `u' ⊥ = ⊥`. As a result, LYG picks the first equation in `u'` as inaccessible to keep alive the pattern-match on the parameter, and the second equation as redundant. Inaccessibility suggests

to the programmer that u' might benefit from a refactor to render the first equation redundant as well (e.g., $u' () = 3$).

Observe that u and u' have completely different warnings, but the only difference between the two functions is whether the second equation uses `True` or `False` in its guard. Moreover, this second equation affects the warnings for *other* equations. This demonstrates that determining whether code is redundant or inaccessible is a non-local problem. Inaccessibility may seem like a corner case, but GHC's users have reported many bugs of this sort (Section 6.2).

2.3.2 Strict Fields. Just like Haskell function parameters, fields of data constructors may hide arbitrary computations as well. However, Haskell programmers can opt into extra strict evaluation by giving a data type strict fields, such as in this example:

```
data Void -- No data constructors; only inhabitant is bottom
data SMaybe a = SJust !a | SNothing
v :: SMaybe Void → Int
v SNothing = 0
v (SJust _) = 1 -- Redundant!
```

The “!” in the definition of `SJust` makes the constructor strict, so $(SJust \perp) = \perp$ semantically, in contrast to the regular lazy `Just` constructor.

Curiously, the strict field semantics of `SJust` makes the second equation of v redundant! Since \perp is the only inhabitant of type `Void`, the only inhabitants of `SMaybe Void` are `SNothing` and \perp . The former will match on the first equation; the latter will make the first equation diverge. In neither case will execution flow to the second equation, so it is redundant and can be deleted.

2.3.3 Bang Patterns. Strict data-constructor fields are one mechanism for adding extra strictness in ordinary Haskell, but GHC adds another in the form of *bang patterns*. When a value v is matched against a bang pattern `!pat`, first v is evaluated to weak-head normal form (WHNF), a step that might diverge, and then v is matched against `pat`. Here is a variant of v , this time using the standard, lazy `Maybe` data type:

```
v' :: Maybe Void → Int
v' Nothing = 0
v' (Just !_) = 1 -- Not redundant, but GRHS is inaccessible
```

The inhabitants of the type `Maybe Void` are \perp , `Nothing`, and $(Just \perp)$. The input \perp makes the first equation diverge; `Nothing` matches on the first equation; and $(Just \perp)$ makes the second equation diverge because of the bang pattern. Therefore, none of the three inhabitants will result in the right-hand side of the second equation being reached. Note that the second equation is inaccessible, but not redundant (Section 2.3.1).

2.4 Type-Equality Constraints

Besides strictness, another way for pattern matches to be rendered unreachable is by way of *equality constraints*. A popular method for introducing equalities between types is matching on GADTs [Xi et al. 2003]. The following examples demonstrate the interaction between GADTs and coverage checking:

<code>data T a b where</code>	<code>g1 :: T Int b → b → Int</code>	<code>g2 :: T a b → T a b → Int</code>
<code>T1 :: T Int Bool</code>	<code>g1 T1 False = 0</code>	<code>g2 T1 T1 = 0</code>
<code>T2 :: T Char Bool</code>	<code>g1 T1 True = 1</code>	<code>g2 T2 T2 = 1</code>

When $g1$ matches against `T1`, the b in the type `T Int b` is known to be a `Bool`, which is why matching




Fig. 1. Bird's eye view of pattern match checking

Meta variables	Pattern syntax
x, y, z, f, g, h Term variables	$defn ::= \overline{clause}$
a, b, c Type variables	$clause ::= f \overline{pat} \overline{match}$
K Data constructors	$pat ::= x \mid _ \mid K \overline{pat} \mid x@pat \mid !pat \mid expr \rightarrow pat$
P Pattern synonyms	$match ::= = \overline{expr} \mid \overline{grhs}$
T Type constructors	$grhs ::= \mid \overline{guard} = expr$
l Literal	$guard ::= pat \leftarrow expr \mid expr \mid let x = expr$
$expr$ Expressions	

Fig. 2. Source syntax: A desugared Haskell

the second argument against False or True will typecheck. Phrased differently, matching against T1 brings into scope an *equality constraint* between the types b and Bool. GHC has a powerful type inference engine that is equipped to reason about type equalities of this sort [Vytiñotis et al. 2011].

Just as important as the code used in the $g1$ function is the code that is *not* used in $g1$. One might wonder if $g1$ not matching its first argument against T2 is an oversight. In fact, the exact opposite is true: matching on T2 would be rejected by the typechecker. This is because T2 is of type T Char Bool, but the first argument to $g1$ must be of type T Int b . Matching against T2 would be tantamount to saying that Int and Char are the same type, which is not the case. As a result, $g1$ is exhaustive even though it does not match on all of T's data constructors.

The presence of type equalities is not always as clear-cut as it is in $g1$. Consider the more complex $g2$ function, which matches on two arguments of the type T $a b$. While matching the arguments against T1 T1 or T2 T2 is possible, it is not possible to match against T1 T2 or T2 T1. To see why, suppose the first argument is matched against T1, giving rise to an equality between a and Int. If the second argument were then matched against T2, we would have that a equals Char. By the transitivity of type equality, we would have that Int equals Char. This cannot be true, so matching against T1 T2 is impossible (and similarly for T2 T1).

Concluding that $g2$ is exhaustive requires some non-trivial reasoning about equality constraints. In GHC, the same engine that typechecks GADT pattern matches is also used to rule out cases made unreachable by type equalities, and LYG adopts a similar approach. Besides GHC's current coverage checker [Karachalias et al. 2015], there are a variety of other coverage checking algorithms that account for GADTs, including those for OCaml [Garrigue and Normand 2011], Dependent ML [Xi 1998a,b, 2003], and Stardust [Dunfield 2007].

3 LOWER YOUR GUARDS: A NEW COVERAGE CHECKER

In this section, we describe our new coverage checking algorithm, LYG. Figure 1 depicts a high-level overview, which divides into three steps:

Guard syntax

$$\begin{array}{ll}
k, n, m \in \mathbb{N} & \gamma \in \text{TyCt} ::= \tau_1 \sim \tau_2 \mid \dots \\
K \in \text{Con} & p \in \text{Pat} ::= _ \mid K \bar{p} \mid \dots \\
x, y, a, b \in \text{Var} & g \in \text{Grd} ::= \text{let } x : \tau = e \\
\tau, \sigma \in \text{Type} ::= a \mid \dots & \mid K \bar{a} \bar{y} \bar{y} : \bar{\tau} \leftarrow x \\
e \in \text{Expr} ::= x \mid K \bar{\tau} \bar{y} \bar{e} \mid \dots & \mid !x
\end{array}$$

Clause tree syntax

$$\begin{array}{l}
t \in \text{Gdt} ::= \rightarrow k \mid \bigcup_{t_2}^{t_1} \mid \dashv g \dashv t \\
u \in \text{Ant} ::= \rightarrow \Theta k \mid \bigcup_{u_2}^{u_1} \mid \dashv \Theta \not\vdash \dashv u
\end{array}$$

Refinement type syntax

$$\begin{array}{ll}
\Gamma ::= \emptyset \mid \Gamma, x : \tau \mid \Gamma, a & \text{Context} \\
\varphi ::= \checkmark \mid \times \mid K \bar{a} \bar{y} \bar{y} : \bar{\tau} \leftarrow x \mid x \not\approx K \mid x \approx \perp \mid x \not\approx \perp \mid \text{let } x = e & \text{Literals} \\
\Phi ::= \varphi \mid \Phi \wedge \Phi \mid \Phi \vee \Phi & \text{Formula} \\
\Theta ::= \langle \Gamma \mid \Phi \rangle & \text{Refinement type}
\end{array}$$

Fig. 3. IR syntax

- First, we desugar the complex source Haskell syntax (cf. [Figure 2](#)) into a **guard tree** $t \in \text{Gdt}$ ([Section 3.1](#)). The language of guard trees is tiny but expressive, and allows the subsequent passes to be entirely independent of the source syntax. LYG can readily be adapted to other languages simply by changing the desugaring algorithm.
- Next, the resulting guard tree is then processed by two different functions ([Section 3.2](#)). The function $\mathcal{A}(t)$ produces an **annotated tree** $u \in \text{Ant}$, which has the same general branching structure as t but describes which clauses are accessible, inaccessible, or redundant. The function $\mathcal{U}(t)$, on the other hand, returns a *refinement type* Θ [[Rushby et al. 1998](#); [Xi and Pfenning 1998](#)] that describes the set of *uncovered values*, which are not matched by any of the clauses.
- Finally, an error-reporting pass generates comprehensible error messages ([Section 3.3](#)). Again there are two things to do. The function \mathcal{R} processes the annotated tree produced by \mathcal{A} to explicitly identify the accessible, inaccessible, or redundant clauses. The function $\mathcal{G}(\Theta)$ produces representative *inhabitants* of the refinement type Θ (produced by \mathcal{U}) that describes the uncovered values.

LYG's main contribution when compared to other coverage checkers, such as GHC's implementation of GMTM, is its incorporation of many small improvements and insights, rather than a single defining breakthrough. In particular, LYG's advantages are:

- Achieving modularity by clearly separating the source syntax ([Figure 2](#)) from the intermediate language ([Figure 3](#)).
- Correctly accounting for strictness in identifying redundant and inaccessible code ([Section 8.5](#)).
- Using detailed term-level reasoning ([Figures 6 to 8](#)), which GMTM does not.

$$\mathcal{D}(\text{defn}) = \text{Gdt}, \quad \mathcal{D}(\text{clause}) = \text{Gdt}, \quad \mathcal{D}(\text{grhs}) = \text{Gdt}$$

k_{rhs} is the index of the right hand side rhs

$$\begin{aligned} \mathcal{D}(\text{clause}_1 \dots \text{clause}_n) &= \begin{array}{l} \top \\ \vdots \\ \mathcal{D}(\text{clause}_n) \end{array} \\ \mathcal{D}(f \text{ pat}_1 \dots \text{pat}_n = \text{rhs}) &= \neg \top \mathcal{D}(x_1, \text{pat}_1) \dots \mathcal{D}(x_n, \text{pat}_n) \rightarrow k_{rhs} \\ \mathcal{D}(f \text{ pat}_1 \dots \text{pat}_n \text{ grhs}_1 \dots \text{grhs}_m) &= \neg \top \mathcal{D}(x_1, \text{pat}_1) \dots \mathcal{D}(x_n, \text{pat}_n) \begin{array}{l} \top \\ \vdots \\ \mathcal{D}(\text{grhs}_m) \end{array} \\ \mathcal{D}(| \text{guard}_1 \dots \text{guard}_n = \text{rhs}) &= \neg \top \mathcal{D}(\text{guard}_1) \dots \mathcal{D}(\text{guard}_n) \rightarrow k_{rhs} \end{aligned}$$

$$\mathcal{D}(\text{guard}) = \overline{\text{Grd}}, \quad \mathcal{D}(x, \text{pat}) = \overline{\text{Grd}}$$

$$\begin{aligned} \mathcal{D}(\text{pat} \leftarrow \text{expr}) &= \text{let } x = \text{expr}, \mathcal{D}(x, \text{pat}) \quad x \text{ fresh} \\ \mathcal{D}(\text{expr}) &= \text{let } y = \text{expr}, \mathcal{D}(y, \text{True}) \quad y \text{ fresh} \\ \mathcal{D}(\text{let } x = \text{expr}) &= \text{let } x = \text{expr} \\ \mathcal{D}(x, y) &= \text{let } y = x \\ \mathcal{D}(x, _) &= \epsilon \\ \mathcal{D}(x, K \text{ pat}_1 \dots \text{pat}_n) &= !x, K \ y_1 \dots y_n \leftarrow x, \mathcal{D}(y_1, \text{pat}_1), \dots, \mathcal{D}(y_n, \text{pat}_n) \quad y_i \text{ fresh } (\dagger) \\ \mathcal{D}(x, y@pat) &= \text{let } y = x, \mathcal{D}(y, \text{pat}) \\ \mathcal{D}(x, !pat) &= !x, \mathcal{D}(x, \text{pat}) \\ \mathcal{D}(x, \text{expr} \rightarrow \text{pat}) &= \text{let } y = \text{expr } x, \mathcal{D}(y, \text{pat}) \quad y \text{ fresh} \end{aligned}$$

Fig. 4. Desugaring from source language to Gdt

- Using *negative information* to sidestep serious performance issues in GMTM without changing the worst-case complexity (Section 8.4). This also enables graceful degradation (Section 5.2) and the ability to handle COMPLETE sets properly (Section 5.3).
- Fixing various bugs present in GMTM, both in the paper [Karachalias et al. 2015] and in GHC's implementation thereof (Section 6.2).

3.1 Desugaring to Guard Trees

The first step is to desugar the source language into the language of guard trees. The syntax of the source language is given in Figure 2. Definitions *defn* consist of a list of *clauses*, each of which has a list of *patterns*, and a list of *guarded right-hand sides* (GRHSs). Patterns include variables and constructor patterns, of course, but also a representative selection of extensions: wildcards, as-patterns, bang patterns, and view patterns. We explore several other extensions in Section 4.

The language of guard trees Gdt is much smaller; its graphical syntax is given in Figure 3. All of the syntactic redundancy of the source language is translated into a minimal form very similar to pattern guards. We start with an example:

$$\begin{aligned} f \text{ (Just (!xs, _)) } ys@Nothing &= \text{True} \\ f \text{ Nothing } (g \rightarrow \text{True}) &= \text{False} \end{aligned}$$

This desugars to the following guard tree (where the x_i represent f 's arguments):

$$\begin{array}{l} \top \!|x_1, \text{Just } t_1 \leftarrow x_1, \!|t_1, (t_2, t_3) \leftarrow t_1, \!|t_2, \text{let } xs = t_2, \text{let } ys = x_2, \!|ys, \text{Nothing} \leftarrow ys \longrightarrow 1 \\ \sqcup \!|x_1, \text{Nothing} \leftarrow x_1, \text{let } t_4 = g \ x_2, \!|t_4, \text{True} \leftarrow t_4 \longrightarrow 2 \end{array}$$

The first line says “evaluate x_1 ; then match x_1 against `Just` t_1 ; then evaluate t_1 ; then match t_1 against (t_2, t_3) ” and so on. If any of those matches fail, we fall through into the second line. Note that we write $\neg |g_1, \dots, g_n \neg t$ instead of $\neg |g_1 \neg \dots \neg |g_n \neg t$ for notational convenience.

Informally, matching a guard tree may *succeed* (binding the variables bound in the tree), *fail*, or *diverge*. Referring to the syntax of guard trees in [Figure 3](#), matching is defined as follows:

- Matching a guard tree $\longrightarrow k$ succeeds, and selects the k 'th right hand side of the pattern match group.
- Matching a guard tree $\begin{array}{l} \top t_1 \\ \sqcup t_2 \end{array}$ means matching against t_1 ; if that succeeds, the overall match succeeds; if not, match against t_2 .
- Matching a guard tree $\neg |x \neg t$ evaluates x ; if that diverges the match diverges; if not match t .
- Matching a guard tree $\neg |K \bar{a} \bar{y} \bar{y} \leftarrow x \neg t$ matches x against constructor K . If the match succeeds, bind \bar{a} to the type components, \bar{y} to the constraint components and \bar{y} to the term components, then match t . If the constructor match fails, then the entire match fails.
- Matching a guard tree $\neg |\text{let } x = e \neg t$ binds x (lazily) to e , and matches t .

See ?? for a formal account of this semantics. The desugaring algorithm, \mathcal{D} , is given in [Figure 4](#). It is a straightforward recursive descent over the source syntax, with a little bit of administrative bureaucracy to account for renaming. It also generates an abundance of fresh temporary variables; in practice, the implementation of \mathcal{D} can be smarter than this by looking at the pattern (which might be a variable match or as-pattern) when choosing a name for a temporary variable. In that case, it is important that every binder in the source language has a unique name.

Notice that both “structural” pattern-matching in the source language (e.g. the match on `Nothing` in the second equation), and view patterns (e.g. $g \rightarrow \text{True}$) can readily be compiled to a single form of matching in guard trees. The same holds for pattern guards. For example, consider this (stylistically contrived) definition of *liftEq*, which is inexhaustive:

$$\begin{array}{l} \text{liftEq Nothing Nothing} = \text{True} \\ \text{liftEq } mx \quad (\text{Just } y) \mid \text{Just } x \leftarrow mx, x == y = \text{True} \\ \quad \quad \quad \mid \text{otherwise} \quad \quad \quad = \text{False} \end{array}$$

It desugars thus:

$$\begin{array}{l} \top \!|mx, \text{Nothing} \leftarrow mx, \!|my, \text{Nothing} \leftarrow my \longrightarrow 1 \\ \sqcup \!|my, \text{Just } y \leftarrow my \begin{array}{l} \top \!|mx, \text{Just } x \leftarrow mx, \text{let } t = x == y, \!|t, \text{True} \leftarrow t \longrightarrow 2 \\ \sqcup \!|otherwise, \text{True} \leftarrow otherwise \longrightarrow 3 \end{array} \end{array}$$

Notice that the pattern guard `(Just $x \leftarrow mx$)` and the boolean guard `($x == y$)` have both turned into the same constructor-matching construct in the guard tree.

In equation (†) of [Figure 4](#) we generate an explicit bang guard $\!|x$ to reflect the fact that pattern matching against a data constructor requires evaluation. However, Haskell’s `newtype` declarations introduce data constructors that are *not* strict, so their desugaring is just like (†) but with no $\!|x$ ([Section 4.7](#)). From this point onwards, then, strictness is expressed *only* through bang guards $\!|x$, while constructor guards $K \ a \ b \leftarrow y$ are not considered strict.

In a way there is nothing very deep here, but it took us a surprisingly long time to come up with the language of guard trees. We recommend it!

Operations on Θ

$$\begin{aligned} \langle \Gamma \mid \Phi \rangle \wedge \varphi &= \langle \Gamma \mid \Phi \wedge \varphi \rangle \\ \langle \Gamma \mid \Phi_1 \rangle \cup \langle \Gamma \mid \Phi_2 \rangle &= \langle \Gamma \mid \Phi_1 \vee \Phi_2 \rangle \end{aligned}$$

Checking guard trees

$$\begin{aligned} &\boxed{\mathcal{U}(\Theta, t) = \Theta} \\ \mathcal{U}(\langle \Gamma \mid \Phi \rangle, \rightarrow n) &= \langle \Gamma \mid \times \rangle \\ \mathcal{U}(\Theta, \bigsqcup_{t_1, t_2}^{t_1, t_2}) &= \mathcal{U}(\mathcal{U}(\Theta, t_1), t_2) \\ \mathcal{U}(\Theta, \neg !x \text{---} t) &= \mathcal{U}(\Theta \hat{\wedge} (x \neq \perp), t) \\ \mathcal{U}(\Theta, \neg \text{let } x = e \text{---} t) &= \mathcal{U}(\Theta \hat{\wedge} (\text{let } x = e), t) \\ \mathcal{U}(\Theta, \neg K \bar{a} \bar{y} \bar{y} : \bar{\tau} \leftarrow x \text{---} t) &= (\Theta \hat{\wedge} (x \neq K)) \cup \mathcal{U}(\Theta \hat{\wedge} (K \bar{a} \bar{y} \bar{y} : \bar{\tau} \leftarrow x), t) \\ &\boxed{\mathcal{A}(\Theta, t) = u} \\ \mathcal{A}(\Theta, \rightarrow n) &= \rightarrow \Theta n \\ \mathcal{A}(\Theta, \bigsqcup_{t_1, t_2}^{t_1, t_2}) &= \bigsqcup_{\mathcal{A}(\Theta, t_1), \mathcal{A}(\mathcal{U}(\Theta, t_1), t_2)} \\ \mathcal{A}(\Theta, \neg !x \text{---} t) &= \text{---} \Theta \hat{\wedge} (x \approx \perp) \not\leftarrow \mathcal{A}(\Theta \hat{\wedge} (x \neq \perp), t) \\ \mathcal{A}(\Theta, \neg \text{let } x = e \text{---} t) &= \mathcal{A}(\Theta \hat{\wedge} (\text{let } x = e), t) \\ \mathcal{A}(\Theta, \neg K \bar{a} \bar{y} \bar{y} : \bar{\tau} \leftarrow x \text{---} t) &= \mathcal{A}(\Theta \hat{\wedge} (K \bar{a} \bar{y} \bar{y} : \bar{\tau} \leftarrow x), t) \end{aligned}$$

Fig. 5. Coverage checking

3.2 Checking Guard Trees

The next step in [Figure 1](#) is to transform the guard tree into an *annotated tree*, Ant , and an *uncovered set*, Θ . Taking the latter first, the uncovered set describes all the input values of the match that are not covered by the match. We use the language of *refinement types* to describe this set (see [Figure 3](#)). A refinement type $\Theta = \langle x_1:\tau_1, \dots, x_n:\tau_n \mid \Phi \rangle$ denotes the vector of values $x_1 \dots x_n$ that satisfy the predicate Φ . For example:

$$\begin{aligned} \langle x:\text{Bool} \mid \checkmark \rangle &\text{ denotes } \{\perp, \text{True}, \text{False}\} \\ \langle x:\text{Bool} \mid x \neq \perp \rangle &\text{ denotes } \{\text{True}, \text{False}\} \\ \langle x:\text{Bool} \mid x \neq \perp \wedge \text{True} \leftarrow x \rangle &\text{ denotes } \{\text{True}\} \\ \langle mx:\text{Maybe Bool} \mid mx \neq \perp \wedge \text{Just } x \leftarrow mx \wedge x \neq \perp \rangle &\text{ denotes } \{\text{Just True}, \text{Just False}\} \end{aligned}$$

The syntax of Φ is given in [Figure 3](#). It consists of a collection of *literals* φ , combined with conjunction and disjunction. Unconventionally, however, a literal may bind one or more variables, and those bindings are in scope in conjunctions to the right. This can readily be formalised by giving a type system for Φ , and we do so in ??, where we define satisfiability of Φ in formal detail. The literal \checkmark means “true”, as illustrated above; while \times means “false”, so that $\langle \Gamma \mid \times \rangle$ denotes the empty set \emptyset .

The uncovered set function $\mathcal{U}(\Theta, t)$, defined in [Figure 5](#), computes a refinement type describing the values in Θ that are not covered by the guard tree t . It is defined by a simple recursive descent over the guard tree, using the operation $\Theta \hat{\wedge} \varphi$ (also defined in [Figure 5](#)) to extend Θ with an extra literal φ .

While \mathcal{U} finds a refinement type describing values that are *not* matched by a guard tree (its set of *Uncovered* values), the function \mathcal{A} finds refinements describing values that *are* matched by a

guard tree, or that cause matching to diverge. It does so by producing an *annotated tree* (hence \mathcal{A} nnotate), whose syntax is given in Figure 3. An annotated tree has the same general structure as the guard tree from whence it came: in particular the top-to-bottom compositions \sqsupset are in the same places. But in an annotated tree, each $\rightarrow \Theta k$ leaf is annotated with a refinement type Θ describing the input values that will lead to that right-hand side; and each $\rightarrow \Theta \zeta$ node is annotated with a refinement type that describes the input values on which matching will diverge. Once again, \mathcal{A} can be defined by a simple recursive descent over the guard tree (Figure 5), but note that the second equation uses \mathcal{U} as an auxiliary function¹.

3.3 Reporting Errors

The final step in Figure 1 is to report errors. First, let us focus on reporting missing equations. Consider the following definition

```
data T = A | B | C
f (Just A) = True
```

If t is the guard tree obtained from f , the expression $\mathcal{U}(\langle x : \text{Maybe } T \mid \checkmark \rangle, t)$ will produce this refinement type describing values that are not matched:

$$\Theta_f = \langle x : \text{Maybe } T \mid x \neq \perp \wedge (x \neq \text{Just} \vee (\text{Just } y \leftarrow x \wedge y \neq \perp \wedge (y \neq A \vee (A \leftarrow y \wedge \times)))) \rangle$$

This is not very helpful to report to the user. It would be far preferable to produce one or more concrete *inhabitants* of Θ_f to report, something like this:

```
Missing equations for function 'f':
  f Nothing = ...
  f (Just B) = ...
  f (Just C) = ...
```

Generating these inhabitants is the main challenge. It is done by $\mathcal{G}(\Theta)$ in Figure 6, which we discuss next in Section 3.4. But first notice that, by calling the very same function \mathcal{G} , we can readily define the function \mathcal{R} , which reports a triple of (accessible, inaccessible, Redundant) GRHSs, as needed in our overall pipeline (Figure 1). \mathcal{R} is defined in Figure 6:

- Having reached a leaf $\rightarrow \Theta k$, if the refinement type Θ is uninhabited ($\mathcal{G}(\Theta) = \emptyset$), then no input values can cause execution to reach right-hand side k , and it is redundant.
- Having reached a node $\rightarrow \Theta \zeta$ t , if Θ is inhabited there is a possibility of divergence. Now suppose that all the GRHSs in t are redundant. Then we should pick the first of them and mark it as inaccessible.
- The case for \sqsupset_u^t follows by congruence: just combine the classifications of t and u .

To illustrate the second case consider u' from Section 2.3.1 and its annotated tree:

$$\begin{array}{l}
 u' () \mid \text{False} = 1 \\
 \quad \mid \text{False} = 2 \\
 u' _ = 3
 \end{array}
 \rightsquigarrow
 \begin{array}{c}
 \begin{array}{c} \Theta_1 \zeta \end{array} \rightarrow \begin{array}{c} \Theta_2 \ 1 \\ \Theta_3 \ 2 \\ \Theta_4 \ 3 \end{array}
 \end{array}$$

Θ_2 and Θ_3 are uninhabited (because of the False guards). But we cannot delete both GRHSs as redundant, because that would make the call $u' \perp$ return 3 rather than diverging. Rather, we want to report the first GRHSs as inaccessible, leaving all the others as redundant.

¹Our implementation avoids this duplicated work – see Section 5.1 – but the formulation in Figure 5 emphasises clarity over efficiency.

Collect accessible (\bar{k}), inaccessible (\bar{n}) and Redundant (\bar{m}) GRHSs

$$\boxed{\mathcal{R}(u) = (\bar{k}, \bar{n}, \bar{m})}$$

$$\mathcal{R}(\rightarrow \Theta n) = \begin{cases} (\epsilon, \epsilon, n), & \text{if } \mathcal{G}(\Theta) = \emptyset \\ (n, \epsilon, \epsilon), & \text{otherwise} \end{cases}$$

$$\mathcal{R}(\sqsubseteq_u^t) = (\bar{k} \bar{k}', \bar{n} \bar{n}', \bar{m} \bar{m}') \text{ where } \begin{cases} (\bar{k}, \bar{n}, \bar{m}) = \mathcal{R}(t) \\ (\bar{k}', \bar{n}', \bar{m}') = \mathcal{R}(u) \end{cases}$$

$$\mathcal{R}(\rightarrow \Theta \zeta \rightarrow t) = \begin{cases} (\epsilon, m, \bar{m}'), & \text{if } \mathcal{G}(\Theta) \neq \emptyset \text{ and } \mathcal{R}(t) = (\epsilon, \epsilon, m \bar{m}') \\ \mathcal{R}(t), & \text{otherwise} \end{cases}$$

Normalised refinement type syntax

$$\begin{array}{lll} \nabla & ::= & \times \mid \langle \Gamma \parallel \Delta \rangle & \text{Normalised refinement type} \\ \Delta & ::= & \emptyset \mid \Delta, \delta & \text{Set of constraints} \\ \delta & ::= & \gamma \mid x \approx K \bar{a} \bar{y} \mid x \not\approx K \mid x \approx \perp \mid x \not\approx \perp \mid x \approx y & \text{Constraints} \end{array}$$

Generate inhabitants of Θ

$$\boxed{\mathcal{G}(\Theta) = \mathcal{P}(\bar{p})}$$

$$\mathcal{G}(\langle \Gamma \mid \Phi \rangle) = \{\mathcal{E}(\nabla, \text{dom}(\Gamma)) \mid \nabla \in \mathcal{N}(\langle \Gamma \parallel \emptyset \rangle, \Phi)\}$$

Normalise Φ into ∇ s

$$\boxed{\mathcal{N}(\nabla, \Phi) = \mathcal{P}(\nabla)}$$

$$\mathcal{N}(\nabla, \varphi) = \begin{cases} \{\langle \Gamma' \parallel \Delta' \rangle\} & \text{where } \langle \Gamma' \parallel \Delta' \rangle = \nabla \oplus_{\varphi} \varphi \\ \emptyset & \text{otherwise} \end{cases}$$

$$\mathcal{N}(\nabla, \Phi_1 \wedge \Phi_2) = \bigcup \{\mathcal{N}(\nabla', \Phi_2) \mid \nabla' \in \mathcal{N}(\nabla, \Phi_1)\}$$

$$\mathcal{N}(\nabla, \Phi_1 \vee \Phi_2) = \mathcal{N}(\nabla, \Phi_1) \cup \mathcal{N}(\nabla, \Phi_2)$$

Expand variables to Pat with ∇

$$\boxed{\mathcal{E}(\nabla, x) = p, \quad \mathcal{E}(\nabla, \bar{x}) = \bar{p}}$$

$$\mathcal{E}(\nabla, \bar{x}) = \overline{\mathcal{E}(\nabla, x)}$$

$$\mathcal{E}(\langle \Gamma \parallel \Delta \rangle, x) = \begin{cases} K \mathcal{E}(\langle \Gamma \parallel \Delta \rangle, \bar{y}) & \text{if } \Delta(x) \approx K \bar{a} \bar{y} \in \Delta \\ - & \text{otherwise} \end{cases}$$

Finding the representative of a variable in Δ

$$\boxed{\Delta(x) = y}$$

$$\Delta(x) = \begin{cases} \Delta(y) & x \approx y \in \Delta \\ x & \text{otherwise} \end{cases}$$

Fig. 6. Generating inhabitants of Θ via ∇

3.4 Generating Inhabitants of a Refinement Type

Thus far, all our functions have been very simple, syntax-directed transformations, but they all ultimately depend on the single function \mathcal{G} , which does the real work. That is our new focus. As [Figure 6](#) shows, $\mathcal{G}(\Theta)$ takes a refinement type $\Theta = \langle \Gamma \mid \Phi \rangle$ and returns a (possibly-empty) set of patterns \bar{p} (syntax in [Figure 3](#)) that give the shape of values that inhabit Θ . We do this in two steps:

- Flatten Θ into a set of *normalised refinement types* ∇ , by the call $\mathcal{N}(\langle \Gamma \parallel \emptyset \rangle, \Phi)$; see [Section 3.6](#).
- For each such ∇ , expand Γ into a list of patterns, by the call $\mathcal{E}(\nabla, \text{dom}(\Gamma))$; see [Section 3.5](#).

A normalised refinement type ∇ is either empty (\times) or of the form $\langle \Gamma \parallel \Delta \rangle$. It is similar to a refinement type $\Theta = \langle \Gamma \mid \Phi \rangle$, but it takes a much more restricted form ([Figure 6](#)):

- Δ is simply a conjunction of literals δ ; there are no disjunctions. Instead, disjunction reflects in the fact that \mathcal{N} returns a *set* of normalised refinement types.

Beyond these syntactic differences, we enforce the following invariants on a $\nabla = \langle \Gamma \parallel \Delta \rangle$:

- 11 *Mutual compatibility*: No two constraints in Δ should *conflict* with each other, where $x \approx \perp$ conflicts with $x \neq \perp$, and $x \approx K _ _$ conflicts with $x \neq K$, for all x .
- 12 *Inhabitation*: If $x:\tau \in \Gamma$ and τ reduces to a data type under type constraints in Δ , there must be at least one constructor K (or \perp) which x can be instantiated to without contradicting I1; see [Section 3.7](#).
- 13 *Triangular form*: A $x \approx y$ constraint implies absence of any other constraint mentioning x in its left-hand side.
- 14 *Single solution*: There is at most one positive constructor constraint $x \approx K \bar{a} \bar{y}$ for a given x .

Invariants I1 and I2 prevent Δ being self-contradictory, so that ∇ (which denotes a set of values) is uninhabited. We use $\nabla = \times$ to represent an uninhabited refinement type. Invariants I3 and I4 require Δ to be in solved form, from which it is easy to “read off” a value that inhabits it – this reading-off step is performed by \mathcal{E} ([Section 3.5](#)).

The structure is directly analogous to the structure of the standard unification algorithm. In unification we start with a set of equalities between types (analogous to Θ) and, by unification, normalise it to a substitution (analogous to ∇). That substitution can itself be regarded as a set of equalities, but in a restricted form. And indeed our normalisation algorithm (described in [Section 3.6](#)) is a form of generalised unification.

Notice that we allow Δ to contain variable/variable equalities $x \approx y$, providing a function $\Delta(x)$ (defined in [Figure 6](#)) that follows these indirections to find the “representative” of x in Δ . A perfectly viable alternative would be to omit such indirections from Δ and instead aggressively substitute them away.

3.5 Expanding a Normalised Refinement Type to a Pattern

Expanding a match variable x under ∇ to a pattern, by calling \mathcal{E} in [Figure 6](#), is straightforward and overloaded to operate similarly on multiple match variables. When there is a solution like $\Delta(x) \approx \text{Just } y$ in Δ for the match variable x of interest, recursively expand y and wrap it in a `Just`. Invariant I4 guarantees that there is at most one such solution and \mathcal{E} is well-defined. When there is no solution for x , return `_`. See [Section 5.4](#) for how we improve on that in the implementation by taking negative information into account.

3.6 Normalising a Refinement Type

Normalisation, carried out by \mathcal{N} in [Figure 6](#), is largely a matter of repeatedly adding a literal φ to a normalised type, thus $\nabla \oplus_{\varphi} \varphi$. This function is where all the work is done, in [Figure 7](#). It does so by expressing a literal φ in terms of simpler constraints δ , and calling out to \oplus_{δ} to add the

Add a formula literal to ∇ $\nabla \oplus_{\varphi} \varphi = \nabla$

$$\begin{aligned} \nabla \oplus_{\varphi} \times &= \times & (1) \\ \nabla \oplus_{\varphi} \checkmark &= \nabla & (2) \\ \langle \Gamma \parallel \Delta \rangle \oplus_{\varphi} K \bar{a} \bar{y} \bar{y} : \bar{\tau} \leftarrow x &= \langle \Gamma, \bar{a}, \bar{y} : \bar{\tau} \parallel \Delta \rangle \oplus_{\delta} \bar{y} \oplus_{\delta} \overline{y' \neq \perp} \oplus_{\delta} x \approx K \bar{a} \bar{y} & (3) \\ &\quad \text{where } \bar{y}' \subseteq \bar{y} \text{ bind strict fields} \\ \langle \Gamma \parallel \Delta \rangle \oplus_{\varphi} \text{let } x : \tau = K \bar{\sigma} \bar{y} \bar{e} &= \langle \Gamma, x : \tau, \bar{a} \parallel \Delta \rangle \oplus_{\delta} \bar{a} \sim \bar{\sigma} \oplus_{\delta} x \neq \perp \oplus_{\delta} x \approx K \bar{a} \bar{y} & (4) \\ &\quad \oplus_{\varphi} \overline{\text{let } y : \tau' = e} \quad \text{where } \bar{a} \bar{y} \text{ fresh, } \bar{e} : \tau' \\ \langle \Gamma \parallel \Delta \rangle \oplus_{\varphi} \text{let } x : \tau = y &= \langle \Gamma, x : \tau \parallel \Delta \rangle \oplus_{\delta} x \approx y & (5) \\ \langle \Gamma \parallel \Delta \rangle \oplus_{\varphi} \text{let } x : \tau = e &= \langle \Gamma, x : \tau \parallel \Delta \rangle & (6) \\ \langle \Gamma \parallel \Delta \rangle \oplus_{\varphi} \varphi &= \langle \Gamma \parallel \Delta \rangle \oplus_{\delta} \varphi & (7) \end{aligned}$$

Add a constraint to ∇ $\nabla \oplus_{\delta} \delta = \nabla$

$$\begin{aligned} \times \oplus_{\delta} \delta &= \times & (8) \\ \langle \Gamma \parallel \Delta \rangle \oplus_{\delta} \gamma &= \begin{cases} \langle \Gamma \parallel (\Delta, \gamma) \rangle & \text{if type checker deems } \gamma \text{ compatible with } \Delta \\ & \text{and } \forall x \in \text{dom}(\Gamma) : \langle \Gamma \parallel (\Delta, \gamma) \rangle \vdash x \text{ inh} \\ \times & \text{otherwise} \end{cases} & (9) \\ \langle \Gamma \parallel \Delta \rangle \oplus_{\delta} x \approx K \bar{a} \bar{y} &= \begin{cases} \langle \Gamma \parallel \Delta \rangle \oplus_{\delta} \overline{a \sim b} \oplus_{\delta} \overline{y \approx z} & \text{if } \Delta(x) \approx K \bar{b} \bar{z} \in \Delta \\ \times & \text{if } \Delta(x) \approx K' \bar{b} \bar{z} \in \Delta \text{ and } K \neq K' \\ \times & \text{if } \Delta(x) \neq K \in \Delta \\ \langle \Gamma \parallel (\Delta, \Delta(x) \approx K \bar{a} \bar{y}) \rangle & \text{otherwise} \end{cases} & (10) \\ \langle \Gamma \parallel \Delta \rangle \oplus_{\delta} x \neq K &= \begin{cases} \times & \text{if } \Delta(x) \approx K \bar{a} \bar{y} \in \Delta \\ \times & \text{if not } \langle \Gamma \parallel (\Delta, \Delta(x) \neq K) \rangle \vdash x \text{ inh} \\ \langle \Gamma \parallel (\Delta, \Delta(x) \neq K) \rangle & \text{otherwise} \end{cases} & (11) \\ \langle \Gamma \parallel \Delta \rangle \oplus_{\delta} x \approx \perp &= \begin{cases} \times & \text{if } \Delta(x) \neq \perp \in \Delta \\ \langle \Gamma \parallel (\Delta, \Delta(x) \approx \perp) \rangle & \text{otherwise} \end{cases} & (12) \\ \langle \Gamma \parallel \Delta \rangle \oplus_{\delta} x \neq \perp &= \begin{cases} \times & \text{if } \Delta(x) \approx \perp \in \Delta \\ \times & \text{if not } \langle \Gamma \parallel (\Delta, \Delta(x) \neq \perp) \rangle \vdash x \text{ inh} \\ \langle \Gamma \parallel (\Delta, \Delta(x) \neq \perp) \rangle & \text{otherwise} \end{cases} & (13) \\ \langle \Gamma \parallel \Delta \rangle \oplus_{\delta} x \approx y &= \begin{cases} \langle \Gamma \parallel \Delta \rangle & \text{if } x' = y' \\ \langle \Gamma \parallel ((\Delta \setminus x'), x' \approx y') \rangle \oplus_{\delta} (\Delta|_{x'} [y'/x']) & \text{otherwise} \end{cases} & (14) \\ &\quad \text{where } x' = \Delta(x) \text{ and } y' = \Delta(y) \end{aligned}$$

$\Delta \setminus x = \Delta$	$\Delta _{x= \Delta}$
$\emptyset \setminus x = \emptyset$	$\emptyset _x = \emptyset$
$(\Delta, x \approx K \bar{a} \bar{y}) \setminus x = \Delta \setminus x$	$(\Delta, x \approx K \bar{a} \bar{y}) _x = \Delta _x, x \approx K \bar{a} \bar{y}$
$(\Delta, x \neq K) \setminus x = \Delta \setminus x$	$(\Delta, x \neq K) _x = \Delta _x, x \neq K$
$(\Delta, x \approx \perp) \setminus x = \Delta \setminus x$	$(\Delta, x \approx \perp) _x = \Delta _x, x \approx \perp$
$(\Delta, x \neq \perp) \setminus x = \Delta \setminus x$	$(\Delta, x \neq \perp) _x = \Delta _x, x \neq \perp$
$(\Delta, \delta) \setminus x = (\Delta \setminus x), \delta$	$(\Delta, \delta) _x = \Delta _x$

Fig. 7. Adding a constraint to the normalised refinement type ∇

simpler constraints to ∇ . \mathcal{N} , \oplus_φ and \oplus_δ all work on the principle that if the incoming ∇ satisfies the Invariants I1 to I4 from [Section 3.4](#), then either the resulting ∇ is \times or it satisfies I1 to I4.

In Equation (3), a pattern guard extends the context and adds suitable type constraints and a positive constructor constraint arising from the binding. Equation (4) of \oplus_φ performs some limited, but important reasoning about let bindings: it flattens possibly nested constructor applications, such as let $x = \text{Just True}$, and asserts that such constructor applications can't be \perp . Note that Equation (6) simply discards let bindings that cannot be expressed in ∇ ; we'll see an extension in [Section 4.3](#) that avoids this information loss.

That brings us to the prime unification procedure, \oplus_δ . When adding $x \approx \text{Just } y$, Equation (10), the unification procedure will first look for a solution for x with *that same constructor*. Let's say there is $\Delta(x) \approx \text{Just } u \in \Delta$. Then \oplus_δ operates on the transitively implied equality $\text{Just } y \approx \text{Just } u$ by equating type and term variables with new constraints, i.e. $y \approx u$. The original constraint, although not conflicting, is not added to the normalised refinement type because of I3.

If there is a solution involving a different constructor like $\Delta(x) \approx \text{Nothing}$ or if there was a negative constructor constraint $\Delta(x) \not\approx \text{Just}$, the new constraint is incompatible with the existing solution. Otherwise, the constraint is compatible and is added to Δ .

Adding a negative constructor constraint $x \not\approx \text{Just}$ is quite similar (Equation (11)), except that we have to make sure that x still satisfies I2, which is checked by the $\nabla \vdash \Delta(x)$ inh judgment (cf. [Section 3.7](#)) in [Figure 8](#). Handling positive and negative constraints involving \perp is analogous.

Adding a type constraint γ (Equation (9)) entails calling out to the type checker to assert that the constraint is consistent with existing type constraints. Afterwards, we have to ensure I2 is upheld for *all* variables in the domain of Γ , because the new type constraint could have rendered a type empty. To demonstrate why this is necessary, imagine we have $\langle x : a \parallel x \not\approx \perp \rangle$ and try to add $a \sim \text{Void}$. Although the type constraint is consistent, x in $\langle x : a \parallel x \not\approx \perp, a \sim \text{Void} \rangle$ is no longer inhabited. There is room for being smart about which variables we have to re-check: For example, we can exclude variables whose type is a non-GADT data type.

Equation (14) of \oplus_δ equates two variables ($x \approx y$) by merging their equivalence classes. Consider the case where x and y aren't in the same equivalence class. Then $\Delta(y)$ is arbitrarily chosen to be the new representative of the merged equivalence class. To uphold I3, all constraints mentioning $\Delta(x)$ have to be removed and renamed in terms of $\Delta(y)$ and then re-added to Δ , one of which in turn might uncover a contradiction.

3.7 Testing for Inhabitation

The process for adding a constraint to a normalised type above (which turned out to be a unification procedure in disguise) makes use of an *inhabitation test* $\nabla \vdash x$ inh, depicted in [Figure 8](#). This tests whether there are any values of x that satisfy ∇ . If not, ∇ does not uphold I2. For example, the conjunction $x \not\approx \text{Just}, x \not\approx \text{Nothing}, x \not\approx \perp$ does not satisfy I2, because no value of x satisfies all those constraints.

The \vdash_{BOT} judgment of $\nabla \vdash x$ inh tries to instantiate x to \perp to conclude that x is inhabited. \vdash_{INST} instantiates x to one of its data constructors. That will only work if its type ultimately reduces to a data type under the type constraints in ∇ . Rule $\vdash_{\text{NO CPL}}$ will accept unconditionally when its type is not a data type, i.e. for $x : \text{Int} \rightarrow \text{Int}$.

Note that the outlined approach is complete in the sense that $\nabla \vdash x$ inh is derivable (if and) only if x is actually inhabited in ∇ , because that means we don't have any ∇ s floating around in the checking process that actually aren't inhabited and trigger false positive warnings. But that also means that the \vdash inh relation is undecidable! Consider the following example:

$$\begin{array}{c}
\text{Test if } x \text{ is inhabited considering } \nabla \quad \boxed{\nabla \vdash x \text{ inh}} \\
\frac{\langle \Gamma \parallel \Delta \rangle \oplus_{\delta} x \approx \perp \neq \times}{\langle \Gamma \parallel \Delta \rangle \vdash x \text{ inh}} \vdash_{\text{BOT}} \quad \frac{x : \tau \in \Gamma \quad \text{Cons}(\langle \Gamma \parallel \Delta \rangle, \tau) = \perp}{\langle \Gamma \parallel \Delta \rangle \vdash x \text{ inh}} \vdash_{\text{NoCPL}} \\
\\
\frac{x : \tau \in \Gamma \quad K \in \text{Cons}(\langle \Gamma \parallel \Delta \rangle, \tau) \quad \text{Inst}(\langle \Gamma \parallel \Delta \rangle, x, K) \neq \times}{\langle \Gamma \parallel \Delta \rangle \vdash x \text{ inh}} \vdash_{\text{INST}} \\
\\
\text{Find data constructors of } \tau \quad \boxed{\text{Cons}(\langle \Gamma \parallel \Delta \rangle, \tau) = \bar{K}} \\
\text{Cons}(\langle \Gamma \parallel \Delta \rangle, \tau) = \begin{cases} \bar{K} & \tau = T \bar{\sigma} \text{ and } T \text{ data type with constructors } \bar{K} \\ & \text{(after normalisation according to the type constraints in } \Delta \text{)} \\ \perp & \text{otherwise} \end{cases} \\
\\
\text{Instantiate } x \text{ to data constructor } K \quad \boxed{\text{Inst}(\nabla, x, K) = \nabla} \\
\text{Inst}(\langle \Gamma \parallel \Delta \rangle, x, K) = \langle \Gamma, \bar{a}, \bar{y} : \bar{\sigma} \parallel \Delta \rangle \oplus_{\delta} \tau_x \sim \tau \oplus_{\delta} \bar{y} \oplus_{\delta} x \approx K \bar{a} \bar{y} \oplus_{\delta} \bar{y}' \neq \perp \\
\text{where } K : \forall \bar{a}. \bar{y} \Rightarrow \bar{\sigma} \rightarrow \tau, \bar{a} \bar{y} \text{ fresh, } x : \tau_x \in \Gamma, \bar{y}' \subseteq \bar{y} \text{ bind strict fields}
\end{array}$$

Fig. 8. Testing for inhabitation

```

data T = MkT !T
f :: SMaybe T → ()
f SNothing = ()

```

This is exhaustive, because T is an uninhabited type. Upon adding the constraint $x \neq \text{SNothing}$ on the match variable x via \oplus_{δ} , we perform an inhabitation test, which tries to instantiate the SJust constructor via \vdash_{INST} . That implies adding (via \oplus_{δ}) the constraints $x \approx \text{SJust } y, y \neq \perp$, the latter of which leads to an inhabitation test on y . That leads to instantiation of the MkT constructor, which leads to constraints $y \approx \text{MkT } z, z \neq \perp$, and so on for z etc.. An infinite chain of fruitless instantiation attempts!

In practice, we implement a fuel-based approach that conservatively assumes that a variable is inhabited after n such iterations (we have $n = 100$ for list-like constructors and $n = 1$ otherwise) and consider supplementing that with a simple termination analysis to detect uninhabited data types like T in the future.

4 EXTENSIONS

LYG is well equipped to handle the fragment of Haskell it was designed to handle. But GHC (and other languages, for that matter) extends Haskell in non-trivial ways. This section exemplifies easy accommodation of new language features and measures to increase precision of the checking process, demonstrating the modularity and extensibility of our approach.

4.1 Long-Distance Information

Coverage checking should also work for **case** expressions and nested function definitions, like

```

f True = 1
f x    = ... (case x of { False → 2; True → 3 }) ...

```

GMTM and unextended LYG will not produce any warnings for this definition. But the reader can easily make the “long distance connection” that the last GRHS of the `case` expression is redundant! That follows by context-sensitive reasoning, knowing that x was already matched against `True`.

In terms of LYG, the input values of the second GRHS of f , described by $\Theta_2 = \langle x : \text{Bool} \mid x \neq \perp, x \neq \text{True} \rangle$, encode the information we are after: we just have to start checking the `case` expression starting from Θ_2 as the initial set of reaching values instead of $\langle x : \text{Bool} \mid \checkmark \rangle$. We already need Θ_2 to determine whether the second GRHS of f is accessible, so long-distance information comes almost for free.

4.2 Empty Case

As can be seen in [Figure 2](#), Haskell function definitions need to have at least one clause. That leads to an awkward situation when pattern matching on empty data types, like `Void`:

```
absurd1 _ = undefined           absurd1, absurd2, absurd3 :: Void → a
absurd2 !_ = undefined         absurd3 x = case x of { }
```

`absurd1` returns `undefined` when called with \perp , thus masking the original \perp with the error thrown by `undefined`. `absurd2` would diverge alright, but LYG will report its GRHS as inaccessible! Hence GHC provides an extension, called `EmptyCase`, that allows the definition of `absurd3` above. Such a `case` expression without any alternatives evaluates its argument to WHNF and crashes when evaluation returns.

It is quite easy to see that `Gdt` lacks expressive power to desugar `EmptyCase` into, since all leaves in a guard tree need to have corresponding GRHSs. Therefore, we need to introduce empty alternatives \bullet_{Gdt} to `Gdt` and \bullet_{Ant} to `Ant`. This is how they affect the checking process:

$$\mathcal{U}(\Theta, \bullet_{\text{Gdt}}) = \Theta \quad \mathcal{A}(\Theta, \bullet_{\text{Gdt}}) = \bullet_{\text{Ant}}$$

Since `EmptyCase`, unlike regular `case`, evaluates its scrutinee to WHNF *before* matching any of the patterns, the set of reaching values is refined with a $x \neq \perp$ constraint *before* traversing the guard tree, thus checking starts starts with $\mathcal{U}(\langle \Gamma \mid x \neq \perp \rangle, \bullet_{\text{Gdt}})$.

4.3 View Patterns

Our source syntax had support for view patterns to start with (cf. [Figure 2](#)). And even the desugaring we gave as part of the definition of \mathcal{D} in [Figure 4](#) is accurate. But this desugaring alone is insufficient for the checker to conclude that `safeLast` from [Section 2.2.1](#) is an exhaustive definition! To see why, let’s look at its guard tree:

```
├─ let y1 = reverse x1, !y1, Nothing ← y1 ───────────→ 1
└─ let y2 = reverse x1, !y2, Just t1 ← y2, !t1, (t2, t3) ← t1 ─→ 2
```

As far as LYG is concerned, the matches on both y_1 and y_2 are non-exhaustive. But that’s actually too conservative: Both bind the same value! By making the connection between y_1 and y_2 , the checker could infer that the match was exhaustive.

This can be fixed by maintaining equivalence classes of semantically equivalent expressions in Δ , similar to what we already do for variables. We simply extend the syntax of δ and change the last `let` case of Θ_φ . Then we can handle the new constraint in Θ_δ , as follows:

$$\delta = \dots \mid e \approx x \quad \langle \Gamma \parallel \Delta \rangle \Theta_\varphi \text{ let } x : \tau = e = \langle \Gamma, x : \tau \parallel \Delta \rangle \Theta_\delta e \approx x$$

$$\langle \Gamma \parallel \Delta \rangle \Theta_\delta e \approx x = \begin{cases} \langle \Gamma \parallel \Delta \rangle \Theta_\delta x \approx y, & \text{if } e' \approx y \in \Delta \text{ and } e \equiv_\Delta e' \\ \langle \Gamma \parallel \Delta, e \approx \Delta(x) \rangle, & \text{otherwise} \end{cases}$$

Where \equiv_{Δ} is (an approximation to) semantic equivalence modulo substitution under Δ . A clever data structure is needed to answer queries of the form $e \approx _ \in \Delta$, efficiently. In our implementation, we use a trie to index expressions rapidly [Peyton Jones and Graf 2023] and sacrifice reasoning modulo Δ in doing so. Plugging in an SMT solver to decide \equiv_{Δ} would be more precise, but certainly less efficient.

4.4 Pattern Synonyms

To accommodate checking of pattern synonyms P , we first have to extend the source syntax and IR syntax by adding the syntactic concept of a *ConLike*:

$$\begin{array}{ll} cl ::= K \mid P & P \in \text{PS} \\ pat ::= x \mid _ \mid \text{cl } \overline{pat} \mid x@pat \mid \dots & C \in \text{CL} ::= K \mid P \\ & p \in \text{Pat} ::= _ \mid C \overline{p} \mid \dots \end{array}$$

Assuming every definition encountered so far is changed to handle ConLikes C instead of data constructors K , everything should work fine. So why introduce the new syntactic variant in the first place? Consider

pattern $P = ()$

pattern $Q = ()$

$n = \text{case } P \text{ of } Q \rightarrow 1; P \rightarrow 2$

If P and Q were data constructors, the first alternative of the **case** would be redundant, because P cannot match Q . But pattern synonyms are quite different: a value produced by P might match a pattern Q , as indeed is the case in this example.

Our solution is a conservative one: we weaken the test that sends ∇ to \times of Equation (10) in the definition of \oplus_{δ} dealing with positive ConLike constraints $x \approx C \overline{a} \overline{y}$:

$$\langle \Gamma \parallel \Delta \rangle \oplus_{\delta} x \approx C \overline{a} \overline{y} = \begin{cases} \langle \Gamma \parallel \Delta \rangle \oplus_{\delta} \overline{a \sim b} \oplus_{\delta} \overline{y \approx z} & \text{if } \Delta(x) \approx C \overline{b} \overline{z} \in \Delta \\ \times & \text{if } \Delta(x) \approx C' \overline{b} \overline{z} \in \Delta \text{ and } C \cap C' = \emptyset \\ \times & \text{if } \Delta(x) \not\approx C \in \Delta \\ \langle \Gamma \parallel (\Delta, \Delta(x) \approx C \overline{a} \overline{y}) \rangle & \text{otherwise} \end{cases}$$

where the suggestive notation $C \cap C' = \emptyset$ is only true iff C and C' are distinct data constructors.

Note that the slight relaxation means that the constructed ∇ might violate $I4$, specifically when $C \cap C' \neq \emptyset$. In practice that condition only matters for the well-definedness of \mathcal{E} , which in case of multiple solutions (i.e. $x \approx P, x \approx Q$) has to commit to one of them for the purposes of reporting warnings. Fixing that requires a bit of boring engineering.

Another subtle point appears in rule (\dagger) in Figure 4: should we or should we not add a bang guard for pattern synonyms? There is no way to know without breaking the abstraction offered by the synonym. In effect, its strictness or otherwise is part of its client-visible semantics. In our implementation, we have (thus far) compromised by assuming that all pattern synonyms are strict for the purposes of coverage checking [GHC issue 2019m].

4.5 COMPLETE Pragmas

In a sense, every algebraic data type defines its own builtin COMPLETE set, consisting of all its data constructors, so the coverage checker already manages a single COMPLETE set.

We have $\vdash \text{INST}$ from Figure 8 currently making sure that this COMPLETE set is in fact inhabited. We also have $\vdash \text{NO CPL}$ that handles the case when we can't find *any* COMPLETE set for the given

type (think $x :: \text{Int} \rightarrow \text{Int}$). The obvious way to generalise this is by looking up all COMPLETE sets attached to a type and check that none of them is completely covered:

$$\frac{\langle \Gamma \parallel \Delta \rangle \oplus_{\delta} x \approx \perp \neq \times}{\langle \Gamma \parallel \Delta \rangle \vdash x \text{ inh}} \vdash_{\text{BOT}} \quad \frac{x : \tau \in \Gamma \quad \text{Cons}(\langle \Gamma \parallel \Delta \rangle, \tau) = \overline{C_1, \dots, C_{n_i}^i}}{\langle \Gamma \parallel \Delta \rangle \vdash x \text{ inh}} \vdash_{\text{INST}}$$

$$\text{Cons}(\langle \Gamma \parallel \Delta \rangle, \tau) = \begin{cases} \overline{C_1, \dots, C_{n_i}^i} & \tau = T \bar{\sigma}; \text{ T type constructor with COMPLETE sets } \overline{C_1, \dots, C_{n_i}^i} \\ & \text{(after normalisation according to the type constraints in } \Delta \text{)} \\ \epsilon & \text{otherwise} \end{cases}$$

Cons was changed to return a list of all available COMPLETE sets, and \vdash_{INST} tries to find an inhabiting ConLike in each one of them in turn. Note that $\vdash_{\text{NO CPL}}$ is gone, because it coincides with \vdash_{INST} for the case where the list returned by Cons was empty. The judgment has become simpler and more general at the same time! A worry is that checking against multiple COMPLETE sets so frequently is computationally intractable. We will worry about that in [Section 5.3](#).

4.6 Literals

The source syntax in [Figure 9](#) deliberately left out literal patterns l . Literals are very similar to nullary data constructors, with one caveat: they don't come with a builtin COMPLETE set. Before [Section 4.5](#), that would have meant quite a bit of hand waving and complication to the $\vdash \text{ inh}$ judgment. Now, literals can be handled like disjoint pattern synonyms (i.e. $l_1 \cap l_2 = \emptyset$ for any two literals l_1, l_2) without a COMPLETE set!

We can even handle overloaded literals, but we will find ourselves in a similar situation as with pattern synonyms:

```
instance Num () where
  fromInteger _ = ()
  n = case (0 :: ()) of 1 → 1; 0 → 2
```

Considering overloaded literals to be disjoint would mean marking the first alternative as redundant, which is unsound. Hence we regard overloaded literals as possibly overlapping, so they behave exactly like nullary pattern synonyms without a COMPLETE set.

4.7 Newtypes

In Haskell, a newtype declares a new type that is completely isomorphic to, but distinct from, an existing type. For example:

```
newtype NT a = MkNT [a]
dup :: NT a → NT a
dup (MkNT xs) = MkNT (xs + xs)
```

Here the type $\text{NT } a$ is isomorphic to $[a]$. We convert to and fro using the “data constructor” MkNT , either as in a term or in a pattern.

To a first approximation, programmers interact with a newtype as if it was a data type with a single constructor with a single field. But the pattern-matching semantics of newtypes are different! Here are three key examples that distinguish newtypes from data types. Functions $g1, g2$ match on a *newtype* N , while functions $h1, h2$ match on a *data type* D :

$$\begin{aligned}
cl ::= K \mid \boxed{N} \quad & N \in \text{NT} \\
& C \in K \mid \boxed{N} \\
\mathcal{D}(x, N \text{ pat}_1 \dots \text{ pat}_n) = N \ y_1 \dots y_n \leftarrow x, \mathcal{D}(y_1, \text{ pat}_1), \dots, \mathcal{D}(y_n, \text{ pat}_n) \\
\Delta_{\text{NT}}(x) = & \begin{cases} \Delta_{\text{NT}}(y) & x \approx y \in \Delta \text{ or } x \approx N \bar{a} \ y \in \Delta \\ x & \text{otherwise} \end{cases} \\
\langle \Gamma \parallel \Delta \rangle \oplus_\varphi \text{ let } x:\tau = K \bar{\sigma} \bar{y} \bar{e} = \dots \text{ as before } \dots & \quad (4a) \\
\langle \Gamma \parallel \Delta \rangle \oplus_\varphi \text{ let } x:\tau = N \bar{\sigma} \bar{e} = \langle \Gamma, x:\tau, \bar{a} \parallel \Delta \rangle \oplus_\delta \bar{a} \sim \bar{\sigma} \oplus_\delta x \approx N \bar{a} \ y \oplus_\varphi \text{ let } y:\tau' = e & \quad (4b) \\
& \text{where } \bar{a} \ y \# \Gamma, e:\tau' \\
\langle \Gamma \parallel \Delta \rangle \oplus_\delta x \approx K \bar{a} \bar{y} = \dots \text{ as before } \dots & \quad (10a) \\
\langle \Gamma \parallel \Delta \rangle \oplus_\delta \boxed{x \approx N \bar{a} \ y} = \begin{cases} \langle \Gamma \parallel \Delta \rangle \oplus_\delta \bar{a} \sim \bar{b} \oplus_\delta y \approx z & \text{if } x' \approx N \bar{b} \ z \in \Delta \\ \langle \Gamma \parallel \Delta \rangle & \text{if } x' = \Delta_{\text{NT}}(y') \\ \langle \Gamma \parallel ((\Delta \setminus x'), x' \approx N \bar{a} \ y') \rangle \oplus_\delta (\Delta|_{x'} [y'/x']) & \text{otherwise} \end{cases} & \quad (10b) \\
& \text{where } x' = \Delta(x) \text{ and } y' = \Delta(y) \\
\langle \Gamma \parallel \Delta \rangle \oplus_\delta x \neq K = \dots \text{ as before } \dots & \quad (11a) \\
\langle \Gamma \parallel \Delta \rangle \oplus_\delta x \neq N = \times & \quad (11b) \\
\langle \Gamma \parallel \Delta \rangle \oplus_\delta x \approx \perp = \begin{cases} \times & \text{if } \Delta_{\text{NT}}(x) \neq \perp \in \Delta \\ \langle \Gamma \parallel (\Delta, \Delta_{\text{NT}}(x) \approx \perp) \rangle & \text{otherwise} \end{cases} & \quad (12) \\
\langle \Gamma \parallel \Delta \rangle \oplus_\delta x \neq \perp = \begin{cases} \times & \text{if } \Delta_{\text{NT}}(x) \approx \perp \in \Delta \\ \times & \text{if not } \langle \Gamma \parallel (\Delta, \Delta_{\text{NT}}(x) \neq \perp) \rangle \vdash x \text{ inh} \\ \langle \Gamma \parallel (\Delta, \Delta_{\text{NT}}(x) \neq \perp) \rangle & \text{otherwise} \end{cases} & \quad (13)
\end{aligned}$$

Fig. 9. Extending coverage checking to handle newtypes

newtype $N \ a = \text{MkN } a$

$g1 :: N \ \text{Void} \rightarrow \text{Bool} \rightarrow \text{Int}$

$g1 \ _ \quad \text{True} = 1$

$g1 \ (\text{MkN } _) \ \text{True} = 2 \quad \text{-- Redundant}$

$g1 \ !_ \quad \text{True} = 3 \quad \text{-- Inaccessible}$

$g2 :: N \ () \rightarrow \text{Bool} \rightarrow \text{Int}$

$g2 \ !(\text{MkN } _) \ \text{True} = 1$

$g2 \ (\text{MkN } !_) \ \text{True} = 2 \quad \text{-- Redundant}$

$g2 \ _ \quad _ = 3$

data $D \ a = \text{MkD } a$

$h1 :: D \ \text{Void} \rightarrow \text{Bool} \rightarrow \text{Int}$

$h1 \ _ \quad \text{True} = 1$

$h1 \ (\text{MkD } _) \ \text{True} = 2 \quad \text{-- Inaccessible}$

$h1 \ !_ \quad \text{True} = 3 \quad \text{-- Redundant}$

$h2 :: D \ () \rightarrow \text{Bool} \rightarrow \text{Int}$

$h2 \ !(\text{MkD } _) \ \text{True} = 1$

$h2 \ (\text{MkD } !_) \ \text{True} = 2 \quad \text{-- Inaccessible}$

$h2 \ _ \quad _ = 3$

If the first equation of $h1$ fails to match (because the second argument is False), the second equation may diverge when matching against $(\text{MkD } _)$ or may fail (because of the False), so the equation is inaccessible. The third equation is redundant. But for a newtype, the second equation of $g1$ will not evaluate the argument when matching against $(\text{MkN } _)$ and hence is redundant. The third equation will evaluate the first argument, which is surely bottom, so matching will diverge and the equation is inaccessible. A perhaps surprising consequence is that the definition of $g1$ is exhaustive,

because after N Void was deprived of its sole inhabitant $\perp = \text{MkN } \perp$ by the third GRHS, there is nothing left to match on (similarly for $h1$). Analogous subtle reasoning justifies the difference in warnings for $g2$ and $h2$.

Figure 9 outlines a solution that handles all these cases correctly:

- A newtype pattern match $N \text{ pat}_1 \dots \text{pat}_n$ is lazy: it does not force evaluation. So, compared to data constructor matches, the desugaring function \mathcal{D} omits the $!x$. Additionally, Equation (4) of \oplus_φ , responsible for reasoning about **let** bindings, has a special case for newtypes that omits the $x \neq \perp$ constraint.
- Similar in spirit to $\Delta(x)$, which chases variable equality constraints $x \approx y$, we now also occasionally need to look through positive newtype constructor constraints $x \approx N \bar{a} y$ with $\Delta_{\text{NT}}(x)$.
- The most important usage of $\Delta_{\text{NT}}(x)$ is in the changed Equations (12) and (13) of \oplus_δ , where we now check \perp constraints modulo $\Delta_{\text{NT}}(x)$.
- Equation (10) (previously handling $x \approx K \bar{a} \bar{y}$) and Equation (11) (previously handling $x \neq K$) have been split to account for newtype constructors.
- The first case of the new Equation (10b) handles any existing positive newtype constructor constraints in Δ , as with Equation (10). Take note that negative newtype constructor constraints may never occur in Δ because of Equation (11b), as explained in the next paragraph. The remaining two cases are reminiscent of Equation (14) ($x \approx y$). Provided there are neither positive nor negative newtype constructor constraints involving x , any remaining \perp constraints are moved from $\Delta(x)$ to the new representative $\Delta'_{\text{NT}}(x)$, which will be $\Delta'_{\text{NT}}(y)$ in the returned Δ' .
- The new Equation (11b) handles negative newtype constructor constraints by immediately rejecting. The reason it doesn't consider \perp as an inhabitant is that for \perp to be an inhabitant, it must be an inhabitant of the newtype's field. For that, we must have $x \approx K y$ for some y , which contradicts with the very constraint we want to add!

To see how these changes facilitate correct warnings for newtype matches, first consider the changed invariant I3 which ensures $\Delta_{\text{NT}}(x)$ is a well-defined function like $\Delta(x)$:

I3 *Triangular form*: Constraints of the form $x \approx y$ and $x \approx N \bar{a} y$ imply absence of any other constraint mentioning x in its left-hand side.

We want Δ to uphold the semantic equation $\perp \equiv N\perp$. In particular, whenever we have $x \approx N \bar{a} y$, we want $x \approx \perp$ iff $y \approx \perp$ (similarly for $x \neq \perp$). Equations (10b), (12) and (13) facilitate just that, modulo $\Delta_{\text{NT}}(x)$. Finally, a new invariant I5 relates positive newtype constructor equalities to \perp constraints:

I5 *Newtype erasure*: Whenever $x \approx N \bar{a} y \in \Delta$, we have $x \approx \perp \in \Delta$ if and only if $y \approx \perp \in \Delta$, and $x \neq \perp \in \Delta$ if and only if $y \neq \perp \in \Delta$.

An alternative design might take inspiration in the coercion semantics of GHC Core, a typed intermediate language of GHC based on System F, and compose coercions attached to \approx . However, that would entail deep changes to syntax as well as to the definition of \mathcal{E} to recover the newtype constructor patterns visible in source syntax.

4.8 Strictness, Divergence and Other Side-Effects

Instead of extending the source language, let's discuss ripping out a language feature for a change! So far, we have focused on Haskell as the source language, which is lazy by default. Although the difference in evaluation strategy of the source language becomes irrelevant after desugaring, it

raises the question of how much our approach could be simplified if we targeted a source language that was strict by default, such as OCaml, Lean, Idris, Rust, Python or C#.


On first thought, it is tempting to simply drop all parts related to laziness from the formalism, such as $!x$ from Grd and $\text{—} \zeta \text{—}$ from Ant. Actually, Ant and \mathcal{R} could vanish altogether and \mathcal{A} could just collect the redundant GRHS directly! Since there wouldn't be any bang guards, there is no reason to have $x \approx \perp$ and $x \not\approx \perp$ constraints either. Most importantly, the \vdash_{BOT} judgment form has to go, because \perp does not inhabit any types anymore.

And compiler writers for total languages such as Lean, Idris or Agda would live happily after: Talking about $x \not\approx \perp$ constraints made no sense there to begin with. Not so with OCaml or Rust, which are strict, non-total languages and allow arbitrary side-effects in expressions. Here's an example in OCaml:

```
let rec f p x =
  match x with
  | []                → []
  | hd :: _ when p hd && x = [] → [hd]
  | _ :: tl           → f p tl;;
```

Is the second clause redundant? It depends on whether p performs a side-effect, such as throwing an exception, diverging, or even releasing a mutex. We may not say without knowing the definition of p , so the second clause has an inaccessible RHS but is not redundant. It's a similar situation as in a lazy language, although the fact that side-effects only matter in the guard of a match clause (where we can put arbitrary expressions) makes the issue much less prominent.

We could come up with a desugaring function for OCaml that desugars the pattern match above to the following guard tree:



Compared to Haskell, note the lack of a bang guard on the match variable x . Instead, there's now a bang guard on t , the new temporary that stands for $p \text{ hd}$. The bang guard will keep alive the second clause of the guard tree and our algorithm would not classify the second clause as redundant, although it will be flagged as inaccessible. Since the RHS of a **let** guard, such as $p \text{ hd}$, might have arbitrary side-effects, equational reasoning is lost and we may no longer identify $p \text{ hs} \approx t$ as in Section 4.3.

Zooming out a bit more, desugaring of Haskell pattern matches using bang guards $!x$ can be understood as forcing *one specific effect*, namely divergence. In this work, we have given this side-effect a first-class treatment in our formalism in order to get accurate coverage warnings in a lazy language.

4.9 Or-patterns

Since this work appeared at ICFP in 2020, GHC 9.12 accumulated a new extension to the pattern language: *Or-patterns*². Or-patterns are an established language feature in many other languages such as OCaml and Python, and can be used as follows:

```
data LogLevel = Debug | Info | Error
notifyAdmin :: LogLevel → Bool
```

²<https://github.com/ghc-proposals/ghc-proposals/blob/master/proposals/0522-or-patterns.rst>

$$\begin{aligned}
pat & ::= \dots \mid \boxed{pat_1; pat_2} \\
t \in \text{Gdt} & ::= \dots \mid \boxed{d} \text{---} t \\
d \in \text{GrdDag} & ::= g \mid \boxed{(d_1 \text{ or } d_2)} \mid (d_1, d_2) \\
\mathcal{D}(x, (pat_1; pat_2)) & = (\mathcal{D}(x, pat_1) \text{ or } \mathcal{D}(x, pat_2)) \\
\boxed{C(\Theta, d) = \Theta} \\
C(\Theta, (d_1 \text{ or } d_2)) & = C(\Theta, d_1) \cup C(\mathcal{U}(\Theta, d_1), d_2) \\
C(\Theta, (d_1, d_2)) & = C(C(\Theta, d_1), d_2) \\
C(\Theta, !x) & = \Theta \dot{\wedge} (x \neq \perp) \\
C(\Theta, \text{let } x = e) & = \Theta \dot{\wedge} (\text{let } x = e) \\
C(\Theta, K \bar{a} \bar{y} \bar{y} : \bar{\tau} \leftarrow x) & = \Theta \dot{\wedge} (K \bar{a} \bar{y} \bar{y} : \bar{\tau} \leftarrow x) \\
\boxed{\mathcal{U}(\Theta, t) = \Theta, \quad \mathcal{U}(\Theta, d) = \Theta} \\
\mathcal{U}(\langle \Gamma \mid \Phi \rangle, \rightarrow n) & = \langle \Gamma \mid \times \rangle \\
\mathcal{U}(\Theta, \bigsqcup_{t_2}^{t_1}) & = \mathcal{U}(\mathcal{U}(\Theta, t_1), t_2) \\
\mathcal{U}(\Theta, \text{---}d\text{---}t) & = \mathcal{U}(\Theta, d) \cup \mathcal{U}(C(\Theta, d), t) \\
\mathcal{U}(\Theta, (d_1 \text{ or } d_2)) & = \mathcal{U}(\mathcal{U}(\Theta, d_1), d_2) \\
\mathcal{U}(\Theta, (d_1, d_2)) & = \mathcal{U}(\Theta, d_1) \cup \mathcal{U}(C(\Theta, d_1), d_2) \\
\mathcal{U}(\langle \Gamma \mid \Phi \rangle, !x) & = \langle \Gamma \mid \times \rangle \\
\mathcal{U}(\langle \Gamma \mid \Phi \rangle, \text{let } x = e) & = \langle \Gamma \mid \times \rangle \\
\mathcal{U}(\Theta, K \bar{a} \bar{y} \bar{y} : \bar{\tau} \leftarrow x) & = \Theta \dot{\wedge} (x \neq K) \\
\boxed{\mathcal{A}(\Theta, t) = u, \quad \mathcal{A}(\Theta, d) = \Theta} \\
\mathcal{A}(\Theta, \rightarrow n) & = \rightarrow \Theta n \\
\mathcal{A}(\Theta, \bigsqcup_{t_2}^{t_1}) & = \bigsqcup \mathcal{A}(\Theta, t_1) \\
& \quad \bigsqcup \mathcal{A}(\mathcal{U}(\Theta, t_1), t_2) \\
\mathcal{A}(\Theta, \text{---}d\text{---}t) & = \text{---} \mathcal{A}(\Theta, d) \dot{\wedge} \text{---} \mathcal{A}(C(\Theta, d), t) \\
\mathcal{A}(\Theta, (d_1 \text{ or } d_2)) & = \mathcal{A}(\Theta, d_1) \cup \mathcal{A}(\mathcal{U}(\Theta, d_1), d_2) \\
\mathcal{A}(\Theta, (d_1, d_2)) & = \mathcal{A}(\Theta, d_1) \cup \mathcal{A}(C(\Theta, d_1), d_2) \\
\mathcal{A}(\Theta, !x) & = \Theta \dot{\wedge} (x \approx \perp) \\
\mathcal{A}(\langle \Gamma \mid \Phi \rangle, \text{let } x = e) & = \langle \Gamma \mid \times \rangle \\
\mathcal{A}(\langle \Gamma \mid \Phi \rangle, K \bar{a} \bar{y} \bar{y} : \bar{\tau} \leftarrow x) & = \langle \Gamma \mid \times \rangle
\end{aligned}$$

Fig. 10. Extending coverage checking to handle Or-patterns

notifyAdmin Error = True
notifyAdmin (Debug; Info) = False

Here, the second clause matches when either Debug or Info matches the parameter. When the programmer later adds a new data constructor Warning to LogLevel, LYG should report the match in *notifyAdmin* as inexhaustive. This coverage warning prompts the programmer to make a conscious decision about which value should be returned for *notifyAdmin* Warning. That is far better

than the alternative of using a wildcard match for the last clause: doing so would silently define `notifyAdmin Warning = False`.

Or-patterns were an interesting real-world benchmark to see how well LYG scales to new language features. Previously in Figure 4, if one part of a pattern failed to match, the whole pattern would fail. As a result, the desugaring function \mathcal{D} could map a pattern into a (conjunctive) list of guards ($\overline{\text{Grd}}$), which was then exploded into a nesting of $\dashv\vdash g \dashv\vdash t$ forms, suitable for a single recursive definition of \mathcal{U} and \mathcal{A} . However, with Or-patterns, we need a way to encode (disjunctive) first-match semantics in the result of $\mathcal{D}(x, pat)$. Such first-match semantics is currently exclusive to the $\sqsubset_{t_2}^{t_1}$ guard tree form. So one way to desugar Or-patterns would be to desugar patterns into full guard trees instead of lists of guards. That would be akin to *exploding* each Or-pattern into two clauses. We would get the equality

$$\mathcal{D}(f (pat_a; pat_b) pat_c = rhs) = \sqsubset_{\mathcal{D}(x_1, pat_b), \mathcal{D}(x_2, pat_c) \rightarrow k_{rhs}}^{\mathcal{D}(x_1, pat_a), \mathcal{D}(x_2, pat_c) \rightarrow k_{rhs}},$$

thus duplicating the desugaring of pat_c . It is easy to see how a sequence of Or-patterns may lead to an exponential number of duplications of pat_c , leading to unacceptable checking performance. Hence we propose a different solution: *Guard DAGs* (directed-acyclic graphs).

Figure 10 defines the structure of guard DAGs (GrdDag) inductively. Now consider the function

$f :: \text{Ordering} \rightarrow \text{Ordering} \rightarrow \text{Int}$

$f (\text{LT}; \text{EQ}) (\text{EQ}; \text{GT}) = 1$

$f _ _ = 2$

The desugaring to guard trees according to Figure 10 is


$$\sqsubset_{\rightarrow 2}^{\rightarrow 1} (((!x_1, \text{LT} \leftarrow x_1) \text{ or } (!x_1, \text{EQ} \leftarrow x_1)), ((!x_2, \text{EQ} \leftarrow x_2) \text{ or } (!x_2, \text{GT} \leftarrow x_2)))$$

We define matching as follows:

- Matching g means matching a single guard $g \in \text{Grd}$, which was done by $\dashv\vdash g \dashv\vdash t$ previously. However, the new $\dashv\vdash d \dashv\vdash t$ form stores a guard DAG d instead of a single guard g .
- Matching a parallel composition (d_1 or d_2) means matching against d_1 ; if that succeeds, the overall match succeeds; if not, match against d_2 .
- Matching a sequential composition (d_1, d_2) means matching against d_1 ; if that succeeds, match against d_2 . If either match fails, the whole match fails.

Matching parallel composition (d_1 or d_2) is much like matching $\sqsubset_{t_2}^{t_1}$, and matching sequential composition (d_1, d_2) is much like matching $\dashv\vdash d \dashv\vdash t$.

A clearer, non-flat visualisation of the guard DAG of the first clause could be



This visualisation acknowledges that GrdDag really models labelled *series-parallel graphs* [Eppstein 1992], a very specific kind of DAG with a straightforward encoding as an algebraic data type: every guard g induces a series-parallel graph with a single edge from source to sink; conjunction (d_1, d_2) corresponds to series composition of graphs for d_1 and d_2 ; and disjunction (d_1 or d_2) corresponds to parallel composition of graphs for d_1 and d_2 .

Although the redefinition of coverage checking functions in Figure 10 is much more expansive in size than the original definition in Figure 5, we will see that the encoded logic is derivative.

There is a new function C that computes the *covered* set of a guard dag d . This function was previously inlined into the recursive call sites of \mathcal{U} and \mathcal{A} ; it computes the set of Θ reaching t in $\dashv\vdash g \dashv\vdash t$. It is no longer possible to inline it because the $\dashv\vdash d \dashv\vdash t$ form now carries a guard DAG d with nested structure; hence we need a separate function and a small refactor.

As expected, computing the uncovered set of parallel composition (d_1 or d_2) is much the same as for the $\sqcap_{t_2}^{t_1}$ form, and similarly for sequential composition (d_1, d_2) and the $\dashv\vdash d \dashv\vdash t$ form. Similarly, the uncovered set for the $\dashv\vdash n$ form is the same as that of the irrefutable guards $!x$ and let $x = e$. For the purposes of \mathcal{U} , we could have written a small function from Gdt to GrdDag to share this duplicate code. Our actual implementation in GHC (Section 5) simply re-uses polymorphic combinators.

The changes to \mathcal{A} are similar in nature. The use of \cup in $\mathcal{A}(\Theta, (g_1, g_2))$ may be unexpected, since usually sequential composition leads to conjunction \wedge , not disjunction \cup . Nevertheless, \cup is the correct choice, because it follows directly from the previous definition of $\mathcal{A}(\Theta, \dashv\vdash g_1, g_2 \dashv\vdash t)$ and how the resulting $\dashv\vdash \Theta_1 \dashv\vdash \dashv\vdash \Theta_2 \dashv\vdash u$ annotations are used in \mathcal{R} (Figure 6): u can be redundant only if there is no inhabitant in $\Theta_1 \cup \Theta_2$; otherwise it is inaccessible.

It is reassuring to know that extending coverage checking in Figure 5 with Or-patterns is derivative and compatible with all the other proposed extensions, although it takes a slight refactoring. On the other hand, years of maintaining LYG have shown that the entire complexity rests in the inhabitation test (Figure 8). We did not need to touch that; and neither did we need to adjust Figure 6 or later: this is compelling evidence that the core of our approach is quite extensible and robust.

5 IMPLEMENTATION

Our implementation of LYG has been part of GHC since the 8.10 release in 2020, including all extensions in Section 4 (except for strict-by-default source syntax). Our implementation accumulates quite a few tricks that go beyond the pure formalism. This section is dedicated to describing these.

Warning messages need to reference source syntax in order to be comprehensible by the user. At the same time, coverage checks involving GADTs need a type checked program, so the only reasonable design is to run the coverage checker between type checking and desugaring to GHC Core, a typed intermediate representation lacking the connection to source syntax. We perform coverage checking in the same tree traversal as desugaring.

5.1 Interleaving \mathcal{U} and \mathcal{A}

The set of reaching values is an argument to both \mathcal{U} and \mathcal{A} . Given a particular set of input values and a guard tree, one can see by a simple inductive argument that both \mathcal{U} and \mathcal{A} are always called at the same arguments! Hence for an implementation it makes sense to compute both results together, if only for not having to recompute the results of \mathcal{U} again in \mathcal{A} .

But there's more: Looking at the last clause of \mathcal{U} in Figure 5, we can see that we syntactically duplicate Θ every time we have a pattern guard. In the worst case, that can amount to exponential growth of the refinement predicate and for the time to prove it empty!

What we really want is to summarise a Θ into a more compact canonical form before doing these kinds of *splits*. But that's exactly what ∇ is! Therefore, in our implementation we don't pass around and annotate refinement types, but the result of calling \mathcal{N} on them directly.

You can see the resulting definition in Figure 11. The readability is clouded by unwrapping of pairs. \mathcal{UA} requires that each ∇ individually is non-empty, i.e. not \times . This invariant is maintained by adding φ constraints through \oplus_φ , which filters out any ∇ that would become empty.

$$\boxed{\bar{\nabla} \dot{\oplus}_\varphi \varphi = \bar{\nabla}}$$

$$\epsilon \dot{\oplus}_\varphi \varphi = \epsilon$$

$$(\nabla_1 \dots \nabla_n) \dot{\oplus}_\varphi \varphi = \begin{cases} ((\Gamma \parallel \Delta)) (\nabla_2 \dots \nabla_n \dot{\oplus}_\varphi \varphi) & \text{if } \langle \Gamma \parallel \Delta \rangle = \nabla \oplus_\varphi \varphi \\ (\nabla_2 \dots \nabla_n) \dot{\oplus}_\varphi \varphi & \text{otherwise} \end{cases}$$

$$\boxed{\mathcal{UA}(\bar{\nabla}, t) = (\bar{\nabla}, u)}$$

$$\begin{aligned}
\mathcal{UA}(\bar{\nabla}, \rightarrow n) &= (\epsilon, \rightarrow \bar{\nabla} n) \\
\mathcal{UA}(\bar{\nabla}, \sqsubset_{t_2}^{t_1}) &= (\bar{\nabla}_2, \sqsubset_{u_2}^{u_1}) \text{ where } (\bar{\nabla}_1, u_1) = \mathcal{UA}(\bar{\nabla}, t_1) \\
&\quad (\bar{\nabla}_2, u_2) = \mathcal{UA}(\bar{\nabla}_1, t_2) \\
\mathcal{UA}(\bar{\nabla}, \dashv!x \dashv t) &= \dashv \bar{\nabla} \dot{\oplus}_\varphi (x \approx \perp) \dashv \dashv u \\
&\quad \text{where } (\bar{\nabla}', u) = \mathcal{UA}(\bar{\nabla} \dot{\oplus}_\varphi (x \not\approx \perp), t) \\
\mathcal{UA}(\bar{\nabla}, \dashv! \text{let } x = e \dashv t) &= \mathcal{UA}(\bar{\nabla} \dot{\oplus}_\varphi (\text{let } x = e), t) \\
\mathcal{UA}(\bar{\nabla}, \dashv! K \bar{a} \bar{y} \bar{y} : \bar{\tau} \leftarrow x \dashv t) &= ((\bar{\nabla} \dot{\oplus}_\varphi (x \not\approx K)) \bar{\nabla}', u) \\
&\quad \text{where } (\bar{\nabla}', u) = \mathcal{UA}(\bar{\nabla} \dot{\oplus}_\varphi (K \bar{a} \bar{y} \bar{y} : \bar{\tau} \leftarrow x), t)
\end{aligned}$$

Fig. 11. Fast coverage checking

5.2 Throttling for Graceful Degradation

Even with the tweaks from Section 5.1, checking certain pattern matches remains NP-hard [Sekar et al. 1995]. Naturally, there will be cases where we have to conservatively approximate in order not to slow down compilation too much. Consider the following example and its corresponding guard tree:

data $T = A \mid B; f1, f2 :: \text{Int} \rightarrow T$

$$\begin{array}{l}
g - \\
| A \leftarrow f1\ 1, A \leftarrow f2\ 1 = () \\
| A \leftarrow f1\ 2, A \leftarrow f2\ 2 = () \\
\dots \\
| A \leftarrow f1\ N, A \leftarrow f2\ N = ()
\end{array}
\quad
\begin{array}{l}
\vdash \text{let } a_1 = f1\ 1, !a_1, A \leftarrow a_1, \text{let } b_1 = f2\ 1, !b_1, A \leftarrow b_1 \longrightarrow 1 \\
\vdash \text{let } a_2 = f1\ 2, !a_2, A \leftarrow a_2, \text{let } b_2 = f2\ 2, !b_2, A \leftarrow b_2 \longrightarrow 2 \\
\vdash \dots \longrightarrow \dots \\
\vdash \text{let } a_N = f1\ N, !a_N, A \leftarrow a_N, \text{let } b_N = f2\ N, !b_N, A \leftarrow b_N \longrightarrow N
\end{array}$$

Each of the N GRHS can fall through in two distinct ways: By failure of either pattern guard involving $f1$ or $f2$. Initially, we start out with a single input ∇ . After the first equation it will split into two sub- ∇ s, after the second into four, and so on. This exponential pattern repeats N times, and leads to horrible performance!

Instead of *refining* ∇ with the pattern guard, leading to a split, we could just continue with the original ∇ , thus forgetting about the $a_1 \not\approx A$ or $b_1 \not\approx A$ constraints. In terms of the modeled refinement type, ∇ is still a superset of both refinements, and thus a sound overapproximation.

In our implementation, we call this *throttling*: We limit the number of reaching ∇ s to a constant. Whenever a split would exceed this limit, we continue with the original reaching ∇ s, a conservative estimate, instead. Intuitively, throttling corresponds to *forgetting* what we matched on in that

particular subtree. Throttling is refreshingly easy to implement! Only the last clause of \mathcal{UA} , where splitting is performed, needs to change:

$$\mathcal{UA}(\bar{V}, \neg \mid K \bar{a} \bar{y} \bar{y} : \bar{\tau} \leftarrow x \text{ --- } t) = \left(\left[\bar{V} \hat{\oplus}_\varphi (x \neq K) \right]_{\bar{V}}, u \right) \\ \text{where } (\bar{V}', u) = \mathcal{UA}(\bar{V} \hat{\oplus}_\varphi (K \bar{a} \bar{y} \bar{y} : \bar{\tau} \leftarrow x), t)$$

where the new throttling operator $[-]_{\bar{V}}$ is defined simply as

$$\left[\bar{V} \right]_{\bar{V}'} = \begin{cases} \bar{V} & \text{if } |\{\bar{V}\}| \leq K \\ \bar{V}' & \text{otherwise} \end{cases}$$

with K being an arbitrary constant. We use 30 as an arbitrary limit in our implementation (dynamically configurable via a command-line flag) without noticing any false positives in terms of exhaustiveness warnings outside of the test suite.

It is worth noting that Or-patterns (Section 4.9) introduce a function C to compute the covered set of a guard DAG d , and its case for $(d_1$ or $d_2)$ splits the incoming Θ as well; hence we throttle there as well to ensure graceful degradation.

5.3 Maintaining Residual COMPLETE Sets

Our implementation tries hard to make the inhabitation test as efficient as possible. For example, we represent Δ s by a mapping from variables to their positive and negative constraints for easier indexing. But there are also asymptotical improvements. Consider the following function:

```
data T = A1 | ... | A1000      f A1    = 1
pattern P = ...                f A2    = 2
{-# COMPLETE A1, P #-}        ...
                                f A1000 = 1000
```

f is exhaustively defined. To see that we need to perform an inhabitation test for the match variable x after the last clause. The test will conclude that the builtin COMPLETE set was completely overlapped. But in order to conclude that, our algorithm tries to instantiate x (via \vdash INST) to each of its 1000 constructors and try to add a positive constructor constraint! What a waste of time, given that we could just look at the negative constraints on x *before* trying to instantiate x . But asymptotically it shouldn't matter much, since we're doing this only once at the end.

Except that is not true, because we also perform redundancy checking! At any point in f 's definition there might be a match on P , after which all remaining clauses would be redundant by the user-supplied COMPLETE set. Therefore, we have to perform the expensive inhabitation test *after every clause*, involving $O(n)$ instantiations each.

Clearly, we can be smarter about that! Indeed, we cache *residual* COMPLETE sets in our implementation: Starting from the full COMPLETE sets, we delete ConLikes from them whenever we add a new negative constructor constraint, maintaining the invariant that each of the sets is inhabited by at least one constructor. Note how we never need to check the same constructor twice (except after adding new type constraints), thus we have an amortised $O(n)$ instantiations for the whole checking process.

5.4 Reporting Uncovered Patterns

The expansion function \mathcal{E} in Figure 6 exists purely for presenting uncovered patterns to the user. It doesn't account for negative information, however, which can lead to surprising warnings. Consider a definition like $b \text{ True} = ()$. The computed uncovered set of b is the normalised refinement type

$\nabla_b = \langle x : \text{Bool} \parallel x \neq \perp, x \neq \text{True} \rangle$, which crucially contains no positive information on x ! As a result, $\mathcal{E}(\nabla_b) = _$ and only the very unhelpful wildcard pattern $_$ will be reported as uncovered.

Our implementation does better and shows that this is just a presentational matter. It splits ∇_b on all possible constructors of `Bool`, immediately rejecting the refinement $\nabla_b \oplus_\delta x \approx \text{True}$ due to $x \neq \text{True} \in \nabla_b$. What remains is the refinement $\nabla_b \oplus_\delta x \approx \text{False} = \langle x : \text{Bool} \parallel x \neq \perp, x \neq \text{True}, x \approx \text{False} \rangle$, which has the desired positive information for which \mathcal{E} will happily report `False` as the uncovered pattern.

Additionally, our implementation formats negative information on opaque data types such as `Int` and `Char`, since idiomatic use would match on literals rather than on GHC-specific data constructors. For example, coverage checking `f 0 = ()` will report something like this:

```
Missing equations for function 'f':
  f x = ... where 'x' is not one of {0}
```

5.5 Syntax-specific Guard Tree Types

Since we submitted our work to ICFP in 2020, we continued to improve and refactor the implementation of LYG in GHC. Many of the changes were incremental improvements and bug fixes that are not easy to present without a lot of context, but one particularly important innovation³ was the introduction of syntax-specific instances of guard trees, such as

```
type SrcInfo = String -- approximately; identifies the k in rhsk
data PmMatch p = PmMatch { pm_pats :: p, pm_grhss :: [PmGRHS p] }
data PmGRHS p = PmGRHS { pg_grds :: p, pg_rhs :: SrcInfo }
```

These types are in structural correspondence to the `match` and `grhs` constructs in Figure 2 from whence they desugar. Prior to coverage checking, type parameter p is instantiated to lists of guards `Grd` (resp. `GrdDag` after Or-patterns were introduced, Section 4.9), and coverage checking elaborates this list into `RedSets`, carrying Θ s encoding covered and diverging input values.

Of course, the meaning of `PmMatch` and `PmGRHS` is in terms of the desugaring into unrestricted guard trees `Gdt`, as before. However, with the new encoding it became much easier to extract covered sets for long-distance information (Section 4.1), because the `pm_grhss` field has the same number of elements as there are `grhs` in a `match` and simple `Data.List.zip` suffices to bring covered sets and `grhs` together.

6 EVALUATION

To put the new coverage checker to the test, we performed a survey of real-world Haskell code using the `head.hackage` repository⁴. `head.hackage` contains a sizable collection of libraries and minimal patches necessary to make them build with a development version of GHC. We identified those libraries which compiled without coverage warnings using GHC 8.8.3 (which uses GMTM as its checking algorithm) but emitted warnings when compiled using our LYG version of GHC.

Of the 361 libraries in `head.hackage`, seven of them revealed coverage issues that only LYG warned about. Two of the libraries, `pandoc` and `pandoc-types`, have cases that were flagged as redundant due to LYG's improved treatment of guards and term equalities.

One library, `geniplate-mirror`, has a case that was redundant by way of long-distance information. Another library, `generic-data`, has a case that is redundant due to bang patterns.

The last three libraries—`Cabal`, `HsYAML`, and `network`—were the most interesting. `HsYAML` in particular defines this function:

³<https://gitlab.haskell.org/ghc/ghc/-/commit/1207576ac0cfdd3fe1ea00b5505f7c874613451e>

⁴<https://gitlab.haskell.org/ghc/head.hackage/commit/30a310fd8033629e1cbb5a9696250b22db5f7045>

Table 1. The relative compile-time performance of GHC 8.8.3 (which implements GMTM) and HEAD (which implements LYG) on test cases designed to stress-test coverage checking.

	Time (milliseconds)			Megabytes allocated		
	8.8.3	HEAD	% change	8.8.3	HEAD	% change
T11276	1.16	1.69	45.7%	1.86	2.39	28.6%
T11303	28.1	18.0	-36.0%	60.2	39.9	-33.8%
T11303b	1.15	0.39	-65.8%	1.65	0.47	-71.8%
T11374	4.62	3.00	-35.0%	6.16	3.20	-48.1%
T11822	1,060	16.0	-98.5%	2,010	27.9	-98.6%
T11195	2,680	22.3	-99.2%	3,080	39.5	-98.7%
T17096	7,470	16.6	-99.8%	17,300	35.4	-99.8%
PmSeriesS	44.5	2.58	-94.2%	52.9	6.19	-88.3%
PmSeriesT	48.3	6.86	-85.8%	61.4	17.6	-71.4%
PmSeriesV	131	4.54	-96.5%	139	9.53	-93.2%

```
go' _ _ _ xs | False = error (show xs)
```

```
go' _ _ _ xs = err xs
```

The first clause is clearly unreachable, and LYG now flags it as such. However, the authors of HsYAML likely left in this clause because it is useful for debugging purposes. One can comment out the second clause and remove the False guard to quickly try out a code path that prints a more detailed error message. Moreover, leaving the first clause in the code ensures that it is typechecked and less susceptible to bitrotting over time.

In order to support this use case in HsYAML, a primitive definition `considerAccessible = False` was added in GHC 9.2, to be used instead of False above and signalling to GHC that the first clause should not get marked as redundant. The unreachable code in Cabal and network is of a similar caliber and could benefit from `considerAccessible` as well.

6.1 Performance Tests

To compare the efficiency of GMTM and LYG quantitatively, we collected a series of test cases from GHC's test suite that are designed to test the compile-time performance of coverage checking. Table 1 lists each of these 11 test cases. Test cases with a T prefix are taken from user-submitted bug reports about the poor performance of GMTM. Test cases with a PmSeries prefix are adapted from Maranget [2007], which presents several test cases that caused GHC to exhibit exponential running times during coverage checking.

We compiled each test case with GHC 8.8.3, which uses GMTM as its checking algorithm, and GHC HEAD (a pre-release of GHC 8.10), which uses LYG. We measured (1) the time spent in the desugarer, the phase of compilation in which coverage checking occurs, and (2) how many megabytes were allocated during desugaring. Table 1 shows these figures as well as the percent change going from 8.8.3 to HEAD. Most cases exhibit a noticeable improvement under LYG, with the exception of T11276. Investigating T11276 suggests that the performance of GHC's equality constraint solver has become more expensive in HEAD [GHC issue 2020c], and these extra costs outweigh the performance benefits of using LYG. This performance bug was fixed in GHC 9.0⁵.

Note that for typical code (rather than for regression tests), time spent doing coverage checking is dwarfed by the time the rest of the desugarer takes. A very desirable property for a static analysis that is irrelevant to the compilation process!

⁵<https://gitlab.haskell.org/ghc/ghc/-/commit/fd7ea0fee92a60f9658254cc4fe3abdb4ff299b1>

6.2 GHC Issues

Implementing LYG in GHC has fixed over 30 bug reports related to coverage checking. These include:

- Better compile-time performance [GHC issue 2015a, 2016e, 2019a,b]
- More accurate warnings for empty case expressions [GHC issue 2015b, 2017f, 2018e,g, 2019c]
- More accurate warnings due to LYG’s desugaring [GHC issue 2016c,d, 2017d, 2018a, 2020d]
- More accurate warnings due to improved term-level reasoning [GHC issue 2016a, 2017a, 2018b,c,d,h, 2019d,e,h]
- More accurate warnings due to tracking long-distance information [GHC issue 2019k, 2020a,b]
- Improved treatment of COMPLETE sets [GHC issue 2016b, 2017b,c,e,g, 2018j, 2019f,g,i]
- Better treatment of strictness, bang patterns, and newtypes [GHC issue 2018f,i, 2019j,l]

7 SOUNDNESS

The evaluation in Section 6 yields compelling evidence that LYG is *sound*. That is, in terms of the formalism, LYG *overapproximates*—but never underapproximates—the set of reaching values passed to \mathcal{U} and \mathcal{A} . As a result, LYG will never fail to report uncovered clauses (no false negatives), but it may report false positives. Similarly, LYG will never report accessible clauses as redundant (no false positives), but it may fail to report clauses which are redundant when the code involved is too close to “undecidable territory”.

Remarkably, the symbolic checking process involving \mathcal{U} , \mathcal{A} and \mathcal{R} does not overapproximate at all. To our knowledge, LYG overapproximates only in these three mechanisms:

- LYG can run out of fuel for inhabitation testing (Section 3.7).
- Throttling (Section 5.2) is useful when implementing LYG efficiently.
- LYG forgoes non-trivial semantic analysis of expressions. LYG can recognize identical patterns or subexpressions, but it stops short of anything more sophisticated, such as interprocedural analysis or SMT-style reasoning (Section 8.2.1).

But what does it actually *mean* for a value to match a particular part of a guard tree, such as a right-hand side $\rightarrow k$? In what precise sense does LYG— or does not — overapproximate this supposed *semantics*?

Since this work appeared at ICFP 2020, Dieterichs [2021] worked out both a formal semantics as well as a mechanised correctness proof in Lean 3 for the coverage checking pass from guard trees into uncovered set and annotated trees.⁶ He shows that \mathcal{U} , \mathcal{A} and \mathcal{R} preserve key semantic properties of the guard trees under analysis, provided that function $\mathcal{G}(\Theta)$ for generating inhabitants indeed overapproximates Θ . We will briefly summarise the correctness results here. For that, we need to define a plausible formal semantics for guard trees and refinement predicates.

7.1 Semantics

We have described the semantics of guard trees and guards informally in Section 3.1. Figure 12 formalises this intuition, describing the semantics of guard trees by a function $\text{Gdt}[\![t]\!]_{\rho}$ that, given a guard tree t and an environment ρ describing a vector of values to match against, returns

- success(k) when ρ is a vector of values that will reach RHS k when matched against t .
- fail when ρ is a vector of values that is not covered by t .
- diverge when ρ is a vector of values that will lead to divergence when matched against t .

Likewise, the valuation $\text{Grd}[\![g]\!]_{\rho}$ returns success(ρ') when the vector of values ρ matches guard g , extending ρ with new bindings into ρ' . The semantics of expressions $\text{Expr}[\![e]\!]_{\rho}$ maps into the

⁶Types and type constraints are ignored; their interaction is largely a black box to our approach anyway.

Semantics of guard trees

$$\begin{aligned}
 d \in \mathbb{D} &= \perp \mid K \bar{d} \mid \dots \\
 \rho \in \text{Env} &::= [x \mapsto \bar{d}] \\
 r \in \text{Res}[\square] &::= \text{success}(\square) \mid \text{fail} \mid \text{diverge}
 \end{aligned}$$

$\text{Expr}[[e]]_\rho \in \mathbb{D}$,	$\text{Grd}[[g]]_\rho \in \text{Res}[\rho]$,	$\text{Gdt}[[t]]_\rho \in \text{Res}[k]$
$\text{Expr}[[K \bar{e}]]_\rho$	$=$	$K \overline{\text{Expr}[[e]]_\rho}$
$\text{Expr}[[e]]_\rho$	$=$	\dots
$\text{Grd}[[\text{let } x = e]]_\rho$	$=$	$\rho[x \mapsto \text{Expr}[[e]]_\rho]$
$\text{Grd}[[K \bar{y} \leftarrow x]]_\rho$	$=$	$\begin{cases} \text{success}(\rho[y \mapsto \bar{d}]) & \text{if } \rho(x) = K \bar{d} \\ \text{fail} & \text{otherwise} \end{cases}$
$\text{Grd}[[!x]]_\rho$	$=$	$\begin{cases} \text{diverge} & \text{if } \rho(x) = \perp \\ \text{success}(\rho) & \text{otherwise} \end{cases}$
$\text{Gdt}[[\rightarrow k]]_\rho$	$=$	$\text{success}(k)$
$\text{Gdt}[[\sqcap_{t_2}^{t_1}]]_\rho$	$=$	$\begin{cases} \text{Gdt}[[t_2]]_\rho & \text{if } \text{Gdt}[[t_1]]_\rho = \text{fail} \\ \text{Gdt}[[t_1]]_\rho & \text{otherwise} \end{cases}$
$\text{Gdt}[[\dashv g \dashv t]]_\rho$	$=$	$\begin{cases} \text{Gdt}[[t]]_{\rho'} & \text{if } \text{Grd}[[g]]_\rho = \text{success}(\rho') \\ \text{Grd}[[g]]_\rho & \text{otherwise} \end{cases}$

Semantics of refinement types

$$\boxed{\rho \models (\varphi, \rho), \quad \rho \models \Theta}$$

$$\begin{array}{c}
 \frac{}{\rho \models (\checkmark, \rho)} \quad \frac{\rho(x) = K \bar{d}}{\rho \models (K \bar{y} \leftarrow x, \rho[y \mapsto \bar{d}])} \quad \frac{\rho(x) \neq K \bar{d}}{\rho \models (x \not\approx K, \rho)} \\
 \\
 \frac{\rho(x) = \perp}{\rho \models (x \approx \perp, \rho)} \quad \frac{\rho(x) \neq \perp}{\rho \models (x \not\approx \perp, \rho)} \quad \frac{}{\rho \models (\text{let } x = e, \rho[x \mapsto \text{Expr}[[e]]_\rho])} \\
 \\
 \frac{\Gamma_1 \vdash \rho_1 \quad \Gamma_2 \vdash \rho_2 \quad \rho_1 \models (\varphi, \rho_2) \quad \rho_2 \models \langle \Gamma \mid \Phi \rangle}{\rho_1 \models \langle \Gamma_1 \mid \varphi \wedge \Phi \rangle} \quad \frac{\rho \models \langle \Gamma \mid \Phi_1 \rangle}{\rho \models \langle \Gamma \mid \Phi_1 \vee \Phi_2 \rangle} \quad \frac{\rho \models \langle \Gamma \mid \Phi_2 \rangle}{\rho \models \langle \Gamma \mid \Phi_1 \vee \Phi_2 \rangle}
 \end{array}$$

Fig. 12. Semantics of guard trees

semantic domain \mathbb{D} , just as the environment ρ . Since this work leaves open a lot of details about the expression fragment of the source language, the semantics leaves open \mathbb{D} and most of $\text{Expr}[[e]]_\rho$ as well, with the exception of postulating a semantics for the data constructor application case.

Refinement types have been introduced by informal examples in [Section 3.2](#), denoting refinement types Θ by sets of vectors of values ρ that satisfy the encoded refinement predicate. [Figure 12](#) finally defines the satisfiability relation by an inductive predicate $\rho \models \Theta$. Thus, whenever a vector of values ρ is part of the set denoted by a refinement type Θ , the inductive predicate must be provable.

The definition of $\rho \models \Theta$ assumes that conjunction \wedge is associated to the right, $\varphi \wedge \Phi$, highlighting an unusual scoping semantics that was previously implicit: Any binding constructs in the φ to

the left of \wedge , such as let $x = e$ or $K \bar{y} \leftarrow x$, introduce names that are subsequently in scope in the Φ to the right of \wedge . In hindsight, we could have picked a different operator symbol to avoid this confusion, for example $(\varphi \text{ in } \Phi)$, such as in [Dieterichs \[2021\]](#). Doing so would however complicate the $(\Theta \wedge \varphi)$ operator a bit. In the absence of types, the postulated judgment $\Gamma \vdash \rho$ merely becomes a scoping check, namely that Γ has the same domain as ρ .

7.2 Formal Soundness Statement

Having stated plausible semantics for the inputs and outputs of \mathcal{U} , we can formulate what it means for \mathcal{U} to be correct, following [Dieterichs \[2021, Section 4.1\]](#) who mechanised the proof in Lean 3.

THEOREM 7.1. *Let $\mathcal{U}(\langle \Gamma \mid \checkmark \rangle, t) = \Theta$. Then $\text{Gdt}\llbracket t \rrbracket_\rho = \text{fail}$ if and only if $\rho \models \Theta$.*

In other words: when Θ is the set of uncovered values of guard tree t as computed by \mathcal{U} , then any vector of values ρ that falls through all clauses of t (i.e., $\text{Gdt}\llbracket t \rrbracket_\rho = \text{fail}$) is in Θ (i.e. $\rho \models \Theta$). In this precise sense, \mathcal{U} is *sound*. Conversely, when \mathcal{U} returns a non-empty refinement type Θ , there exists a vector of values ρ in Θ , and by [Theorem 7.1](#) we have that ρ must also fall through all clauses of t . In this precise sense, \mathcal{U} is *complete*.

Of course, the judgment $\rho \models \Theta$ frequently compares domain values d that ultimately come from evaluating expressions $\text{Expr}\llbracket e \rrbracket_\rho$, rendering the predicate undecidable for many source languages. Thus, any implementation of \mathcal{G} will be sound wrt. $\rho \models \Theta$, but not complete — it will *overapproximate* Θ . [Dieterichs \[2021\]](#) captures this in his `can_prove_empty` definition to abstract over sound implementations of \mathcal{G} . [Dieterichs \[2021, Section 4.2\]](#) proves the following soundness theorem about \mathcal{A} and \mathcal{R} .

THEOREM 7.2. *Let $\mathcal{R}(\mathcal{A}(\langle \Gamma \mid \checkmark \rangle, t)) = (a, i, r)$ and \mathcal{G} sound in the above sense.*

- *If $\text{Gdt}\llbracket t \rrbracket_\rho = \text{success}(k)$, then $k \in a$, i.e., clause k is accessible according to \mathcal{A} and \mathcal{R} .*
- *If $k \in r$ is redundant, then removing clause k from guard tree t does not change the semantics of t , i.e., $\forall \rho. \text{Gdt}\llbracket t \rrbracket_\rho = \text{Gdt}\llbracket \text{remove}(k, t) \rrbracket_\rho$ (where $\text{remove}(k, t)$ is the implied removal operation).*

Perhaps unsurprisingly, proving correct the transformation in the second part of [Theorem 7.2](#) proved far more subtle than the proof for [Theorem 7.1](#). Fortunately, mechanisation provides confidence in its correctness.

8 RELATED WORK

8.1 Comparison with GADTs Meet Their Match

[Karachalias et al. \[2015\]](#) present GADTs Meet Their Match (GMTM), an algorithm which handles many of the subtleties of GADTs, guards, and laziness mentioned in [Section 2](#). Despite this, the GMTM algorithm still gives incorrect warnings in many cases.

8.1.1 GMTM Does Not Consider Laziness in its Full Glory. The formalism in [Karachalias et al. \[2015\]](#) incorporates strictness constraints, but these constraints can only arise from matching against data constructors. GMTM does not consider strict matches that arise from strict fields of data constructors or bang patterns. A consequence of this is that GMTM would incorrectly warn that v ([Section 2.3](#)) is missing a case for `SJust`, even though such a case is unreachable. LYG, on the other hand, more thoroughly tracks strictness when desugaring Haskell programs.

8.1.2 GMTM's Treatment of Guards Is Shallow. GMTM can only reason about guards through an abstract term oracle. Although the algorithm is parametric over the choice of oracle, in practice the implementation of GMTM in GHC uses an extremely simple oracle that can only reason about guards in a limited fashion. More sophisticated uses of guards, such as in this variation of the `safeLast` function from [Section 2.2.1](#), will cause GMTM to emit erroneous warnings:

safeLast2 xs

```
| (x : _) ← reverse xs = Just x
| []      ← reverse xs = Nothing
```

While GMTM’s term oracle is customisable, it is not as simple to customize as one might hope. The formalism in [Karachalias et al. \[2015\]](#) represents all guards as $p \leftarrow e$, where p is a pattern and e is an expression. This is a straightforward, syntactic representation, but it also makes it more difficult to analyse when e is a complicated expression. This is one of the reasons why it is difficult for GMTM to accurately give warnings for the *safeLast* function, since it would require recognizing that both clauses scrutinise the same expression in their view patterns.

LYG makes analysing term equalities simpler by first desugaring guards from the surface syntax to guard trees. The \oplus_φ function, which is roughly a counterpart to GMTM’s term oracle, can then reason about terms arising from patterns. While \oplus_φ is already more powerful than a trivial term oracle, its real strength lies in the fact that it can easily be extended, as LYG’s treatment of view patterns ([Section 4.3](#)) demonstrates. While GMTM’s term oracle could be improved to accomplish the same thing, it is unlikely to be as straightforward of a process as extending \oplus_φ .

8.2 Comparison with Similar Coverage Checkers

8.2.1 Structural and Semantic Pattern Matching Analysis in Haskell. [Kalvoda and Kerckhove \[2019\]](#) implement a variation of GMTM that leverages an SMT solver to give more accurate coverage warnings for programs that use guards. For instance, their implementation can conclude that the *signum* function from [Section 2.1](#) is exhaustive. This is something that LYG cannot do out of the box, although it would be possible to extend \oplus_φ with SMT-like reasoning about booleans and linear integer arithmetic.

8.2.2 Warnings for Pattern Matching. [Maranget \[2007\]](#) presents a coverage checking algorithm for OCaml that can identify clauses that are not *useful*, i.e. *useless*. While OCaml is a strict language, the algorithm can be adapted to handle languages with non-strict semantics such as Haskell. In a lazy setting, uselessness corresponds to our notion of unreachable clauses. [Maranget](#) does not distinguish inaccessible clauses from redundant ones; thus clauses flagged as useless (such as the first two clauses of u' in [Section 2.3.1](#)) generally can’t be deleted without changing (lazy) program semantics.

8.2.3 Case Trees in Dependently Typed Languages. *Case trees* [[Augustsson 1985](#)] are a standard way of compiling pattern matches to efficient code. Much like LYG’s guard trees, case trees present a simplified representation of pattern matching. Several compilers for dependently typed languages also use case trees as coverage checking algorithms, as a well typed case tree can guarantee that it covers all possible cases. Case trees play an integral role in coverage checking in Agda [[Cockx and Abel 2018](#); [Norell 2007](#)] and the Equations plugin for Coq [[Sozeau 2010](#); [Sozeau and Mangin 2019](#)]. [Oury \[2007\]](#) checks for coverage in a dependently typed setting using sets of inhabitants of data types, which have similarities to case trees.

One could take inspiration from case trees should one wish to extend LYG to support dependent types. Our implementation of LYG in GHC can already handle quasi-dependently typed code, such as the *singletons* library [[Eisenberg and Stolarek 2014](#); [Eisenberg and Weirich 2012](#)], so we expect that it can be adapted to full dependent types. One key change that would be required is extending equation (9) in [Figure 7](#) to reason about term constraints in addition to type constraints. GHC’s constraint solver already has limited support for term-level reasoning as part of its *DataKinds* language extension [[Yorgey et al. 2012](#)], so the groundwork is present.

8.2.4 Refinement Type–Based Totality Checking in Liquid Haskell. In addition to LYG, Liquid Haskell uses refinement types to perform a limited form of exhaustivity checking [Vazou et al. 2014, 2017]. While exhaustiveness checks are optional in ordinary Haskell, they are mandatory for Liquid Haskell, as proofs written in Liquid Haskell require user-defined functions to be total (and therefore exhaustive) in order to be sound. For example, consider this non-exhaustive function:

```
fibPartial :: Integer → Integer
fibPartial 0 = 0
fibPartial 1 = 1
```

When compiled, GHC fills out this definition by adding an extra `fibPartial _ = error "undefined"` clause. Liquid Haskell leverages this by giving `error` the refinement type:

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

As a result, attempting to use `fibPartial` in a proof will fail to verify unless the user can prove that `fibPartial` is only ever invoked with the arguments 0 or 1.

8.3 Other Representations of Constraints

8.3.1 Leveraging Existing Constraint Solvers. LYG represents Φ constraints using logical predicates that are tailor-made for LYG’s purposes. One could instead imagine encoding Φ constraints in a more standard logic and then using an “off-the-shelf” constraint solver to check them. This would render Figure 7 and the arguably rather intricate Sections 3.6 and 3.7 unnecessary, and it allows the checker to benefit from improvements to the solver without any further maintenance burden.

Encoding Φ constraints into another logic would have its downsides, however. The \oplus_{φ} function is able to reason about LYG-oriented predicates rather efficiently, but other constraint solvers (e.g., STM solvers) might incur significant constant factors. Moreover, elaborating from one logic to another could inhibit programmers from forming a mental model of how coverage checking works.

8.3.2 Refinement Types versus Predicates. Refinement types Θ and predicates Φ are very similar. The main difference between the two is that refinement types carry a typing context Γ that is used for inhabitation testing. Predicates are quite fully featured on their own, however, as they can bind variables that scope over conjunctions. The scoping semantics of predicates allows \mathcal{U} and \mathcal{A} to be purely syntactic transformations, and in fact, they could be modified to take Φ as an argument rather than Θ .

Making \mathcal{U} and \mathcal{A} operate over Θ or Φ is ultimately a design choice. We have opted to operate over Θ mainly because we find it more intuitive to think about coverage checking as refining a vector of values as it falls from one match to the next. In our opinion, that intuition is more easily expressed with refinement types than predicates alone.

8.4 Positive and Negative Information

LYG’s use of positive and negative constructor constraints is inspired by Sestoft [1996], which uses positive and negative information to implement a pattern-match compiler for ML. Sestoft utilises positive and negative information to generate decision trees that avoid scrutinizing the same terms repeatedly. This insight is equally applicable to coverage checking and is one of the primary reasons for LYG’s efficiency.

Besides efficiency, the accuracy of redundancy warnings involving COMPLETE sets hinge on negative constraints. To see why this isn’t possible in other checkers that only track positive information, such as those of Karachalias et al. [2015] (Section 8.1) and Marangé [2007] (Section 8.2.2), consider the following example:

<pre> pattern True' = True {-# COMPLETE True', False #-} </pre>	<pre> <i>f</i> False = 1 <i>f</i> True' = 2 <i>f</i> True = 3 </pre>
--	---

GMTM would have to commit to a particular COMPLETE set when encountering the match on False, without any semantic considerations. Choosing {True', False} here will mark the third GRHS as redundant, while choosing {True, False} won't. GHC's implementation used to try each COMPLETE set in turn and would disambiguate using a complicated metric based on the number and kinds of warnings the choice of each set would generate [GHC team 2020], which was broken still [GHC issue 2017g].

Negative constraints make LYG efficient in other places too, such as in this example:

<pre> data T = A1 ... A1000 </pre>	<pre> <i>h</i> A1 _ = 1 <i>h</i> _ A1 = 2 </pre>
---	--

In *h*, GMTM would split the value vector (which is like LYG's Δ s without negative constructor constraints) into 1000 alternatives over the first match variable, and then *each* of the 999 value vectors reaching the second GRHS into another 1000 alternatives over the second match variable. Negative constraints allow LYG to compress the 999 value vectors falling through into a single one indicating that the match variable can no longer be A1.

8.5 Strict Fields in Inhabitation Testing

The `Inst` function in Figure 8 takes strict fields into account during inhabitation testing, which is essential to conclude that the v function from Section 2.3 is exhaustive. This trick was pioneered by Oury [2007], who uses it to check for unreachable cases in the presence of dependent types. Coverage checkers for strict and total programming languages usually implement inhabitation testing, but sometimes with less-than-perfect results. As two data points, we decided to see how OCaml and Idris, two call-by-value languages that check for pattern-match coverage⁷, would fare when checking functions like v :

<pre> (*OCaml*) type void = ;; let <i>v</i> (None : void option) : int = 0;; let <i>v'</i> (o : void option) : int = <i>match</i> o <i>with</i> None → 0 Some _ → 1;; </pre>	<pre> -- Idris <i>v</i> : Maybe Void → Int <i>v</i> Nothing = 0 <i>v'</i> : Maybe Void → Int <i>v'</i> Nothing = 0 <i>v'</i> (Just _) = 1 </pre>
---	--

Both OCaml 4.10.0 and Idris 1.3.2 correctly mark their respective versions of v as exhaustive. OCaml also correctly warns that the `Some` case in v' is unreachable, while Idris emits no warnings for v' at all.

Section 3.7 also contains an example of a function f that LYG will fail to recognize as exhaustive due to LYG's conservative, fuel-based approach to inhabitation testing. Porting f to OCaml and Idris reveals that both languages will also conservatively claim that f is non-exhaustive:

⁷Idris has separate compile-time and runtime semantics, the latter of which is call by value.

<pre>(*OCaml*) type t = MkT of t;; let f (None : t option) : int = 0;;</pre>	<pre>-- Idris data T : Type where MkT : T → T f : Maybe T → Int f Nothing = 0</pre>
--	---

Indeed, the warning that OCaml produces will cite `Some (MkT (MkT (MkT (MkT (MkT _)))))` as a case that is not matched, which suggests that OCaml may also be using a fuel-based approach. We believe these examples show that inhabitation testing is something that programming language implementors have discovered independently, but with varying degrees of success in putting into practice. We hope that LYG can bring this knowledge into wider use.

9 CONCLUSION

In this paper, we describe Lower Your Guards, a coverage checking algorithm that distills rich pattern matching into simple guard trees. Guard trees are amenable to analyses that are not easily expressible in coverage checkers that work over structural pattern matches. The last 4 years of continued maintenance of GHC’s implementation offer a compelling retrospective: the approach scales well to new language features, causes very few functional bug reports in practice, and offers robust performance.

ACKNOWLEDGMENTS

We would like to thank the anonymous ICFP reviewers for their feedback, as well as Henning Dieterichs, Martin Hecker, Sylvain Henry, Philipp Krüger, Luc Maranget and Sebastian Ullrich.

REFERENCES

- Lennart Augustsson. 1985. Compiling pattern matching. In *Functional Programming Languages and Computer Architecture*, Jean-Pierre Jouannaud (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 368–381.
- Jesper Cockx and Andreas Abel. 2018. Elaborating Dependent (Co)Pattern Matching. *Proc. ACM Program. Lang.* 2, ICFP, Article 75 (July 2018), 30 pages. <https://doi.org/10.1145/3236770>
- Henning Dieterichs. 2021. Formal Verification of Pattern Matching Analyses. Supervised by Sebastian Graf.
- Joshua Dunfield. 2007. *A Unified System of Type Refinements*. Ph.D. Dissertation. Carnegie Mellon University. CMU-CS-07-129.
- Richard A. Eisenberg and Jan Stolarek. 2014. Promoting Functions to Type Families in Haskell. In *Proceedings of the 2014 ACM SIGPLAN Symposium on Haskell (Gothenburg, Sweden) (Haskell ’14)*. Association for Computing Machinery, New York, NY, USA, 95–106. <https://doi.org/10.1145/2633357.2633361>
- Richard A. Eisenberg and Stephanie Weirich. 2012. Dependently Typed Programming with Singletons. In *Proceedings of the 2012 Haskell Symposium (Copenhagen, Denmark) (Haskell ’12)*. ACM, New York, NY, USA, 117–130. <https://doi.org/10.1145/2364506.2364522>
- David Eppstein. 1992. Parallel recognition of series-parallel graphs. *Information and Computation* 98, 1 (1992), 41–55. [https://doi.org/10.1016/0890-5401\(92\)90041-D](https://doi.org/10.1016/0890-5401(92)90041-D)
- Jacques Garrigue and Jacques Le Normand. 2011. Adding GADTs to OCaml: the direct approach. In *Workshop on ML*.
- GHC issue. 2015a. New pattern-match check can be non-performant. <https://gitlab.haskell.org/ghc/ghc/issues/11195>
- GHC issue. 2015b. No non-exhaustive pattern match warning given for empty case analysis. <https://gitlab.haskell.org/ghc/ghc/issues/10746>
- GHC issue. 2016a. In a record-update construct:ghc-stage2: panic! (the ‘impossible’ happened). <https://gitlab.haskell.org/ghc/ghc/issues/12957>
- GHC issue. 2016b. Inaccessible RHS warning is confusing for users. <https://gitlab.haskell.org/ghc/ghc/issues/13021>
- GHC issue. 2016c. Pattern coverage checker ignores dictionary arguments. <https://gitlab.haskell.org/ghc/ghc/issues/12949>
- GHC issue. 2016d. Pattern match incompleteness / inaccessibility discrepancy. <https://gitlab.haskell.org/ghc/ghc/issues/11984>
- GHC issue. 2016e. Representation of value set abstractions as trees causes performance issues. <https://gitlab.haskell.org/ghc/ghc/issues/11528>

- GHC issue. 2017a. -Woverlapping-patterns warns on wrong patterns for Int. <https://gitlab.haskell.org/ghc/ghc/issues/14546>
- GHC issue. 2017b. COMPLETE sets don't work at all with data family instances. <https://gitlab.haskell.org/ghc/ghc/issues/14059>
- GHC issue. 2017c. COMPLETE sets nerf redundant pattern-match warnings. <https://gitlab.haskell.org/ghc/ghc/issues/13965>
- GHC issue. 2017d. Incorrect pattern match warning on nested GADTs. <https://gitlab.haskell.org/ghc/ghc/issues/14098>
- GHC issue. 2017e. Pattern match checker mistakenly concludes pattern match on pattern synonym is unreachable. <https://gitlab.haskell.org/ghc/ghc/issues/14253>
- GHC issue. 2017f. Pattern synonym exhaustiveness checks don't play well with EmptyCase. <https://gitlab.haskell.org/ghc/ghc/issues/13717>
- GHC issue. 2017g. Wildcard patterns and COMPLETE sets can lead to misleading redundant pattern-match warnings. <https://gitlab.haskell.org/ghc/ghc/issues/13363>
- GHC issue. 2018a. -Wincomplete-patterns gets confused when combining GADTs and pattern guards. <https://gitlab.haskell.org/ghc/ghc/issues/15385>
- GHC issue. 2018b. Bogus -Woverlapping-patterns warning with OverloadedStrings. <https://gitlab.haskell.org/ghc/ghc/issues/15713>
- GHC issue. 2018c. Compiling a function with a lot of alternatives bottlenecks on insertIntHeap. <https://gitlab.haskell.org/ghc/ghc/issues/14667>
- GHC issue. 2018d. Completeness of View Patterns With a Complete Set of Output Patterns. <https://gitlab.haskell.org/ghc/ghc/issues/15884>
- GHC issue. 2018e. EmptyCase thinks pattern match involving type family is not exhaustive, when it actually is. <https://gitlab.haskell.org/ghc/ghc/issues/14813>
- GHC issue. 2018f. Erroneous "non-exhaustive pattern match" using nested GADT with strictness annotation. <https://gitlab.haskell.org/ghc/ghc/issues/15305>
- GHC issue. 2018g. Inconsistency w.r.t. coverage checking warnings for EmptyCase under unsatisfiable constraints. <https://gitlab.haskell.org/ghc/ghc/issues/15450>
- GHC issue. 2018h. Inconsistent pattern-match warnings when using guards versus case expressions. <https://gitlab.haskell.org/ghc/ghc/issues/15753>
- GHC issue. 2018i. nonVoid is too conservative w.r.t. strict argument types. <https://gitlab.haskell.org/ghc/ghc/issues/15584>
- GHC issue. 2018j. "Pattern match has inaccessible right hand side" with TypeRep. <https://gitlab.haskell.org/ghc/ghc/issues/14851>
- GHC issue. 2019a. 67-pattern COMPLETE pragma overwhelms the pattern match checker. <https://gitlab.haskell.org/ghc/ghc/issues/17096>
- GHC issue. 2019b. Add Luke Maranget's series in "Warnings for Pattern Matching". <https://gitlab.haskell.org/ghc/ghc/issues/17264>
- GHC issue. 2019c. `case (x :: Void) of _ -> ()` should be flagged as redundant. <https://gitlab.haskell.org/ghc/ghc/issues/17376>
- GHC issue. 2019d. GHC thinks pattern match is exhaustive. <https://gitlab.haskell.org/ghc/ghc/issues/16289>
- GHC issue. 2019e. Incorrect non-exhaustive pattern warning with PatternSynonyms. <https://gitlab.haskell.org/ghc/ghc/issues/16129>
- GHC issue. 2019f. Minimality of missing pattern set depends on constructor declaration order. <https://gitlab.haskell.org/ghc/ghc/issues/17386>
- GHC issue. 2019g. Panic during tyConAppArgs. <https://gitlab.haskell.org/ghc/ghc/issues/17112>
- GHC issue. 2019h. Pattern-match checker: True /= False. <https://gitlab.haskell.org/ghc/ghc/issues/17251>
- GHC issue. 2019i. Pattern match checking open unions. <https://gitlab.haskell.org/ghc/ghc/issues/17149>
- GHC issue. 2019j. Pattern match overlap checking doesn't consider -XBangPatterns. <https://gitlab.haskell.org/ghc/ghc/issues/17234>
- GHC issue. 2019k. Pattern match warnings are per Match, not per GRHS. <https://gitlab.haskell.org/ghc/ghc/issues/17465>
- GHC issue. 2019l. PmCheck treats Newtype patterns the same as constructors. <https://gitlab.haskell.org/ghc/ghc/issues/17248>
- GHC issue. 2019m. Strictness of pattern synonym matches and pattern-match checking. <https://gitlab.haskell.org/ghc/ghc/issues/17357>
- GHC issue. 2020a. -Wincomplete-record-updates ignores context. <https://gitlab.haskell.org/ghc/ghc/issues/17783>
- GHC issue. 2020b. Pattern match checker stumbles over reasonably tricky pattern-match. <https://gitlab.haskell.org/ghc/ghc/issues/17703>
- GHC issue. 2020c. Pattern match coverage checker allocates twice as much for trivial program with instance constraint vs. without. <https://gitlab.haskell.org/ghc/ghc/issues/17891>
- GHC issue. 2020d. Pattern match warning emitted twice. <https://gitlab.haskell.org/ghc/ghc/issues/17646>

- GHC team. 2020. COMPLETE pragmas. https://downloads.haskell.org/~ghc/8.8.3/docs/html/users_guide/glasgow_exts.html#pragma-COMPLETE
- Pavel Kalvoda and Tom Sydney Kerckhove. 2019. Structural and semantic pattern matching analysis in Haskell. arXiv:1909.04160 [cs.PL]
- Georgios Karachalias, Tom Schrijvers, Dimitrios Vytiniotis, and Simon Peyton Jones. 2015. *GADTs meet their match (extended version)*. Technical Report. KU Leuven. <https://people.cs.kuleuven.be/~tom.schrijvers/Research/papers/icfp2015.pdf>
- Luc Maranget. 2007. Warnings for pattern matching. *Journal of Functional Programming* 17 (2007), 387–421. Issue 3.
- Simon Marlow et al. 2010. Haskell 2010 Language Report. <https://web.archive.org/web/20100710201151/https://www.haskell.org/onlinereport/haskell2010/haskellch3.html#x8-230003.1>. Accessed: 2010-07-10.
- Ulf Norell. 2007. *Towards a practical programming language based on dependent type theory*. Ph.D. Dissertation. Department of Computer Science and Engineering, Chalmers University of Technology, SE-412 96 Göteborg, Sweden.
- Nicolas Oury. 2007. Pattern Matching Coverage Checking with Dependent Types Using Set Approximations. In *Proceedings of the 2007 Workshop on Programming Languages Meets Program Verification (Freiburg, Germany) (PLPV '07)*. Association for Computing Machinery, New York, NY, USA, 47–56. <https://doi.org/10.1145/1292597.1292606>
- Simon Peyton Jones and Sebastian Graf. 2023. Triemaps that match. arXiv:2302.08775 [cs.PL]
- Matthew Pickering, Gergő Erdi, Simon Peyton Jones, and Richard A. Eisenberg. 2016. Pattern Synonyms. In *Proceedings of the 9th International Symposium on Haskell (Nara, Japan) (Haskell 2016)*. Association for Computing Machinery, New York, NY, USA, 80–91. <https://doi.org/10.1145/2976002.2976013>
- John Rushby, Sam Owre, and Natarajan Shankar. 1998. Subtypes for specifications: Predicate subtyping in PVS. *IEEE Transactions on Software Engineering* 24, 9 (1998), 709–720.
- R. C. Sekar, R. Ramesh, and I. V. Ramakrishnan. 1995. Adaptive Pattern Matching. *SIAM J. Comput.* 24, 6 (Dec. 1995), 1207–1234. <https://doi.org/10.1137/S0097539793246252>
- Peter Sestoft. 1996. ML pattern match compilation and partial evaluation. In *Partial Evaluation*. Springer, 446–464.
- Matthieu Sozeau. 2010. Equations: A Dependent Pattern-Matching Compiler. In *Interactive Theorem Proving*, Matt Kaufmann and Lawrence C. Paulson (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 419–434.
- Matthieu Sozeau and Cyprien Mangin. 2019. Equations Reloaded: High-Level Dependently-Typed Functional Programming and Proving in Coq. *Proc. ACM Program. Lang.* 3, ICFP, Article 86 (July 2019), 29 pages. <https://doi.org/10.1145/3341690>
- 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 (Gothenburg, Sweden) (ICFP '14)*. ACM, 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. 2017. Refinement Reflection: Complete Verification with SMT. *Proc. ACM Program. Lang.* 2, POPL, Article 53 (Dec. 2017), 31 pages. <https://doi.org/10.1145/3158141>
- Dimitrios Vytiniotis, Simon Peyton Jones, Tom Schrijvers, and Martin Sulzmann. 2011. OutsideIn(X): Modular Type Inference with Local Assumptions. *J. Funct. Program.* 21, 4-5 (Sept. 2011), 333–412. <https://doi.org/10.1017/S0956796811000098>
- Hongwei Xi. 1998a. Dead Code Elimination Through Dependent Types. In *Proceedings of the First International Workshop on Practical Aspects of Declarative Languages (PADL '99)*. Springer-Verlag, London, UK, 228–242.
- Hongwei Xi. 1998b. *Dependent Types in Practical Programming*. Ph.D. Dissertation. Carnegie Mellon University.
- Hongwei Xi. 2003. Dependently typed pattern matching. *Journal of Universal Computer Science* 9 (2003), 851–872.
- Hongwei Xi, Chiyen Chen, and Gang Chen. 2003. Guarded Recursive Datatype Constructors. In *Proceedings of the 30th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (New Orleans, Louisiana, USA) (POPL '03)*. ACM, New York, NY, USA, 224–235. <https://doi.org/10.1145/604131.604150>
- 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 (Montreal, Quebec, Canada) (PLDI '98)*. Association for Computing Machinery, New York, NY, USA, 249–257. <https://doi.org/10.1145/277650.277732>
- Brent A. Yorgey, Stephanie Weirich, Julien Cretin, Simon Peyton Jones, Dimitrios Vytiniotis, and José Pedro Magalhães. 2012. Giving Haskell a Promotion. In *Proceedings of the 8th ACM SIGPLAN Workshop on Types in Language Design and Implementation (Philadelphia, Pennsylvania, USA) (TLDI '12)*. ACM, New York, NY, USA, 53–66. <https://doi.org/10.1145/2103786.2103795>