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), 33 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:

$f :: \mathit{Int} \longrightarrow \mathit{Bool}$
$f\ 0 = \mathit{True}$
$f\ 0 = \mathit{False}$

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 $f\ 0 = \mathit{False}$ clause

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 all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.
2475-1421/2020/8-ART107
https://doi.org/10.1145/3408989

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.

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 7 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 7.1). 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).

We discuss the wealth of related work in Section 7.

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 7) 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'` $\leftarrow c1$), 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 → Bool
not b | False ← b = True
      | True ← b  = False
```

```
not2 :: Bool → Bool
not2 False = True
not2 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 ($f \to p$), 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 *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
pattern Nil ← (Text.null → True)
pattern Cons :: Char → Text → Text
pattern Cons x xs ← (Text.uncons → Just (x, xs))
```

```
length :: Text → Int
length Nil = 0
length (Cons _ xs) = 1 + length xs
```

The pattern synonym *Nil* matches precisely when the view pattern *Text.null* → *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. While Haskell is a lazy language, 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 $(SJust\ \bot) = \bot$. Curiously, this makes the second equation of v redundant! Since $\bot$ is the only inhabitant of type *Void*, the only inhabitants of *SMaybe Void* are *SNothing* and $\bot$. 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.

2.3.1 Redundancy Versus Inaccessibility. 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 this:

```
u :: () → Int                 u′ :: () → Int
u () | False = 1              u′ () | False = 1
     | True  = 2                    | False = 2
u _          = 3              u′ _          = 3
```

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', one can never reach the right-hand sides of the equations that return 1 and 2, but they cannot be removed so easily. Using the definition above, $u'\ \bot\ =\ \bot$, but if the first two equations were removed, then $u'\ \bot\ = 3$ because the argument is no longer forced by the () pattern. As a result, LYG warns that the first two equations in u' are *inaccessible*, which suggests to the programmer that u' might benefit from a refactor to avoid this (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 tricky corner case, but GHC's users have reported many bugs of this sort (Section 6.2).

2.3.2 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 ⊥, *Nothing*, and (*Just* ⊥). The input ⊥ makes the first equation diverge; *Nothing* matches on the first equation; and (*Just* ⊥) 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:

```
data T a b where          g1 :: T Int b → b → Int      g2 :: T a b → T a b → Int
  T1 :: T Int Bool        g1 T1 False = 0              g2 T1 T1 = 0
  T2 :: T Char Bool       g1 T1 True = 1               g2 T2 T2 = 1
```

When *g1* matches against *T1*, the *b* in the type *T Int b* is known to be a *Bool*, which is why matching 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 [Vytiniotis 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 2 depicts a high-level overview, which divides into three steps:

- First, we desugar the complex source Haskell syntax (cf. Figure 1) into a **guard tree** $t \in$ Gdt (Section 3.1). The language of guard trees is tiny but expressive, and allows the subsequent

Meta variables

x, y, z, f, g, h	Term variables
a, b, c	Type variables
K	Data constructors
P	Pattern synonyms
T	Type constructors
l	Literal
$expr$	Expressions

Pattern syntax

$$
\begin{array}{rcl}
defn & ::= & \overline{clause} \\
clause & ::= & f\ \overline{pat}\ \overline{match} \\
pat & ::= & x \mid _ \mid K\ \overline{pat} \mid x@pat \mid !pat \mid expr \rightarrow pat \\
match & ::= & =\ expr \mid \overline{grhs} \\
grhs & ::= & \mid\ \overline{guard} = expr \\
guard & ::= & pat \leftarrow expr \mid expr \mid \texttt{let}\ x = expr
\end{array}
$$

Fig. 1. Source syntax: A desugared Haskell

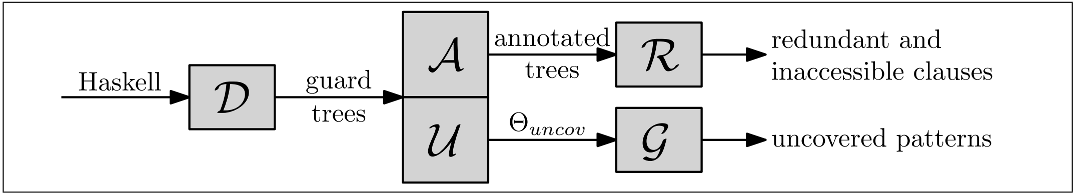

Fig. 2. Bird's eye view of pattern match checking

Guard syntax

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

Clause tree syntax

$$
\begin{array}{rcl}
t \in \text{Gdt} & ::= & \longrightarrow k \mid \genfrac{}{}{0pt}{}{t_1}{t_2} \mid \dashv g \longrightarrow t \\
u \in \text{Ant} & ::= & \longrightarrow \Theta\, k \mid \genfrac{}{}{0pt}{}{u_1}{u_2} \mid \longrightarrow \Theta \lightning \longrightarrow u
\end{array}
$$

Refinement type syntax

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

Fig. 3. IR syntax

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 1) from the intermediate language (Figure 3).
- Correctly accounting for strictness in identifying redundant and inaccessible code (Section 7.5).
- Using detailed term-level reasoning (Figures 6 to 8), which GMTM does not.
- Using *negative information* to sidestep serious performance issues in GMTM without changing the worst-case complexity (Section 7.4). 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 1. 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:

$$
\begin{array}{lll}
f\ (\mathit{Just}\ (!\mathit{xs}, _)) & \mathit{ys}@\mathit{Nothing} & = \mathit{True} \\
f\ \mathit{Nothing} & (g \to \mathit{True}) & = \mathit{False}
\end{array}
$$

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

$$
\begin{array}{l}
\dashv !x_1, \mathit{Just}\ t_1 \leftarrow x_1, !t_1, (t_2, t_3) \leftarrow t_1, !t_2, \mathsf{let}\ \mathit{xs} = t_2, \mathsf{let}\ \mathit{ys} = x_2, !\mathit{ys}, \mathit{Nothing} \leftarrow \mathit{ys} \longrightarrow 1 \\
\dashv !x_1, \mathit{Nothing} \leftarrow x_1, \mathsf{let}\ t_3 = g\ x_2, !t_3, \mathit{True} \leftarrow t_3 \longrightarrow 2
\end{array}
$$

The first line says "evaluate x_1; then match x_1 against $\mathit{Just}\ t_1$; 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 $\dashv g_1, \ldots, g_n — t$ instead of $\dashv g_1 \dashv \cdots \dashv g_n — t$ for notational convenience.

More formally, 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.

$$\boxed{\begin{array}{c}\mathcal{D}(\mathit{defn}) = \mathrm{Gdt}, \quad \mathcal{D}(\mathit{clause}) = \mathrm{Gdt}, \quad \mathcal{D}(\mathit{grhs}) = \mathrm{Gdt} \\ k_{rhs} \text{ is the index of the right hand side } \mathit{rhs}\end{array}}$$

$$\begin{array}{lcl}
\mathcal{D}(\mathit{clause}_1 \ldots \mathit{clause}_n) & = & \begin{array}{l}\top\!\!\!-\ \mathcal{D}(\mathit{clause}_1) \\ \vdash \ldots \\ \llcorner\ \mathcal{D}(\mathit{clause}_n)\end{array} \\[2ex]
\mathcal{D}(f\ \mathit{pat}_1 \ldots \mathit{pat}_n = \mathit{rhs}) & = & \longrightarrow\!\!\!\!\dashv \mathcal{D}(x_1, \mathit{pat}_1) \ldots \mathcal{D}(x_n, \mathit{pat}_n) \longrightarrow k_{rhs} \\
\mathcal{D}(f\ \mathit{pat}_1 \ldots \mathit{pat}_n\ \mathit{grhs}_1 \ldots \mathit{grhs}_m) & = & \longrightarrow\!\!\!\!\dashv \mathcal{D}(x_1, \mathit{pat}_1) \ldots \mathcal{D}(x_n, \mathit{pat}_n) \begin{array}{l}\top\!\!\!-\ \mathcal{D}(\mathit{grhs}_1) \\ \vdash \ldots \\ \llcorner\ \mathcal{D}(\mathit{grhs}_m)\end{array} \\
\mathcal{D}(\mid \mathit{guard}_1 \ldots \mathit{guard}_n = \mathit{rhs}) & = & \longrightarrow\!\!\!\!\dashv \mathcal{D}(\mathit{guard}_1) \ldots \mathcal{D}(\mathit{guard}_n) \longrightarrow k_{rhs}
\end{array}$$

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

$$\begin{array}{lcll}
\mathcal{D}(\mathit{pat} \leftarrow \mathit{expr}) & = & \mathsf{let}\ x = \mathit{expr}, \mathcal{D}(x, \mathit{pat}) & x \text{ fresh} \\
\mathcal{D}(\mathit{expr}) & = & \mathsf{let}\ y = \mathit{expr}, \mathcal{D}(y, \mathit{True}) & y \text{ fresh} \\
\mathcal{D}(\mathtt{let}\ x = \mathit{expr}) & = & \mathsf{let}\ x = \mathit{expr} &
\end{array}$$

$$\begin{array}{lcll}
\mathcal{D}(x, y) & = & \mathsf{let}\ y = x & \\
\mathcal{D}(x, _) & = & \epsilon & \\
\mathcal{D}(x, K\ \mathit{pat}_1 \ldots \mathit{pat}_n) & = & !x, K\ y_1 \ldots y_n \leftarrow x, \mathcal{D}(y_1, \mathit{pat}_1), \ldots, \mathcal{D}(y_n, \mathit{pat}_n) & y_i \text{ fresh } (\dagger) \\
\mathcal{D}(x, y@\mathit{pat}) & = & \mathsf{let}\ y = x, \mathcal{D}(y, \mathit{pat}) & \\
\mathcal{D}(x, !\mathit{pat}) & = & !x, \mathcal{D}(x, \mathit{pat}) & \\
\mathcal{D}(x, \mathit{expr} \rightarrow \mathit{pat}) & = & \mathsf{let}\ y = \mathit{expr}\ x, \mathcal{D}(y, \mathit{pat}) & y \text{ fresh}
\end{array}$$

Fig. 4. $\mathcal{D}$esugaring from source language to Gdt

- Matching a guard tree $\begin{array}{l}\top\!\!\!-\ t_1 \\ \llcorner\ t_2\end{array}$ means matching against t_1; if that succeeds, the overall match succeeds; if not, match against t_2.
- Matching a guard tree $\longrightarrow\!\!\!\!\dashv !x \longrightarrow t$ evaluates x; if that diverges the match diverges; if not match t.
- Matching a guard tree $\longrightarrow\!\!\!\!\dashv K\ \overline{a}\ \overline{\gamma}\ \overline{y} \leftarrow x \longrightarrow t$ matches x against constructor K. If the match succeeds, bind $\overline{a}$ to the type components, $\overline{\gamma}$ to the constraint components and $\overline{y}$ to the term components, then match t. If the constructor match fails, then the entire match fails.
- Matching a guard tree $\longrightarrow\!\!\!\!\dashv \mathsf{let}\ x = e \longrightarrow t$ binds x (lazily) to e, and matches t.

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.

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. $g \rightarrow \mathit{True}$) 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:

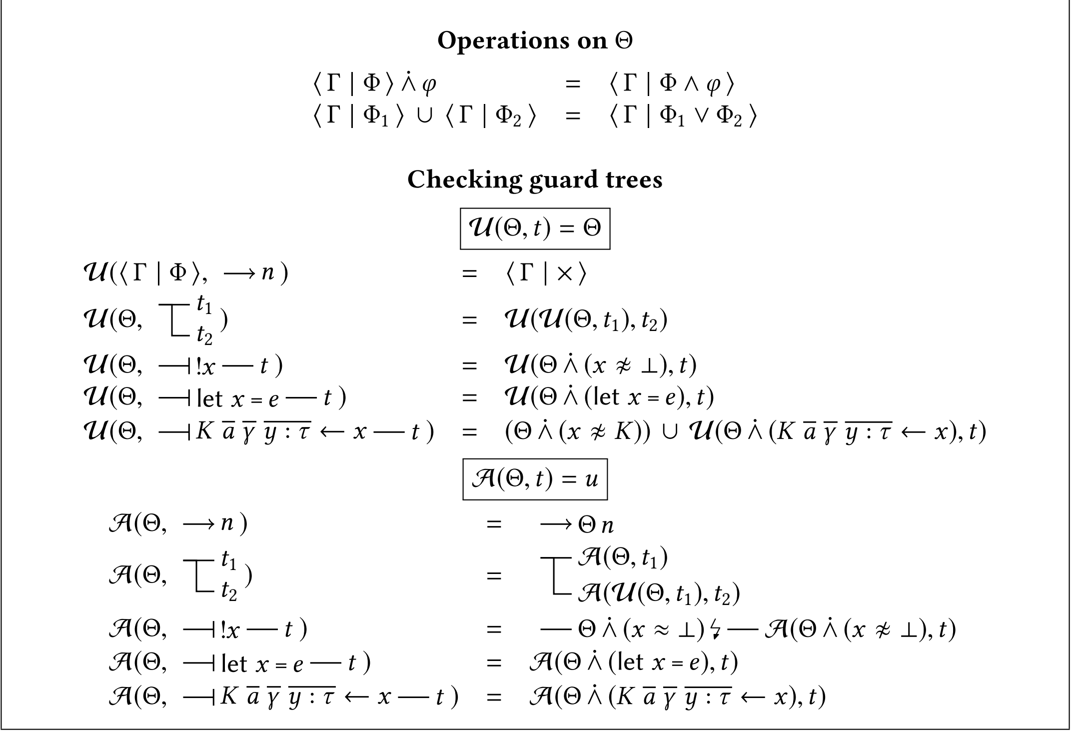

Fig. 5. Coverage checking

```
liftEq Nothing Nothing = True
liftEq mx      (Just y) | Just x ← mx, x == y = True
                        | otherwise           = False
```

It desugars thus:

```
┬┤!mx, Nothing ← mx, !my, Nothing ← my ──────────────────→ 1
└┤!my, Just y ← my ┬┤!mx, Just x ← mx, let t = x == y, !t, True ← t ──→ 2
                   └┤!otherwise, True ← otherwise ───────────→ 3
```

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 $!x$ (see Appendix A). 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!

3.2 Checking Guard Trees

The next step in Figure 2 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 $\Theta = \langle\, x_1{:}\tau_1, \ldots, x_n{:}\tau_n \mid \Phi \,\rangle$ denotes the vector of values $x_1 \ldots x_n$ that satisfy the predicate Φ. For example:

$$\begin{array}{rll} \langle\, x{:}Bool \mid \checkmark \,\rangle & \text{denotes} & \{\bot, True, False\} \\ \langle\, x{:}Bool \mid x \not\approx \bot \,\rangle & \text{denotes} & \{True, False\} \\ \langle\, x{:}Bool \mid True \leftarrow x \,\rangle & \text{denotes} & \{True\} \\ \langle\, mx{:}Maybe\ Bool \mid Just\ x \leftarrow mx, x \not\approx \bot \,\rangle & \text{denotes} & \{Just\ True, Just\ False\} \end{array}$$

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, but we omit that here. 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 $\Theta \mathbin{\dot\wedge} \varphi$ (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 $\mathcal{U}$ncovered 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 ⊤⊏ are in the same places. But in an annotated tree, each $\longrightarrow \Theta\, k$ leaf is annotated with a refinement type Θ describing the input values that will lead to that right-hand side; and each $—\,\Theta\,\lightning\,—$ 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 function[1].

3.3 Reporting Errors

The final step in Figure 2 is to report errors. First, let us focus on reporting missing equations. Consider the following definition

data $T = A \mid B \mid C$
$f\ (Just\ A) = True$

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

$$\Theta_f = \langle\, x{:}Maybe\ T \mid x \not\approx \bot \wedge (x \not\approx Just \vee (Just\ y \leftarrow x \wedge y \not\approx \bot \wedge (y \not\approx A \vee (A \leftarrow y \wedge \times)))) \,\rangle$$

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) = ...
```

$\mathcal{G}$enerating 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

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

Collect accessible ($\overline{k}$), inaccessible ($\overline{n}$) and $\mathcal{R}$edundant ($\overline{m}$) GRHSs

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

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

$$\mathcal{R}(\sqsubset^{t}_{u}) = (\overline{k}\,\overline{k'}, \overline{n}\,\overline{n'}, \overline{m}\,\overline{m'}) \text{ where } \begin{array}{l} (\overline{k}, \overline{n}, \overline{m}) = \mathcal{R}(t) \\ (\overline{k'}, \overline{n'}, \overline{m'}) = \mathcal{R}(u) \end{array}$$

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

Normalised refinement type syntax

∇	$::=$	$\times \mid \langle \Gamma \parallel \Delta \rangle$	Normalised refinement type
Δ	$::=$	$\varnothing \mid \Delta, \delta$	Set of constraints
δ	$::=$	$\gamma \mid x \approx K\ \overline{a}\ \overline{y} \mid x \not\approx K \mid x \approx \bot \mid x \not\approx \bot \mid x \approx y$	Constraints

$\mathcal{G}$enerate inhabitants of Θ

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

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

$\mathcal{N}$ormalise Φ into ∇s

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

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

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

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

$\mathcal{E}$xpand variables to Pat with ∇

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

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

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

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 ∇

the function $\mathcal{R}$, which reports a triple of (accessible, inaccessible, $\mathcal{R}$edundant) GRHSs, as needed in our overall pipeline (Figure 2). $\mathcal{R}$ is defined in Figure 6:

- Having reached a leaf $\longrightarrow \Theta\, k$, 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 $—\Theta ϟ — t$, 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 $\mathcal{R}(t; u)$ is trivial: just combine the classifications of t and u.

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

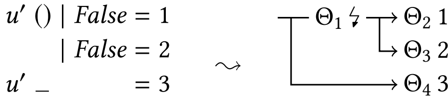

Θ_2 and Θ_3 are uninhabited (because of the *False* guards). But we cannot delete both GRHSs as redundant, because that would make the call $u'\ \bot$ return 3 rather than diverging. Rather, we want to report the first GRHSs as inaccessible, leaving all the others as redundant.

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 $\overline{p}$ (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 $\mathcal{N}(\langle \Gamma \parallel \varnothing \rangle, \Phi)$; see Section 3.6.
- For each such ∇, expand Γ into a list of patterns, by the call $\mathcal{E}(\nabla, \mathsf{dom}(\Gamma))$; 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 $\nabla = \langle \Gamma \parallel \Delta \rangle$:

I1 *Mutual compatibility*: No two constraints in Δ should *conflict* with each other, where $x \approx \bot$ conflicts with $x \not\approx \bot$, and $x \approx K\ _\ _$ conflicts with $x \not\approx K$, for all x.

I2 *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 $\bot$) which x can be instantiated to without contradicting I1; see Section 3.7.

I3 *Triangular form*: A $x \approx y$ constraint implies absence of any other constraint mentioning x in its left-hand side.

I4 *Single solution*: There is at most one positive constructor constraint $x \approx K\ \overline{a}\ \overline{y}$ 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 substition 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

$\mathcal{E}$xpanding 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 $\Delta(x) \approx \mathit{Just}\ y$ 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

$\mathcal{N}$ormalisation, 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 $\oplus_\delta$ to add the simpler constraints to ∇. $\mathcal{N}$, $\oplus_\varphi$ and $\oplus_\delta$ 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 $\oplus_\varphi$ performs some limited, but important reasoning about let bindings: it flattens possibly nested constructor applications, such as $\mathsf{let}\ x = \mathit{Just}\ \mathit{True}$, and asserts that such constructor applications can't be $\bot$. 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, $\oplus_\delta$. When adding $x \approx \mathit{Just}\ y$, Equation (10), the unification procedure will first look for a solution for x with *that same constructor*. Let's say there is $\Delta(x) \approx \mathit{Just}\ u \in \Delta$. Then $\oplus_\delta$ operates on the transitively implied equality $\mathit{Just}\ y \approx \mathit{Just}\ u$ 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 $\Delta(x) \approx \mathit{Nothing}$ or if there was a negative constructor constraint $\Delta(x) \not\approx \mathit{Just}$, the new constraint is incompatible with the existing solution. Otherwise, the constraint is compatible and is added to Δ.

Adding a negative constructor constraint $x \not\approx \mathit{Just}$ is quite similar (Equation (11)), except that we have to make sure that x still satisfies I2, which is checked by the $\nabla \vdash \Delta(x)\ \mathsf{inh}$ judgment (cf. Section 3.7) in Figure 8. Handling positive and negative constraints involving $\bot$ 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 $\langle x : a \,\|\, x \not\approx \bot \rangle$ and try to add $a \sim \mathit{Void}$. Although the type constraint is consistent, x in $\langle x : a \,\|\, x \not\approx \bot, a \sim \mathit{Void} \rangle$ 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 $\oplus_\delta$ 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.

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

$$
\begin{array}{lcll}
\nabla \oplus_\varphi \times & = & \times & (1)\\
\nabla \oplus_\varphi \checkmark & = & \nabla & (2)\\
\langle \Gamma \,\|\, \Delta \rangle \oplus_\varphi K\ \overline{a}\ \overline{\gamma}\ \overline{y{:}\tau} \leftarrow x & = & \langle \Gamma, \overline{a}, \overline{y{:}\tau} \,\|\, \Delta \rangle \oplus_\delta \overline{\gamma} \oplus_\delta \overline{y' \not\approx \bot} \oplus_\delta x \approx K\ \overline{a}\ \overline{y} & (3)\\
 & & \quad \text{where } \overline{y'} \subseteq \overline{y} \text{ bind strict fields} & \\
\langle \Gamma \,\|\, \Delta \rangle \oplus_\varphi \mathsf{let}\ x{:}\tau = K\ \overline{\sigma}\ \overline{\gamma}\ \overline{e} & = & \langle \Gamma, x{:}\tau, \overline{a} \,\|\, \Delta \rangle \oplus_\delta \overline{a \sim \sigma} \oplus_\delta x \not\approx \bot \oplus_\delta x \approx K\ \overline{a}\ \overline{y} & (4)\\
 & & \quad \oplus_\varphi \overline{\mathsf{let}\ y{:}\tau' = e} \qquad \text{where } \overline{a}\,\overline{y} \ \#\ \Gamma, \overline{e{:}\tau'} & \\
\langle \Gamma \,\|\, \Delta \rangle \oplus_\varphi \mathsf{let}\ x{:}\tau = y & = & \langle \Gamma, x{:}\tau \,\|\, \Delta \rangle \oplus_\delta x \approx y & (5)\\
\langle \Gamma \,\|\, \Delta \rangle \oplus_\varphi \mathsf{let}\ x{:}\tau = e & = & \langle \Gamma, x{:}\tau \,\|\, \Delta \rangle & (6)\\
\langle \Gamma \,\|\, \Delta \rangle \oplus_\varphi \varphi & = & \langle \Gamma \,\|\, \Delta \rangle \oplus_\delta \varphi & (7)
\end{array}
$$

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

$$
\begin{array}{lcl}
\times \oplus_\delta \delta & = \times & (8)\\
\langle \Gamma \,\|\, \Delta \rangle \oplus_\delta \gamma & = \begin{cases} \langle \Gamma \,\|\, (\Delta, \gamma) \rangle & \text{if type checker deems } \gamma \text{ compatible with } \Delta \\ & \text{and } \forall x \in \mathrm{dom}(\Gamma) : \langle \Gamma \,\|\, (\Delta, \gamma) \rangle \vdash x \ \mathsf{inh} \\ \times & \text{otherwise} \end{cases} & (9)\\
\langle \Gamma \,\|\, \Delta \rangle \oplus_\delta x \approx K\ \overline{a}\ \overline{y} & = \begin{cases} \times & \text{if } \Delta(x) \approx K'\ \overline{b}\ \overline{z} \in \Delta \text{ and } K \neq K' \\ \times & \text{if } \Delta(x) \not\approx K \in \Delta \\ \langle \Gamma \,\|\, \Delta \rangle \oplus_\delta \overline{a \sim b} \oplus_\delta \overline{y \approx z} & \text{if } \Delta(x) \approx K\ \overline{b}\ \overline{z} \in \Delta \\ \langle \Gamma \,\|\, (\Delta, \Delta(x) \approx K\ \overline{a}\ \overline{y}) \rangle & \text{otherwise} \end{cases} & (10)\\
\langle \Gamma \,\|\, \Delta \rangle \oplus_\delta x \not\approx K & = \begin{cases} \times & \text{if } \Delta(x) \approx K\ \overline{a}\ \overline{y} \in \Delta \\ \times & \text{if not } \langle \Gamma \,\|\, (\Delta, \Delta(x) \not\approx K) \rangle \vdash x \ \mathsf{inh} \\ \langle \Gamma \,\|\, (\Delta, \Delta(x) \not\approx K) \rangle & \text{otherwise} \end{cases} & (11)\\
\langle \Gamma \,\|\, \Delta \rangle \oplus_\delta x \approx \bot & = \begin{cases} \times & \text{if } \Delta(x) \not\approx \bot \in \Delta \\ \langle \Gamma \,\|\, (\Delta, \Delta(x) \approx \bot) \rangle & \text{otherwise} \end{cases} & (12)\\
\langle \Gamma \,\|\, \Delta \rangle \oplus_\delta x \not\approx \bot & = \begin{cases} \times & \text{if } \Delta(x) \approx \bot \in \Delta \\ \times & \text{if not } \langle \Gamma \,\|\, (\Delta, \Delta(x) \not\approx \bot) \rangle \vdash x \ \mathsf{inh} \\ \langle \Gamma \,\|\, (\Delta, \Delta(x) \not\approx \bot) \rangle & \text{otherwise} \end{cases} & (13)\\
\langle \Gamma \,\|\, \Delta \rangle \oplus_\delta x \approx y & = \begin{cases} \langle \Gamma \,\|\, \Delta \rangle & \text{if } x' = y' \\ \langle \Gamma \,\|\, ((\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{array}
$$

$$
\boxed{\Delta \setminus x = \Delta}
\qquad\qquad
\boxed{\Delta|_x = \Delta}
$$

$$
\begin{array}{rcl}
\varnothing \setminus x & = & \varnothing \\
(\Delta, x \approx K\ \overline{a}\ \overline{y}) \setminus x & = & \Delta \setminus x \\
(\Delta, x \not\approx K) \setminus x & = & \Delta \setminus x \\
(\Delta, x \approx \bot) \setminus x & = & \Delta \setminus x \\
(\Delta, x \not\approx \bot) \setminus x & = & \Delta \setminus x \\
(\Delta, \delta) \setminus x & = & (\Delta \setminus x), \delta
\end{array}
\qquad
\begin{array}{rcl}
\varnothing|_x & = & \varnothing \\
(\Delta, x \approx K\ \overline{a}\ \overline{y})|_x & = & \Delta|_x, x \approx K\ \overline{a}\ \overline{y} \\
(\Delta, x \not\approx K)|_x & = & \Delta|_x, x \not\approx K \\
(\Delta, x \approx \bot)|_x & = & \Delta|_x, x \approx \bot \\
(\Delta, x \not\approx \bot)|_x & = & \Delta|_x, x \not\approx \bot \\
(\Delta, \delta)|_x & = & \Delta|_x
\end{array}
$$

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

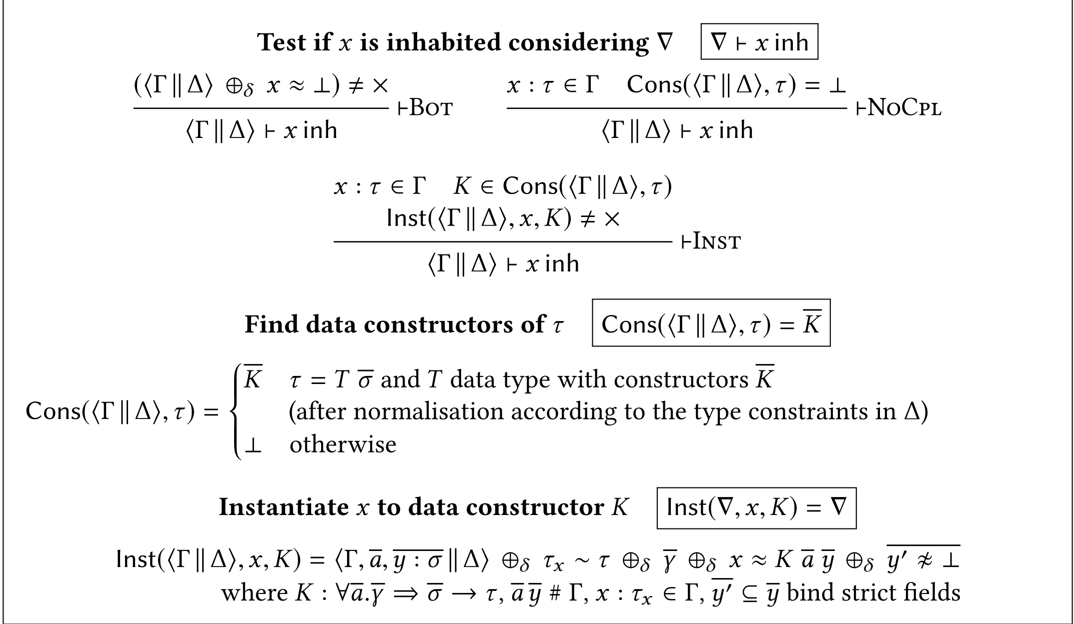

Fig. 8. Testing for inhabitation

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* $\nabla \vdash x\,\mathsf{inh}$, 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 $x \not\approx \mathit{Just}, x \not\approx \mathit{Nothing}, x \not\approx \bot$ does not satisfy I2, because no value of x satisfies all those constraints.

The ⊢Bot judgment of $\nabla \vdash x\,\mathsf{inh}$ tries to instantiate x to $\bot$ to conclude that x is inhabited. ⊢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 ⊢NoCpl will accept unconditionally when its type is not a data type, i.e. for $x : \mathit{Int} \to \mathit{Int}$.

Note that the outlined approach is complete in the sense that $\nabla \vdash x\,\mathsf{inh}$ 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 $\vdash\ \mathsf{inh}$ relation is undecidable! Consider the following example:

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

This is exhaustive, because T is an uninhabited type. Upon adding the constraint $x \not\approx \mathit{SNothing}$ on the match variable x via $\oplus_\delta$, we perform an inhabitation test, which tries to instantiate the *SJust* constructor via ⊢Inst. That implies adding (via $\oplus_\delta$) the constraints $x \approx \mathit{SJust}\ y, y \not\approx \bot$, the latter of which leads to an inhabitation test on y. That leads to instantiation of the MkT constructor, which leads to constraints $y \approx \mathit{MkT}\ z, z \not\approx \bot$, 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.

3.8 A Note on Precision

Using fuel to limit the number of inhabitation tests is one example where LYG sacrifices a small amount of precision in its warnings. It is worth noting that this does not impact LYG's soundness. 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". We can broadly describe three places where LYG overapproximates:

- LYG can run out of fuel for inhabitation testing (Section 3.7).
- Throttling (Section 5.2) is useful when implementing LYG.
- 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 7.2.1).

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

$$
\begin{array}{ll}
f\ \mathit{True} & = 1 \\
f\ x & = \ldots(\mathbf{case}\ x\ \mathbf{of}\ \{\mathit{False} \rightarrow 2; \mathit{True} \rightarrow 3\})\ \ldots
\end{array}
$$

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 $\Theta_2 = \langle\, x : \mathit{Bool} \mid x \not\approx \bot, x \not\approx \mathit{True} \,\rangle$, 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 $\langle\, x : \mathit{Bool} \mid \checkmark \,\rangle$. 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 1, 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*:

$$
\begin{array}{ll}
\mathit{absurd1}\ _\ = \bot & \mathit{absurd1}, \mathit{absurd2}, \mathit{absurd3} :: \mathit{Void} \rightarrow a \\
\mathit{absurd2}\ !_ = \bot & \mathit{absurd3}\ x = \mathbf{case}\ x\ \mathbf{of}\ \{\,\}
\end{array}
$$

absurd1 returns $\bot$ when called with $\bot$, thus masking the original $\bot$ with the error thrown by $\bot$. *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 $\bullet_{\text{Gdt}}$ to Gdt and $\bullet_{\text{Ant}}$ to Ant. This is how they affect the checking process:

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

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 $x \not\approx \bot$ constraint *before* traversing the guard tree, thus checking starts starts with $\mathcal{U}(\langle\, \Gamma \mid x \not\approx \bot \,\rangle, \bullet_{\text{Gdt}})$.

4.3 View Patterns

Our source syntax had support for view patterns to start with (cf. Figure 1). 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:

$$\begin{array}{l} \vdash \text{let } y_1 = \textit{reverse}\ x_1, !y_1, \textit{Nothing} \leftarrow y_1 \longrightarrow 1 \\ \vdash \text{let } y_2 = \textit{reverse}\ x_1, !y_2, \textit{Just}\ t_1 \leftarrow y_2, !t_1, (t_2, t_3) \leftarrow t_1 \longrightarrow 2 \end{array}$$

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 $\oplus_\varphi$. Then we can handle the new constraint in $\oplus_\delta$, as follows:

$$\delta = \ldots \mid e \approx x \qquad \langle \Gamma \,\|\, \Delta \rangle \oplus_\varphi \text{let } x : \tau = e = \langle \Gamma, x : \tau \,\|\, \Delta \rangle \oplus_\delta e \approx x$$

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

Where $\equiv_\Delta$ 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 and sacrifice reasoning modulo Δ in doing so. Plugging in an SMT solver to decide $\equiv_\Delta$ 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*:

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

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

$$\begin{array}{l} \textit{pattern}\ P = () \\ \textit{pattern}\ Q = () \\ n = \textbf{case}\ P\ \textbf{of}\ Q \longrightarrow 1; P \longrightarrow 2 \end{array}$$

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 $\oplus_\delta$ dealing with positive ConLike constraints $x \approx C\ \overline{a}\ \overline{y}$:

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

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 them for the purposes of reporting warnings. Fixing that requires a bit of boring engineering.

Another subtle point appears in rule (†) 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 ⊢Inst from Figure 8 currently making sure that this `COMPLETE` set is in fact inhabited. We also have ⊢NoCpl that handles the case when we can't find *any* `COMPLETE` set for the given type (think $x : Int \to Int$). 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:

$$\frac{(\langle \Gamma \,\|\, \Delta \rangle \oplus_\delta x \approx \bot) \neq \times}{\langle \Gamma \,\|\, \Delta \rangle \vdash x \ \mathsf{inh}}\ \vdash\textsc{Bot} \qquad \frac{x : \tau \in \Gamma \quad \mathsf{Cons}(\langle \Gamma \,\|\, \Delta \rangle, \tau) = \overline{C_1, ..., C_{n_i}}^{\,i} \qquad \overline{\mathsf{Inst}(\langle \Gamma \,\|\, \Delta \rangle, x, C_j) \neq \times}^{\,i}}{\langle \Gamma \,\|\, \Delta \rangle \vdash x \ \mathsf{inh}}\ \vdash\textsc{Inst}$$

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

Cons was changed to return a list of all available `COMPLETE` sets, and ⊢Inst tries to find an inhabiting ConLike in each one of them in turn. Note that ⊢NoCpl is gone, because it coincides with ⊢Inst for the case where the list returned by Cons was empty. The judgment has become simpler and and more general at the same time! Note that checking against multiple `COMPLETE` sets so frequently is computationally intractable. We will worry about that in Section 5.

4.6 Other Extensions

We consider further extensions, including overloaded literals, newtypes, and a strict-by-default (or total) language semantics, in Appendix A.

$$\boxed{\overline{\nabla} \mathbin{\dot{\oplus}_\varphi} \varphi = \overline{\nabla}}$$

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

$$\boxed{\mathcal{UA}(\overline{\nabla}, t) = (\overline{\nabla}, \mathrm{Ant})}$$

$$\begin{array}{lcl}
\mathcal{UA}(\overline{\nabla},\ \longrightarrow n) & = & (\epsilon,\ \longrightarrow \overline{\nabla}\, n) \\
\mathcal{UA}(\overline{\nabla},\ \sqsubset^{t_1}_{t_2}) & = & (\overline{\nabla}_2,\ \sqsubset^{u_1}_{u_2}) \text{ where } \begin{array}{l} (\overline{\nabla}_1, u_1) = \mathcal{UA}(\overline{\nabla}, t_1) \\ (\overline{\nabla}_2, u_2) = \mathcal{UA}(\overline{\nabla}_1, t_2) \end{array} \\
\mathcal{UA}(\overline{\nabla},\ \dashv !x \longrightarrow t) & = & \longrightarrow \overline{\nabla} \mathbin{\dot{\oplus}_\varphi} (x \approx \bot) \lightning \longrightarrow u \\
 & & \quad \text{where } (\overline{\nabla}', u) = \mathcal{UA}(\overline{\nabla} \mathbin{\dot{\oplus}_\varphi} (x \not\approx \bot), t) \\
\mathcal{UA}(\overline{\nabla},\ \dashv \mathrm{let}\ x = e \longrightarrow t) & = & \mathcal{UA}(\overline{\nabla} \mathbin{\dot{\oplus}_\varphi} (\mathrm{let}\ x = e), t) \\
\mathcal{UA}(\overline{\nabla},\ \dashv K\ \overline{a}\ \overline{\gamma}\ \overline{y:\tau} \leftarrow x \longrightarrow t) & = & ((\overline{\nabla} \mathbin{\dot{\oplus}_\varphi} (x \not\approx K))\, \overline{\nabla}', u) \\
 & & \quad \text{where } (\overline{\nabla}', u) = \mathcal{UA}(\overline{\nabla} \mathbin{\dot{\oplus}_\varphi} (K\ \overline{a}\ \overline{\gamma}\ \overline{y:\tau} \leftarrow x), t)
\end{array}$$

Fig. 9. Fast coverage checking

5 IMPLEMENTATION

We have implemented LYG in a to-be-released version of GHC[2], including all extensions in Section 4 (except for strict-by-default source syntax in Appendix A). 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. That can amount to exponential growth of the refinement predicate in the worst case 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 9. 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 $\dot{\oplus}_\varphi$, which filters out any ∇ that would become empty.

[2]The functionality described in this paper will be available in GHC 8.12 and later.

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:

$$
\begin{array}{l}
\textbf{data}\ T = A \mid B; f1, f2 :: Int \rightarrow T \\
g\ _ \\
\quad \mid A \leftarrow f1\ 1,\ \ A \leftarrow f2\ 1\ \ = () \\
\quad \mid A \leftarrow f1\ 2,\ \ A \leftarrow f2\ 2\ \ = () \\
\quad \ldots \\
\quad \mid A \leftarrow f1\ N, A \leftarrow f2\ N = ()
\end{array}
\qquad
\begin{array}{l}
\dashv \text{let } a_1 = f1\ 1, !a_1, A \leftarrow a_1, \text{let } b_1 = f2\ 1, !b_1, A \leftarrow b_1 \longrightarrow 1 \\
\dashv \text{let } a_2 = f1\ 2, !a_2, A \leftarrow a_2, \text{let } b_2 = f2\ 2, !b_2, A \leftarrow b_2 \longrightarrow 2 \\
\dashv \cdots \longrightarrow \cdots \\
\dashv \text{let } a_N = f1\ N, !a_N, A \leftarrow a_N, \text{let } b_N = f2\ N, !b_N, A \leftarrow b_N \longrightarrow \text{N}
\end{array}
$$

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 input ∇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:

$$
\begin{array}{l}
\mathcal{UA}(\overline{\nabla},\ \longrightarrow\!\!\dashv K\ \overline{a}\ \overline{\gamma}\ \overline{y:\tau} \leftarrow x \longrightarrow t\ ) = (\left\lfloor (\overline{\nabla} \mathbin{\dot{\oplus}_\varphi} (x \not\approx K))\ \overline{\nabla}' \right\rfloor_{\overline{\nabla}}, u) \\
\qquad\qquad\qquad\qquad\qquad\qquad \text{where } (\overline{\nabla}', u) = \mathcal{UA}(\overline{\nabla} \mathbin{\dot{\oplus}_\varphi} (K\ \overline{a}\ \overline{\gamma}\ \overline{y:\tau} \leftarrow x), t)
\end{array}
$$

where the new throttling operator $\lfloor _ \rfloor_{_}$ is defined simply as

$$
\left\lfloor \overline{\nabla} \right\rfloor_{\overline{\nabla}'} \quad = \quad \begin{cases} \overline{\nabla} & \text{if } |\{\overline{\nabla}\}| \leqslant K \\ \overline{\nabla}' & \text{otherwise} \end{cases}
$$

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.

5.3 Maintaining Residual `COMPLETE` Sets

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

$$
\begin{array}{l}
\textbf{data}\ T = A1 \mid \ldots \mid A1000 \\
\textit{pattern}\ P = \ldots \\
\ \ \{\text{-\# COMPLETE A1, P \#-}\}
\end{array}
\qquad
\begin{array}{l}
f\ A1 \quad\ \ = 1 \\
f\ A2 \quad\ \ = 2 \\
\ \ \ldots \\
f\ A1000 = 1000
\end{array}
$$

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 ⊢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 $b\ \mathit{True} = ()$. The computed uncovered set of b is the normalised refinement type $\nabla_b = \langle x : \mathit{Bool} \parallel x \not\approx \bot, x \not\approx \mathit{True}\rangle$, 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 $\nabla_b \oplus_\delta x \approx \mathit{True}$ due to $x \not\approx \mathit{True} \in \nabla_b$. What remains is the refinement $\nabla_b \oplus_\delta x \approx \mathit{False} = \langle x : \mathit{Bool} \parallel x \not\approx \bot, x \not\approx \mathit{True}, x \approx \mathit{False}\rangle$, 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}
```

6 EVALUATION

To put the new coverage checker to the test, we performed a survey of real-world Haskell code using the `head.hackage` repository [3]. `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/head.hackage/commit/30a310fd8033629e1cbb5a9696250b22db5f7045

	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%

Fig. 10. 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.

$go'\ _\ _\ _\ xs \mid False = error\ (show\ xs)$
$go'\ _\ _\ _\ xs = err\ xs$

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.

We may consider adding a primitive function *considerAccessible* such that *considerAccessible False* does not get marked as redundant in order to support use cases like `HsYAML`'s. The unreachable code in `Cabal` and `network` is of a similar caliber and would also benefit from *considerAccessible*.

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. Figure 10 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, 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. Figure 10 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.

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!

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 RELATED WORK

7.1 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.

7.1.1 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.

7.1.2 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 $\oplus_\varphi$ function, which is roughly a counterpart to GMTM's term oracle, can then reason about terms arising from patterns. While $\oplus_\varphi$ 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 $\oplus_\varphi$.

7.2 Comparison with Similar Coverage Checkers

7.2.1 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 $\oplus_\varphi$ with SMT-like reasoning about booleans and linear integer arithmetic.

7.2.2 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.

*7.2.3 Case Trees in Dependently Typed Languages. Case tree*s [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.

7.2.4 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:

$$\begin{array}{l} \mathit{fibPartial} :: \mathit{Integer} \to \mathit{Integer} \\ \mathit{fibPartial}\ 0 = 0 \\ \mathit{fibPartial}\ 1 = 1 \end{array}$$

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

$$\mathit{error} :: \{\, v : \mathit{String} \mid \mathit{false} \,\} \to a$$

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.

7.3 Other Representations of Constraints

7.3.1 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 $\oplus_\varphi$ 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.

7.3.2 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.

7.4 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 7.1) and Maranget [2007] (Section 7.2.2), consider the following example:

```
pattern True′ = True                    f False = 1
 {-# COMPLETE True', False #-}          f True′ = 2
                                        f True  = 3
```

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:

```
data T = A1 | ... | A1000               h A1 _  = 1
                                        h _  A1 = 2
```

In h, GMTM would split the value vector (which is like LYG's Δs 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*.

7.5 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 coverage [4], would fare when checking functions like v:

```
(*OCaml*)
type void =|;;
let v (None : void option) : int = 0;;
let v' (o : void option) : int =
  match o with
    None   -> 0
  | Some _ -> 1;;
```

```
-- Idris
v : Maybe Void -> Int
v Nothing = 0

v' : Maybe Void -> Int
v' Nothing = 0
v' (Just _) = 1
```

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:

```
(*OCaml*)
type t = MkT of t;;
let f (None : t option) : int = 0;;
```

```
-- Idris
data T : Type where
  MkT : T -> T

f : Maybe T -> Int
f Nothing = 0
```

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.

8 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. This allows LYG to report more accurate warnings while also avoiding performance issues when checking complex programs.

ACKNOWLEDGMENTS

We would like to thank the anonymous ICFP reviewers for their feedback, as well as anyone else who provided feedback on earlier drafts: Henning Dieterichs, Martin Hecker, Sylvain Henry, Philipp Krüger, Luc Maranget and Sebastian Ullrich.

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

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

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

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.

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

Matthew Pickering, Gergő Érdi, 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, Chiyan 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

A APPENDIX

A.1 Literals

The source syntax in Figure A.11 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 $\vdash$ 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.

A.2 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 *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*, *g3* match on a *newtype N*, while functions *h1*, *h2*, *h3* match on a *data type D*:

$$cl ::= K \mid N \qquad \begin{array}{ll} N \in & \mathsf{NT} \\ C \in & K \mid N \end{array}$$

$$\mathcal{D}(x, N\ pat_1 \ldots pat_n) = N\ y_1 \ldots y_n \leftarrow x, \mathcal{D}(y_1, pat_1), \ldots, \mathcal{D}(y_n, pat_n)$$

$$\Delta_{\mathrm{NT}}(x) \quad = \quad \begin{cases} \Delta_{\mathrm{NT}}(y) & x \approx y \in \Delta \text{ or } x \approx N\ \overline{a}\ y \in \Delta \\ x & \text{otherwise} \end{cases}$$

$$\langle \Gamma \,\|\, \Delta \rangle \oplus_\varphi \mathsf{let}\ x{:}\tau = K\ \overline{\sigma}\ \overline{\gamma}\ \overline{e} = \ldots \text{as before} \ldots \qquad (4a)$$

$$\langle \Gamma \,\|\, \Delta \rangle \oplus_\varphi \mathsf{let}\ x{:}\tau = N\ \overline{\sigma}\ e \quad = \langle \Gamma, x{:}\tau, \overline{a} \,\|\, \Delta \rangle \oplus_\delta \overline{a \sim \sigma} \oplus_\delta x \approx N\ \overline{a}\ y \oplus_\varphi \mathsf{let}\ y{:}\tau' = e \qquad (4b)$$
$$\text{where } \overline{a}\, y \;\#\; \Gamma,\ e{:}\tau'$$

$$\langle \Gamma \,\|\, \Delta \rangle \oplus_\delta x \approx K\ \overline{a}\ \overline{y} \quad = \ldots \text{as before} \ldots \qquad (10a)$$

$$\langle \Gamma \,\|\, \Delta \rangle \oplus_\delta x \approx N\ \overline{a}\ y = \begin{cases} \times & \text{if } x' \not\approx N \in \Delta \\ \langle \Gamma \,\|\, \Delta \rangle \oplus_\delta \overline{a \sim b} \oplus_\delta y \approx z & \text{if } x' \approx N\ \overline{b}\ z \in \Delta \\ \langle \Gamma \,\|\, \Delta \rangle & \text{if } x' = \Delta_{\mathrm{NT}}(y') \\ \langle \Gamma \,\|\, ((\Delta \setminus x'), x' \approx N\ \overline{a}\ y') \rangle \oplus_\delta (\Delta|_{x'}[y'/x']) & \text{otherwise} \end{cases} \qquad (10b)$$
$$\text{where } x' = \Delta(x) \text{ and } y' = \Delta(y)$$

$$\langle \Gamma \,\|\, \Delta \rangle \oplus_\delta x \not\approx C \quad = \ldots \text{similar to before with } C \text{ instead of } K \ldots \qquad (11)$$

$$\langle \Gamma \,\|\, \Delta \rangle \oplus_\delta x \approx \bot \quad = \begin{cases} \times & \text{if } \Delta_{\mathrm{NT}}(x) \not\approx \bot \in \Delta \\ \langle \Gamma \,\|\, (\Delta, \Delta_{\mathrm{NT}}(x) \approx \bot) \rangle & \text{otherwise} \end{cases} \qquad (12)$$

$$\langle \Gamma \,\|\, \Delta \rangle \oplus_\delta x \not\approx \bot \quad = \begin{cases} \times & \text{if } \Delta_{\mathrm{NT}}(x) \approx \bot \in \Delta \\ \times & \text{if not } \langle \Gamma \,\|\, (\Delta, \Delta_{\mathrm{NT}}(x) \not\approx \bot) \rangle \vdash x\ \mathsf{inh} \\ \langle \Gamma \,\|\, (\Delta, \Delta_{\mathrm{NT}}(x) \not\approx \bot) \rangle & \text{otherwise} \end{cases} \qquad (13)$$

Fig. A.11. Extending coverage checking to handle newtypes

```
newtype N a = MkN a
g1 :: N Void → Bool → Int             g2 :: N () → Bool → Int
g1 _         True = 1                 g2 !(MkN _) True = 1
g1 (MkN _)   True = 2  -- Redundant   g2 (MkN !_) True = 2  -- Redundant
g1 !_        True = 3  -- Inaccessible g2 _        _    = 3

data D a = MkD a
h1 :: D Void → Bool → Int             h2 :: D () → Bool → Int
h1 _         True = 1                 h2 !(MkD _) True = 1
h1 (MkD _)   True = 2  -- Inaccessible h2 (MkD !_) True = 2  -- Inaccessible
h1 !_        True = 3  -- Redundant   h2 _        _    = 3
```

If the first equation of *h1* fails to match (because the second argument is *False*), the second equation may diverge when matching against (*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 (*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 $N\ \mathit{Void}$ was deprived of its sole inhabitant $\bot \equiv \mathit{MkN}\ \bot$ 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 A.11 outlines a solution that handles all these cases correctly:

- A newtype pattern match $N\ pat_1 \ldots pat_n$ 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 $\oplus_\varphi$, responsible for reasoning about **let** bindings, has a special case for newtypes that omits the $x \not\approx \bot$ 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 $x \approx N\ \overline{a}\ y$ with $\Delta_{\mathrm{NT}}(x)$.
- The most important usage of $\Delta_{\mathrm{NT}}(x)$ is in the changed Equations (12) and (13) of $\oplus_\delta$, where we now check $\bot$ constraints modulo $\Delta_{\mathrm{NT}}(x)$.
- Equation (10) (previously handling $x \approx K\ \overline{a}\ \overline{y}$) has been split into Equation (10*a*) that handles positive data constructor constraints, as before, and (10*b*), which handles positive newtype constructor constraints.
- The first two cases of the new Equation (10*b*) handle any existing positive or negative constructor constraints in Δ, as with Equation (10). 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 $\bot$ constraints are moved from $\Delta(x)$ to the new representative $\Delta'_{\mathrm{NT}}(x)$, which will be $\Delta'_{\mathrm{NT}}(y)$ in the returned Δ'.

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

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

We want Δ to uphold the semantic equation $\bot \equiv N\bot$. In particular, whenever we have $x \approx N\ \overline{a}\ y$, we want $x \approx \bot$ iff $y \approx \bot$ (similarly for $x \not\approx \bot$). Equations (10b), (12) and (13) facilitate just that, modulo $\Delta_{\mathrm{NT}}(x)$. Finally, a new invariant I5 relates positive newtype constructor equalities to $\bot$ constraints:

I5 *Newtype erasure*: Whenever $x \approx N\ \overline{a}\ y \in \Delta$, we have $x \approx \bot \in \Delta$ if and only if $y \approx \bot \in \Delta$, and $x \not\approx \bot \in \Delta$ if and only if $y \not\approx \bot \in \Delta$.

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.

A.3 Strictness and Totality

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 or Idris (or even Rust).

First off, both OCaml and Idris offer language support for laziness and lazy pattern matches, so the question rather becomes whether the gained simplification is actually worth risking unusable or even unsound warning messages when making use of laziness. If the answer is "No", then there isn't anything to simplify, just relatively more $x \not\approx \bot$ constraints to handle.

Otherwise, in a completely eager language we could simply drop $!x$ from Grd and $—\lightning—$ from Ant. Actually, Ant and $\mathcal{R}$ could go 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 $x \approx \bot$ and $x \not\approx \bot$ constraints either. Most importantly, the ⊢Bot judgment form has to go, because $\bot$ does not inhabit any types anymore.

Note that in a total language such as Agda, reasoning about $x \approx \bot$ makes no sense to begin with! All the same simplifications apply.