

Safe Zero-cost Coercions for Haskell (extended version)¹

Joachim Breitner
Karlsruhe Institute of Technology
breitner@kit.edu

Richard A. Eisenberg
University of Pennsylvania
eir@cis.upenn.edu

Simon Peyton Jones
Microsoft Research
simonpj@microsoft.com

Stephanie Weirich
University of Pennsylvania
sweirich@cis.upenn.edu

Abstract

Generative type abstractions – present in Haskell, OCaml, and other languages – are useful concepts to help prevent programmer errors. They serve to create new types that are distinct at compile time but share a run-time representation with some base type. We present a new mechanism that allows for zero-cost conversions between generative type abstractions and their representations, even when such types are deeply nested. We prove type safety in the presence of these conversions and have implemented our work in GHC.

Categories and Subject Descriptors D.3.3 [Programming Languages]: Language Constructs and Features—abstract data types; F.3.3 [Logics and Meanings of Programs]: Studies of Program Constructs—Type structure

Keywords Haskell; Coercion; Type class; Newtype deriving

1. Introduction

Modular languages support *generative type abstraction*, the ability for programmers to define application-specific types, and rely on the type system to distinguish between these new types and their underlying representations. Type abstraction is a powerful tool for programmers, enabling both flexibility (implementors can change representations) and security (implementors can maintain invariants about representations). Typed languages provide these mechanisms with zero run-time cost – there should be no performance penalty for creating abstractions – using mechanisms such as ML’s module system [MTHM97] and Haskell’s **newtype** declaration [Mar10].

For example, a Haskell programmer might create an abstract type for HTML data, representing them as Strings (Figure 1). Although String values use the same patterns of bits in memory as HTML values, the two types are distinct. That is, a String will not be accepted by a function expecting an HTML. The constructor `Mk` converts a String to an HTML (see function `text`), while using `Mk` in a pattern converts in the other direction (see function `unMk`). By exporting the type `HTML`, but not its data constructor, module `Html` ensures that the type `HTML` is *abstract* – clients cannot make arbitrary strings into `HTML` – and thereby prevent cross-site scripting attacks.

Using **newtype** for abstraction in Haskell has always suffered from an embarrassing difficulty. Suppose in the module

```
module Html( HTML, text, unMk, ... ) where
  newtype HTML = Mk String
  unMk :: HTML → String
  unMk (Mk s) = s
  text :: String → HTML
  text s = Mk (escapeSpecialCharacters s)
```

Figure 1. An abstraction for HTML values

`Html`, the programmer wants to break HTML data into a list of lines:

```
linesH :: HTML → [HTML]
linesH h = map Mk (lines (unMk h))
```

To get the resulting `[HTML]` we are forced to map `Mk` over the list. Operationally, this map is the identity function – the run-time representation of `[String]` is identical to `[HTML]` – *but it will carry a run-time cost nevertheless*. The optimiser in the Glasgow Haskell Compiler (GHC) is powerless to fix the problem, because it works over a *typed* intermediate language; the `Mk` constructor changes the type of its operand, and hence cannot be optimised away. There is nothing that the programmer can do to prevent this run-time cost. What has become of the claim of zero-overhead abstraction?

In this paper we describe a robust, simple mechanism that programmers can use to solve this problem, making the following contributions:

- We describe the design of *safe coercions* (Section 2), which introduces the function

```
coerce :: Coercible a b ⇒ a → b
```

and a new type class `Coercible`. This function performs a zero-cost conversion between two types `a` and `b` that have the same representation. The crucial question becomes *what instances of Coercible exist?* We give a simple but non-obvious strategy (Sections 2.1–2.2), expressed largely in the familiar language of Haskell type classes.

- We formalise `Coercible` by translation into GHC’s intermediate language System FC, augmented with the concept of *roles* (Section 2.2), adapted from prior work [WVPZ11]. Our new contribution is a significant simplification of the roles idea in System FC; we formalise this simpler system

¹This is a substantial revision to published work [BEPW14].

and give the usual proofs of preservation and progress in Section 4.

- Adding safe coercions to the source language raises new issues for abstract types, and for the coherence of type elaboration. We articulate the issues, and introduce *role annotations* to solve them (Section 3).
- It would be too onerous to insist on programmer-supplied role annotations for every type, so we give a *role inference algorithm* in Section 5.
- To support our claim of practical utility, we have implemented the whole scheme in GHC (Section 6), and evaluated it against thousands of Haskell libraries (Section 9).

Our work finally resolves a notorious and long-standing bug in GHC (#1496), which concerns the interaction of new-type coercions with type families (Section 7). While earlier work [WVPZ11] was motivated by the same bug, it was too complicated to implement. Our new approach finds a sweet spot, offering a considerably simpler system in exchange for a minor loss of expressiveness (Sections 8 and 10).

As this work demonstrates, the interactions between type abstraction and advanced type system features, such as type families and GADTs, are subtle. The ability to create and enforce zero-cost type abstraction is not unique to Haskell – notably the ML module system also provides this capability, and more. As a result, OCaml developers are now grappling with similar difficulties. We discuss the connection between roles and OCaml’s variance annotations (Section 8), as well as other related work.

2. The design and interface of Coercible

We begin by focusing exclusively on the programmer’s-eye-view of safe coercions. We need no new syntax; rather, the programmer simply sees a new API, provided in just two declarations:

```
class Coercible a b
coerce :: Coercible a b => a -> b
```

The type class `Coercible` is abstract, i.e. its methods are not visible. It differs from other type classes in a few minor points: The user cannot create manual instances; instances are automatically generated by the compiler; and the visibility of instances is conditional. Generally, users can think of it as a normal type class, which is a nice property of the design.

The key principle is this: *If two types s and t are related by `Coercible s t`, then s and t have bit-for-bit identical run-time representations.* Moreover, as you can see from the type of `coerce`, if `Coercible s t` holds then `coerce` can convert a value of type s to one of type t . And that’s it!

The crucial question, to which we devote the rest of this section and the next, becomes this: exactly when does `Coercible s t` hold? To whet your appetite consider these declarations:

```
newtype Age      = MkAge Int
newtype AgeRange = MkAR (Int,Int)
newtype BigAge   = MkBig Age
```

Here are some coercions that hold, so that a single call to `coerce` suffices to convert between the two types:

- `Coercible Int Age`: we can coerce from `Int` to `Age` at zero cost; this is simply the `MkAge` constructor.
- `Coercible Age Int`: and the reverse; this is pattern matching on `MkAge`.

GHC generates the following instances of `Coercible`:

- (1) `instance Coercible a a`
- (2) For every `newtype NT x = MkNT (T x)`, the instances

```
instance Coercible (T x) b => Coercible (NT x) b
instance Coercible a (T x) => Coercible a (NT x)
```

which are visible if and only if the constructor `MkNT` is in scope.

- (3) For every type constructor `TC r p n`, where
 - r stands for `TC`’s parameters at role representational,
 - p for those at role phantom and
 - n for those at role nominal,

the instance

```
instance Coercible r1 r2 =>
  Coercible (TC r1 p1 n) (TC r2 p2 n)
```

Figure 2. Coercible instances

- `Coercible [Age] [Int]`: lifting the coercion over lists.
- `Coercible (Either Int Age) (Either Int Int)`: lifting the coercion over `Either`.
- `Coercible (Either Int Age) (Either Age Int)`: this is more complicated, because first argument of `Either` must be coerced in one direction, and the second in the other.
- `Coercible (Int -> Age) (Age -> Int)`: all this works over function arrows too.
- `Coercible (Age, Age) AgeRange`: we have to unwrap the pair of `Ages` and then wrap with `MkAR`.
- `Coercible [BigAge] [Int]`: two levels of coercion.

In the rest of this section we will describe how `Coercible` constraints are solved or, equivalently, which instances of `Coercible` exist. (See Figure 2 for a concise summary.)

2.1 Coercing newtypes

Since `Coercible` relates a newtype with its base type, we need `Coercible` instance declarations for every such newtype. The naive `instance Coercible Int Age` does not work well, for reasons explained in the box on page 3, so instead we generate *two* instances for each newtype:

```
instance Coercible a Int => Coercible a Age   — (A1)
instance Coercible Int b => Coercible Age b   — (A2)
```

```
instance Coercible a Age => Coercible a BigAge — (B1)
instance Coercible Age b => Coercible BigAge b — (B2)
```

```
instance Coercible a AgeRange => Coercible a (Int,Int)
instance Coercible AgeRange b => Coercible (Int,Int) b
```

Notice that each instance unwraps just one layer of the newtype, so we call them the “unwrapping instances”.

If we now want to solve, say, a constraint `Coercible s Age`, for any type s , we can use (A1) to reduce it to the simpler goal `Coercible s Int`. A more complicated, two-layer coercion `Coercible BigAge Int` is readily reduced, in two such steps, to `Coercible Int Int`. All we need now is for GHC to have a built-in witness of reflexivity, expressing that any type has the same run-time representation as itself:

```
instance Coercible a a
```

This simple scheme allows coercions that involve arbitrary levels of wrapping or unwrapping, in either direction, with a single call to `coerce`. The solution path is not fully determined, but that does not matter. For example, here are two ways to solve `Coercible BigAge Age`:

```

Coercible BigAge Age
→ Coercible BigAge Int   — By (A1)
→ Coercible Age Int      — By (B2)
→ Coercible Int Int      — By (A2)
→ solved                 — By reflexivity

Coercible BigAge Age
→ Coercible Age Age      — By (B2)
→ solved                 — By reflexivity

```

Since `Coercible` constraints have no run-time behaviour (unlike normal type class constraints), we have no concerns about incoherence; any solution will do.

The newtype-unwrapping instances (i.e., (2) in Figure 2) are available *only if the corresponding newtype data constructor (Mk in our current example) is in scope*; this is required to preserve abstraction, as we explain in Section 3.1.

2.2 Coercing parameters of type constructors

As Figure 2 shows, as well as the unwrapping instances for a **newtype**, we also generate one instance for each type constructor, including data types, newtypes the function type, and built-in data types like tuples. We call this instance the “lifting instance” for the type, because it lifts coercions through the type. The shape of the instance depends on the so-called *roles* of the type constructor. Each type parameter of a type constructor has a role, determined by the way in which the parameter is used in the definition of the type constructor. In practice, the roles of a declared data type are determined by a role inference algorithm (Section 5) and can be modified by role annotations (Section 3.1). Once defined, the roles of a type constructor are the same in every scope, regardless of whether the concrete definition of that type is available in that scope.

Roles, a development of earlier work [WVPZ11] (Section 8), are a new concept for the programmer. In the following subsections, we discuss how the three possible roles, *representational*, *phantom* and *nominal*, ensure that lifting instances do not violate type safety by allowing coercions between types with different run-time representations.

2.2.1 Coercing representational type parameters

The most common role is *representational*. It is the role that is assigned to the type parameters of ordinary newtypes and data types like `Maybe`, the list type and `Either`. The `Coercible` instances for these type constructors are:

```

instance Coercible a b ⇒ Coercible (Maybe a) (Maybe b)
instance Coercible a b ⇒ Coercible [a] [b]
instance (Coercible a1 b1, Coercible a2 b2)
  ⇒ Coercible (Either a1 a2) (Either b1 b2)

```

These instances are just as you would expect: for example, the type `Maybe t1` and `Maybe t2` have the same run-time representation if and only if `t1` and `t2` have the same representation.

Most primitive type constructors also have representational roles for their arguments. For example, the domain and co-domain of arrow types are representational, giving rise to the following `Coercible` instance:

Why a single instance is not enough

Why do we create two instances for every newtype, rather than just the single declaration

```
instance Coercible Int Age
```

to witness the fact that `Int` and `Age` have the same run-time representation?

That would indeed allow us to convert from `Int` to `Age`, using `coerce`, but what about the reverse direction? We then might need a second function

```
uncoerce :: Coercible a b ⇒ b → a
```

although it would be tiresome for the programmer to remember which one to call. Alternatively, perhaps GHC should generate *two* instances:

```
instance Coercible Int Age
instance Coercible Age Int
```

But how would we get from `BigAge` to `Int`? We could try this:

```
down :: BigAge → Int
down x = coerce (coerce x)
```

Our intent here is that each invocation of `coerce` unwraps one “layer” of newtype. But this is not good, because the type inference engine cannot figure out which type to use for the result of the inner `coerce`. To make the code typecheck we would have to add a type signature:

```
down :: BigAge → Int
down x = coerce (coerce x :: Age)
```

Not very nice. Moreover we would prefer to do all this with a *single* call to `coerce`, implying that `Coercible BigAge Int` must hold. That might make us consider adding the instance declaration

```
instance (Coercible a b, Coercible b c) ⇒ Coercible a c
```

to express the transitivity of `Coercible`. But now the problem of the un-specified intermediate type `b` re-appears, and cannot be solved with a type signature.

All of these problems are nicely solved using the instances in Figure 2.

```
instance (Coercible a1 b1, Coercible a2 b2)
  ⇒ Coercible (a1 → a2) (b1 → b2)
```

Likewise, the type `IORef` has a representational parameter, so expressions of type `IORef Int` can be converted to type `IORef Age` for zero cost (and outside of the `IO` monad).

Returning to the introduction, we can use these instances to write `linesH` very directly, thus:

```
linesH :: HTML → [HTML]
linesH = coerce lines
```

In this case, the call to `coerce` gives rise to a constraint `Coercible (String → [String]) (HTML → [HTML])`, which gets simplified to `Coercible String HTML` using the instances for arrow and list types. Then the instance for the newtype `HTML` reduces it to `Coercible String String`, which is solved by the reflexive instance.

2.2.2 Coercing phantom type parameters

A type parameter has a *phantom* role if it does not occur in the definition of the type, or if it does, then only as a phantom parameter of another type constructor. For example, these declarations

```
data Phantom b = Phantom
data NestedPhantom b = L [Phantom b] | SomethingElse
```

both have parameter `b` at a phantom role.

When do the types `Phantom t1` and `Phantom t2` have the same run-time representation? Always! Therefore, we have the instances

```
instance Coercible (Phantom a) (Phantom b)
instance Coercible (NestedPhantom a) (NestedPhantom b)
```

and `coerce` can be used to change the phantom parameter arbitrarily.

2.2.3 Coercing nominal type parameters

In contrast, the *nominal* role induces the strictest preconditions for Coercible instances. This role is assigned to a parameter that possibly affects the run-time representation of a type, commonly because it is passed to a type function. For example, consider the following code

```
type family EncData a where
  EncData String = (ByteString, Encoding)
  EncData HTML = ByteString
```

```
data Encoding = ...
data EncText a = MkET (EncData a)
```

Even though we have `Coercible HTML String`, it would be wrong to derive the instance `Coercible (EncText HTML) (EncText String)`, because these two types have quite different run-time representations! Therefore, there are no instances that change a nominal parameter of a type constructor.

All parameters of a type or data *family* have nominal role, because they could be inspected by the type family instances. For similar reasons, the non-uniform parameters to GADTs are also required to be nominal.

2.2.4 Coercing multiple type parameters

A type constructor can have multiple type parameters, each at a different role. In that case, an appropriate constraint for each type parameter is used:

```
data Params r p n = Con1 (Maybe r) | Con2 (EncData n)
```

yields the instance

```
instance Coercible r1 r2
  => Coercible (Params r1 p1 n) (Params r2 p2 n)
```

This instance expresses that the representational type parameters may change if there is a Coercible instance for them; the phantom type parameters may change arbitrarily; and the nominal type parameters must stay the same.

3. Abstraction and coherence

The purpose of the `HTML` type from the introduction is to prevent accidentally mixing up unescaped strings and `HTML` fragments. Rejecting programs that make this mistake is not a matter of type safety as traditionally construed, but rather of preserving a desired abstraction.

While the previous section described how the Coercible instances ensure that uses of `coerce` are type safe, this section discusses two other properties: *abstraction* and *class coherence*.

3.1 Preserving abstraction

When the constructors of a type are in scope then we can write code semantically equivalent to `coerce` by hand (although it might be less efficient). In this situation, the use of `coerce` should definitely be allowed. However, when the constructors are not in scope, it turns out that we sometimes want the lifting instance, and sometimes we do *not* want it.

The newtype unwrapping instance is directly controlled by the visibility of the constructor and can be used if and only if this is in scope. (See Section 2.1 for how this is accomplished.) For example, since the author of module `Html` did not export `Mk`, a client does not see the unwrapping instances for `HTML`, and the abstraction is preserved.

However, we permit the use of the coercion lifting instance for a type constructor even when the data constructors are not available. For example, built-in types like `IORef` or the function type `(\rightarrow)` do not even have constructors that can be in scope. Nevertheless, coercing from `IORef HTML` to `IORef String` and from `HTML \rightarrow HTML` to `String \rightarrow String` should be allowed.

Therefore the rule for the lifting instance is that it can be used independent of the visibility of constructors. Instead, its form – what coercions it allows – is controlled by the roles of the type constructor’s parameters.

Library authors can control the roles assigned to type constructors using *role annotations*. In many cases, the role inferred by the type checker is sufficient, even for abstract types. Consider a library for non-empty lists:

```
module NonEmptyListLib( NE, singleton, ... ) where
  data NE a = MkNE [a]
  singleton :: a  $\rightarrow$  NE a
  ... etc...
```

The type must be exported abstractly; otherwise, the non-empty property can be broken by its users. Nevertheless lifting a coercion through `NE`, i.e. coercing `NE HTML` to `NE String`, should be allowed. Therefore, the role of `NE`’s parameter should be representational. In this case, the library author does not have to actively set it: As it is the most permissive type-safe role, the role inference algorithm (Section 5.2) already chooses representational.

However, sometimes library authors must restrict the usage of the lifting coercion to ensure that the invariants of their abstract types can be preserved. For example, consider the data type `Map k v`, which implements an efficient finite map from keys of type `k` to values of type `v`, using an internal representation based on a balanced tree, something like this:

```
data Map k v = Leaf | Node k v (Map k v) (Map k v)
```

It would be disastrous if the user were allowed to coerce from `(Map Age v)` to `(Map Int v)`, because a valid tree with regard to the ordering of `Age` might be completely bogus when using the ordering of `Int`.

To prevent that difficulty, the author specifies

```
type role Map nominal representational
```

As explained in Section 2.2, we now have the desirable and useful lifting instance

```
instance Coercible a b => Coercible (Map k a) (Map k b)
```

which allows the coercion from `Map k HTML` to `Map k String`.

Note that in the declaration of `Map` the parameters `k` and `v` are used in exactly the same way, so this distinction cannot be made by the compiler; it can only be specified by the programmer. However, the compiler ensures that programmer-specified role annotations cannot subvert the type system: if the annotation specifies an unsafe role, the compiler will reject the program.

3.2 Preserving class coherence

Another property of Haskell, independent of type-safety, is the coherence of type classes. There should only ever be one class instance for a particular class and type. We call this desirable property *coherence*. Without extra checks, `Coercible` could be used to create incoherence.

Consider this (non-Haskell98) data type, which reifies a `Show` instance as a value:

```
data HowToShow a where
```

```
  MkHTS :: Show a => HowToShow a
```

```
showH :: HowToShow a -> a -> String
showH MkHTS x = show x
```

Here `showH` pattern-matches on a `HowToShow` value, and uses the instance stored inside it to obtain the `show` method. If we are not careful, the following code would break the coherence of the `Show` type class:

```
instance Show HTML where
```

```
  show (Mk s) = "HTML:" ++ show s
```

```
stringShow :: HowToShow String
stringShow = MkHTS
htmlShow :: HowToShow HTML
htmlShow = MkHTS
badShow :: HowToShow HTML
badShow = coerce stringShow
```

```
λ> showH stringShow "Hello"
"Hello"
λ> showH htmlShow (Mk "Hello")
"HTML:Hello"
λ> showH badShow (Mk "Hello")
"Hello"
```

In the final example we were applying `show` to a value of type `HTML`, but the `Show` instance for `String` (coerced to `(Show HTML)`) was used.

To avoid this confusion, the parameters of a type class are all assigned a *nominal* role by default. Accordingly, the parameter of `HowToShow` is also assigned a nominal role by default, preventing the coercion between `(HowToShow HTML)` and `(HowToShow String)`.

4. Ensuring type safety: System FC with roles

Haskell is a large and complicated language. How do we know that the ideas sketched above in source-language terms are actually sound? What, precisely, do roles mean, and when precisely are two types equal? In this section we answer these questions for GHC’s small, statically-typed intermediate language, `GHC Core`. Every Haskell program is translated into `Core`, and we can typecheck `Core` to reassure ourselves that the (large, complicated) front end accepts only good programs.

Metavariables:

x	term	α, β	type	c	coercion
C	axiom	D	data type	N	newtype
F	type family	K	data constructor		
e	$::= \lambda c:\phi.e \mid e \gamma \mid e \triangleright \gamma \mid \dots$				terms
τ, σ	$::= \alpha \mid \tau_1 \tau_2 \mid \forall \alpha:\kappa.\tau \mid H \mid F(\bar{\tau})$				types
κ	$::= * \mid \kappa_1 \rightarrow \kappa_2$				kinds
H	$::= (\rightarrow) \mid (\Rightarrow) \mid (\sim_{\rho}^{\kappa}) \mid T$				type constants
T	$::= D \mid N$				algebraic data types
ϕ	$::= \tau \sim_{\rho}^{\kappa} \sigma$				proposition
γ, η	$::=$				coercions
	$\left\{ \begin{array}{l} \langle \tau \rangle \mid \langle \tau, \sigma \rangle_P \mid \mathbf{sym} \gamma \mid \gamma_1 \circ \gamma_2 \\ H(\bar{\gamma}) \mid F(\bar{\gamma}) \mid \gamma_1 \gamma_2 \mid \forall \alpha:\kappa.\gamma \\ c \mid C(\bar{c}) \\ \mathbf{nth}^i \gamma \mid \mathbf{left} \gamma \mid \mathbf{right} \gamma \mid \gamma @ \tau \\ \mathbf{sub} \gamma \end{array} \right.$				equivalence congruence assumptions decomposition sub-roling
ρ	$::= N \mid R \mid P$				roles
Γ	$::= \emptyset \mid \Gamma, \alpha:\kappa \mid \Gamma, c:\phi \mid \Gamma, x:\tau$				typing contexts
Ω	$::= \emptyset \mid \Omega, \alpha:\rho$				role contexts

Figure 3. An excerpt of the grammar of System FC

`Core` is an implementation of a calculus called `System FC`, itself an extension of the classical Girard/Reynolds `System F`. The version of `FC` that we develop in this paper derives from much prior work.² However, for clarity we give a self-contained description of the system and do not assume familiarity with previous versions.

Figure 3 gives the syntax of `System FC`. The starting point is an entirely conventional lambda calculus in the style of `System F`. We therefore elide most of the syntax of terms e , giving the typing judgement for terms in Appendix C.2. Types τ are also conventional, except that we add (saturated) type-family applications $F(\bar{\tau})$, to reflect their addition to source Haskell [CKP05, CKPM05]. Types are classified by kinds κ in the usual way; the kinding judgement $\Gamma \vdash \tau : \kappa$ on types is conventional and appears in Appendix C.2. To avoid clutter we use only monomorphic kinds, but it is easy to add kind polymorphism along the lines of [YWC⁺12], and our implementation does so.

4.1 Roles and casts

`FC`’s distinctive feature is a type-safe cast $(e \triangleright \gamma)$ (Figure 3), which uses a *coercion* γ to cast a term from one type to another. A coercion γ is a witness or proof of the equality of two types. Coercions are classified by the judgement

$$\Gamma \vdash \gamma : \tau \sim_{\rho}^{\kappa} \sigma$$

given in Figure 4, and pronounced “in type environment Γ the coercion γ witnesses that the types τ and σ both have kind κ , and are equal at role ρ ”. The notion of being “equal at role ρ ” is the important feature of this paper; it is a development of earlier work, as Section 8 describes. There are

² Several versions of `System FC` are described in published work. Some of these variants have had decorations to the `FC` name, such as `FC2` or `FC†`. We do not make these distinctions in the present work, referring instead to all of these systems – in fact, one evolving system – as “`FC`”.

precisely three roles (see Figure 3), written N, R, and P, with the following meaning:

Nominal equality, written \sim_N , is the equality that the type checker reasons about. When a Haskell programmer says that two Haskell types are the “same”, we mean that the types are nominally equal. Thus, we can say that $\text{Int} \sim_N \text{Int}$. Type families introduce new nominal equalities. So, if we have **type instance** $F \text{ Int} = \text{Bool}$, then $F \text{ Int} \sim_N \text{Bool}$.

Representational equality, written \sim_R , holds between two types that share the same run-time representation. Because all types that are nominally equal also share the same representation, nominal equality is a subset of representational equality. Continuing the example from the introduction, $\text{HTML} \sim_R \text{String}$.

Phantom equality, written \sim_P , holds between any two types, whatsoever. It may seem odd that we produce and consume proofs of this “equality”, but doing so keeps the system uniform and easier to reason about. The idea of phantom equality is new in this work, and it allows for zero-cost conversions among types with phantom parameters.

We can now give the typing judgement for type-safe cast:

$$\frac{\Gamma \vdash e : \tau_1 \quad \Gamma \vdash \gamma : \tau_1 \sim_R \tau_2}{\Gamma \vdash e \triangleright \gamma : \tau_2} \text{ TM_CAST}$$

The coercion γ must be a proof of *representational* equality, as witnessed by the R subscript to the result of the coercion typing premise. This makes good sense: we can treat an expression of one type τ_1 as an expression of some other type τ_2 if and only if those types share a representation.

4.2 Coercions

Coercions (Figure 3) and their typing rules (Figure 4) are the heart of System FC. The basic typing judgement for coercions is $\Gamma \vdash \gamma : \tau \sim_\rho^x \sigma$. When this judgement holds, it is easy to prove that τ and σ must have the same kind κ . However, kinds are not very relevant to the focus of this work, and so we often omit the kind annotation in our presentation. It can always be recovered by using the (syntax-directed) kinding judgement on types.

We can understand the typing rules in Figure 4, by thinking about the equalities that they define.

4.2.1 Nominal implies representational

If we have a proof that two types are nominally equal, then they are certainly representationally equal. This intuition is expressed by the **sub** operator, and the rule CO_SUB.

4.2.2 Equality is an equivalence relation

Equality is an equivalence relation at all three roles. Symmetry (rule CO_SYM) and transitivity (CO_TRANS) work for any role ρ . Reflexivity is more interesting: CO_REFL is a proof of nominal equality only. From this we can easily get representational reflexivity using **sub**. But what does “phantom” reflexivity mean? It is a proof term that any two types τ and σ are equal at role P, and we need a new coercion form to express that, written as $\langle \tau, \sigma \rangle_P$ (rule CO_PHANTOM).

4.2.3 Axioms for equality

Each newtype declaration, and each type-family instance, gives rise to an FC *axiom*; newtypes give rise to representa-

$$\boxed{\Gamma \vdash \gamma : \phi}$$

$$\frac{\Gamma \vdash \tau : \kappa}{\Gamma \vdash \langle \tau \rangle : \tau \sim_N \tau} \text{ CO_REFL}$$

$$\frac{\Gamma \vdash \gamma : \sigma \sim_\rho \tau}{\Gamma \vdash \mathbf{sym} \gamma : \tau \sim_\rho \sigma} \text{ CO_SYM}$$

$$\frac{\Gamma \vdash \gamma_1 : \tau_1 \sim_\rho \tau_2 \quad \Gamma \vdash \gamma_2 : \tau_2 \sim_\rho \tau_3}{\Gamma \vdash \gamma_1 \circ \gamma_2 : \tau_1 \sim_\rho \tau_3} \text{ CO_TRANS}$$

$$\frac{\overline{\Gamma \vdash \gamma : \tau \sim_\rho \sigma} \quad \bar{\rho} \text{ is a prefix of } \mathit{roles}(H) \quad \Gamma \vdash H\bar{\tau} : \kappa}{\Gamma \vdash H(\bar{\gamma}) : H\bar{\tau} \sim_R H\bar{\sigma}} \text{ CO_TYCONAPP}$$

$$\frac{\overline{\Gamma \vdash \gamma : \tau \sim_N \sigma} \quad \Gamma \vdash F(\bar{\tau}) : \kappa \quad \Gamma \vdash F(\bar{\sigma}) : \kappa}{\Gamma \vdash F(\bar{\gamma}) : F(\bar{\tau}) \sim_N F(\bar{\sigma})} \text{ CO_TYFAM}$$

$$\frac{\Gamma \vdash \gamma_1 : \tau_1 \sim_\rho \sigma_1 \quad \Gamma \vdash \gamma_2 : \tau_2 \sim_N \sigma_2 \quad \Gamma \vdash \tau_1 \tau_2 : \kappa \quad \Gamma \vdash \sigma_1 \sigma_2 : \kappa}{\Gamma \vdash \gamma_1 \gamma_2 : \tau_1 \tau_2 \sim_\rho \sigma_1 \sigma_2} \text{ CO_APP}$$

$$\frac{\Gamma, \alpha : \kappa \vdash \gamma : \tau \sim_\rho \sigma}{\Gamma \vdash \forall \alpha : \kappa. \gamma : \forall \alpha : \kappa. \tau \sim_\rho \forall \alpha : \kappa. \sigma} \text{ CO_FORALL}$$

$$\frac{\Gamma \vdash \tau : \kappa \quad \Gamma \vdash \sigma : \kappa}{\Gamma \vdash \langle \tau, \sigma \rangle_P : \tau \sim_P \sigma} \text{ CO_PHANTOM}$$

$$\frac{c : \tau \sim_\rho \sigma \in \Gamma}{\Gamma \vdash c : \tau \sim_\rho \sigma} \text{ CO_VAR}$$

$$\frac{C : [\bar{\alpha} : \kappa]. \sigma_1 \sim_\rho \sigma_2 \quad \overline{\Gamma \vdash \tau : \kappa}}{\Gamma \vdash C(\bar{\tau}) : \sigma_1[\bar{\tau}/\bar{\alpha}] \sim_\rho \sigma_2[\bar{\tau}/\bar{\alpha}]} \text{ CO_AXIOM}$$

$$\frac{\Gamma \vdash \gamma : H\bar{\tau} \sim_R H\bar{\sigma} \quad \bar{\rho} \text{ is a prefix of } \mathit{roles}(H) \quad H \text{ is not a newtype}}{\Gamma \vdash \mathbf{nth}^i \gamma : \tau_i \sim_{\rho_i} \sigma_i} \text{ CO_NTH}$$

$$\frac{\Gamma \vdash \gamma : \tau_1 \tau_2 \sim_N \sigma_1 \sigma_2 \quad \Gamma \vdash \tau_1 : \kappa \quad \Gamma \vdash \sigma_1 : \kappa}{\Gamma \vdash \mathbf{left} \gamma : \tau_1 \sim_N \sigma_1} \text{ CO_LEFT}$$

$$\frac{\Gamma \vdash \gamma : \tau_1 \tau_2 \sim_N \sigma_1 \sigma_2 \quad \Gamma \vdash \tau_2 : \kappa \quad \Gamma \vdash \sigma_2 : \kappa}{\Gamma \vdash \mathbf{right} \gamma : \tau_2 \sim_N \sigma_2} \text{ CO_RIGHT}$$

$$\frac{\Gamma \vdash \gamma : \forall \alpha : \kappa. \tau_1 \sim_\rho \forall \alpha : \kappa. \sigma_1 \quad \Gamma \vdash \tau : \kappa}{\Gamma \vdash \gamma @ \tau : \tau_1[\tau/\alpha] \sim_\rho \sigma_1[\tau/\alpha]} \text{ CO_INST}$$

$$\frac{\Gamma \vdash \gamma : \tau \sim_N \sigma}{\Gamma \vdash \mathbf{sub} \gamma : \tau \sim_R \sigma} \text{ CO_SUB}$$

Figure 4. Formation rules for coercions

```

newtype HTML = Mk String

type family F a
type instance F String = Int
type instance F HTML = Bool

data T a = MkT (F a)

```

Figure 5. Congruence and roles example code

tional axioms, and type-family instances give rise to nominal axioms.³ For example, the declarations

```

newtype HTML = Mk String
type family F [a] = Maybe a

```

produce the axioms

$$C_1 : \text{HTML} \sim_R \text{String}$$

$$C_2 : [\alpha : *]. F([\alpha]) \sim_N \text{Maybe } \alpha$$

Axiom C_1 states that `HTML` is *representationally* equal to `String` (since they are distinct types, but share a common representation), while C_2 states that $F([\sigma])$ is *nominally* equal to `Maybe σ` (meaning that the two are considered to be the same type by the type checker). In C_2 , the notation “ $[\alpha : *]$.” binds α in the types being equated. Uses of these axioms are governed by the rule `CO_AXIOM`. Axioms must always appear fully applied, and we assume that they live in a global context, separate from the local context Γ .

4.2.4 Equality can be abstracted

Just as one can abstract over types and values in System F, one can also abstract over equality proofs in FC. To this end, FC terms (Figure 3) include coercion abstraction $\lambda c:\phi.e$ and application $e \gamma$. These are the introduction and elimination forms for the coercion-abstraction arrow (\Rightarrow), just as ordinary value abstraction and application are the introduction and elimination forms for ordinary arrow (\rightarrow) (see Appendix C.2).

A coercion abstraction binds a coercion variable $c:\phi$. These variables can occur only in coercions; see the entirely conventional rule `CO_VAR`. Coercion variables can also be bound in the patterns of a `case` expression, which supports the implementation of generalised algebraic data types (GADTs).

4.2.5 Equality is congruent

Several rules witness that, ignoring roles, equality is *congruent* – for example, if $\sigma \sim_\rho \tau$ then `Maybe σ` \sim_ρ `Maybe τ` . However, the roles in these rules deserve some study, as they are the key to understanding the whole system.

Congruence of type application Before diving into the rules themselves, it is helpful to consider some examples of how we want congruence and roles to interact. Let’s consider the definitions in Figure 5. With these definitions in hand, what equalities should be derivable? (Recall the intuitive meanings of the different roles in Section 4.1.)

1. Should `Maybe HTML` \sim_R `Maybe String` hold? Yes, it should. The type parameter to `Maybe` has a representational role, so it makes sense that two `Maybe`s built out of representationally equal types should be representationally equal.

2. Should `Maybe HTML` \sim_N `Maybe String` hold? Certainly not. These two types are entirely distinct to Haskell programmers and its type checker.
3. Should `T HTML` \sim_R `T String` hold? Certainly not. We can see, by unfolding the definition for `T`, that the representations of the two types are different.
4. Should $\alpha \text{HTML} \sim_R \alpha \text{String}$ hold, for a type variable α ? It depends on the instantiation of α ! If α becomes `Maybe`, then “yes”; if α becomes `T`, then “no”. Since we may be abstracting over α , we do not know which of the two will happen, so we take the conservative stance and say that $\alpha \text{HTML} \sim_R \alpha \text{String}$ does *not* hold.

This last point is critical. The alternative is to express α ’s argument roles in its kind, but that leads to a much more complicated system; see related work in Section 8. A distinguishing feature of this paper is the substantial simplification we obtain by attributing roles only to the arguments to type constants (H , in the grammar), and not to abstracted type variables. We thereby lose a little expressiveness, but we have not found that to be a big problem in practice. See Section 8.1 for an example of an easily fixed problem case.

To support both (1) and (4) requires two coercion forms and corresponding typing rules:

- The coercion form $H(\bar{\gamma})$ has an explicit type constant at its head. This form always proves a representational equality, and it requires input coercions of the roles designated by the roles of H ’s parameters (rule `CO_TYCONAPP`). The *roles* function gives the list of roles assigned to H ’s parameters, as explained in Section 2.2. We allow $\bar{\rho}$ to be a prefix of *roles*(H) to accommodate partially-applied type constants.
- The coercion form $\gamma_1 \gamma_2$ does not have an explicit type constant, so we must use the conservative treatment of roles discussed above. Rule `CO_APP` therefore requires γ_2 to be a nominal coercion, though the role of γ_1 carries through to $\gamma_1 \gamma_2$.

What if we wish to prove a nominal equality such as `Maybe (F String)` \sim_N `Maybe Int`? We can’t use the $H(\bar{\gamma})$ form, which proves only representational equality, but we can use the $\gamma_1 \gamma_2$ form. The leftmost coercion would just be `(Maybe)`.

Congruence of type family application Rule `CO_TYFAM` proves the equality of two type-family applications. It requires nominal coercions among all the arguments. Why? Because type families can inspect their (type) arguments and branch on them. We would not want to be able to prove any equality between `F String` and `F HTML`.

Congruence of polymorphic types The rule `CO_FORALL` works for any role ρ ; polymorphism and roles do not interact.

4.2.6 Equality can be decomposed

If we have a proof of `Maybe σ` \sim_ρ `Maybe τ` , should we be able to get a proof of $\sigma \sim_\rho \tau$, by decomposing the equality? Yes, in this case, but we must be careful here as well.

Rule `CO_NTH` is almost an inverse to `CO_TYCONAPP`. The difference is that `CO_NTH` prohibits decomposing equalities among newtypes. Why? Because `nth` witnesses injectivity and newtypes are not injective! For example, consider these definitions:

```

data Phant a = MkPhant
newtype App a b = MkApp (a b)

```

³For simplicity, we are restricting ourselves to *open* type families. Closed type families [EVPW14] are readily accommodated.

Here, $roles(App) = R, N$. (The roles are inferred during compilation; see Section 5.) Yet, we can see the following chain of equalities:

App Phant Int \sim_R Phant Int \sim_R Phant Bool \sim_R App Phant Bool

By transitivity, we can derive a coercion γ witnessing

App Phant Int \sim_R App Phant Bool

If we could use \mathbf{nth}^2 on γ , we would get Int \sim_N Bool: disaster! We eliminate this possibility by preventing \mathbf{nth} on newtypes.

The rules CO_LEFT and CO_RIGHT are almost inverses to CO_APP. The difference is that both CO_LEFT and CO_RIGHT require and produce only nominal coercions. We need a new newtype to see why this must be so:

newtype EitherInt a = MkEI (Either a Int)

This definition yields an axiom showing that, for all a, EitherInt a \sim_R (Either a Int). Suppose we could apply **left** and **right** to coercions formed from this axiom. Using **left** would get us a proof of EitherInt \sim_R (Either a), which could then be used to show, say, (Either Char) \sim_R (Either Bool) and then (using **nth**) Char \sim_N Bool. Using **right** would get us a proof of a \sim_R Int, for *any* a. These are both clearly disastrous. So, we forbid using these coercion formers on representational coercions.⁴

Thankfully, polymorphism and roles play well together, and the CO_INST rule (inverse to CO_FORALL) shows quite straightforwardly that, if two polytypes are equal, then so are the instantiated types.

There is no decomposition form for type family applications: knowing that $F(\bar{\tau})$ is equal to $F(\bar{\sigma})$ tells us nothing whatsoever about the relationship between $\bar{\tau}$ and $\bar{\sigma}$.

4.3 Role attribution for type constants

In System FC we assume an unwritten global environment of top-level constants: data types, type families, axioms, and so on. For a data type H , for example, this environment will give the kind of H , the types of H 's data constructors, and the roles of H 's parameters. Clearly this global environment must be internally consistent. For example, a data constructor K must return a value of type $D \bar{\tau}$ where D is a data type; K 's type must be well-kinded, and that kind must be consistent with D 's kind.

All of this is standard except for roles. It is essential that the roles of D 's parameters, $roles(D)$, are consistent with D 's definition. For example, it would be utterly wrong for the global environment to claim that $roles(Maybe) = P$, because then we could prove that MaybeInt \sim_R MaybeBool using CO_TYCONAPP.

We use the judgement $\bar{\rho} \models H$, to mean " $\bar{\rho}$ are suitable roles for the parameters of H ", and in our proof of type safety, we assume that $roles(H) \models H$ for all H . The rules for this judgement and two auxiliary judgements appear in Figure 6. Note that this judgement defines a *relation* between roles and data types. Our role inference algorithm (Section 5) determines the most permissible roles for this relation, but

⁴We note in passing that the forms **left** and **right** are present merely to increase expressivity. They are not needed anywhere in the metatheory to prove type soundness. Though originally part of FC, they were omitted in previous versions [WVPZ11] and even in the implementation. Haskell users then found that some desirable program were no longer type-checking. Thus, these forms were re-introduced.

$$\begin{array}{c}
\boxed{\bar{\rho} \models H} \quad \text{"}\bar{\rho} \text{ are appropriate roles for } H\text{."} \\
\forall \bar{\alpha}, \bar{\beta}, \bar{\sigma} \text{ s.t. } K : \forall \bar{\alpha}:\bar{\kappa}. \forall \bar{\beta}:\bar{\kappa}'. \bar{\phi} \Rightarrow \bar{\sigma} \rightarrow D \bar{\alpha} : \\
\forall \tau \text{ s.t. } \tau \in \bar{\sigma} \vee \tau \in \bar{\phi} : \\
\frac{\bar{\alpha}:\bar{\rho}, \bar{\beta}:\bar{N} \vdash \tau : R}{\bar{\rho} \models D} \quad \text{ROLES_DATA} \\
C : [\bar{\alpha}:\bar{\kappa}]. N \bar{\alpha} \sim_R \sigma \quad \frac{\bar{\alpha}:\bar{\rho} \vdash \sigma : R}{\bar{\rho} \models N} \quad \text{ROLES_NEWTYPER} \\
\frac{}{R, R \models (\rightarrow)} \quad \frac{}{R, R \models (\Rightarrow)} \quad \frac{}{\rho, \rho \models (\sim_\rho)} \\
\boxed{\Omega \vdash \tau : \rho} \quad \text{"Assuming } \Omega, \tau \text{ can be used at role } \rho\text{."} \\
\frac{\alpha:\rho' \in \Omega \quad \rho' \leq \rho}{\Omega \vdash \alpha : \rho} \quad \text{RTY_VAR} \\
\frac{\bar{\rho} \text{ is a prefix of } roles(H)}{\Omega \vdash \tau : \rho} \quad \text{RTY_TYCONAPP} \\
\frac{}{\Omega \vdash H \bar{\tau} : R} \quad \text{RTY_TYCON} \\
\frac{}{\Omega \vdash H : N} \quad \text{RTY_TYCON} \\
\frac{\Omega \vdash \tau : \rho \quad \Omega \vdash \sigma : N}{\Omega \vdash \tau \sigma : \rho} \quad \text{RTY_APP} \\
\frac{\Omega, \alpha:N \vdash \tau : \rho}{\Omega \vdash \forall \alpha:\kappa. \tau : \rho} \quad \text{RTY_FORALL} \\
\frac{}{\Omega \vdash \tau : N} \quad \text{RTY_TYFAM} \\
\frac{}{\Omega \vdash F(\bar{\tau}) : \rho} \quad \text{RTY_TYFAM} \\
\frac{}{\Omega \vdash \tau : P} \quad \text{RTY_PHANTOM} \\
\boxed{\rho_1 \leq \rho_2} \quad \text{"}\rho_1 \text{ is a sub-role of } \rho_2\text{."} \\
\frac{}{N \leq \rho} \quad \frac{}{\rho \leq P} \quad \frac{}{\rho \leq \rho}
\end{array}$$

Figure 6. Rules asserting a correct assignment of roles to data types

often other, less permissive roles, such as those specified by role annotations, are also included by this relation.

Start with ROLES_NEWTYPER. Recall that a newtype declaration for N gives rise to an axiom $C : [\bar{\alpha}:\bar{\kappa}]. N \bar{\alpha} \sim_R \sigma$. The rule says that roles $\bar{\rho}$ are acceptable for N if each parameter α_i is used in σ in a way consistent with ρ_i , expressed using the auxiliary judgement $\bar{\alpha}:\bar{\rho} \vdash \sigma : R$.

The key auxiliary judgement $\Omega \vdash \tau : \rho$ checks that the type variables in τ are used in a way consistent with their roles specified in Ω , when considered at role ρ . More precisely, if $\alpha:\rho' \in \Omega$ and if $\sigma_1 \sim_{\rho'} \sigma_2$ then $\tau[\sigma_1/\alpha] \sim_\rho \tau[\sigma_2/\alpha]$. Unlike in many typing judgements, the role ρ (as well as Ω) is an *input* to this judgement, not an output. With this in mind, the rules for the auxiliary judgement are straightforward. For example, RTY_TYFAM says that the argument types of a type family application are used at nominal role. The variable rule, RTY_VAR, allows a variable to be assigned a more restrictive role (via the sub-role judgement) than required, which is needed both for multiple occurrences of the same variable, and to account for role signatures. Note that rules RTY_TYCONAPP and RTY_APP overlap – this judgement is not syntax-directed.

Returning to our original judgement $\bar{\rho} \models H$, `ROLES_DATA` deals with algebraic data types D , by checking roles in each of its data constructors K . The type of a constructor is parameterised by universal type variables $\bar{\alpha}$, existential type variables $\bar{\beta}$, coercions (with types $\bar{\phi}$), and term-level arguments (with types $\bar{\sigma}$). For each constructor, we must examine each proposition ϕ and each term-level argument type σ , checking to make sure that each is used at a representational role. Why check for a representational role specifically? Because *roles* is used in `CO_TYCONAPP`, which produces a representational coercion. In other words, we must make sure that each term-level argument appears at a representational role within the type of each constructor K for `CO_TYCONAPP` to be sound.

Finally (\rightarrow) and (\Rightarrow) have representational roles: functions care about representational equality but never branch on the nominal identity of a type. (For example, functions always treat `HTML` and `String` identically.) We also see that the roles of the arguments to an equality proposition match the role of the proposition. This fact comes from the congruence of the respective equality relations.

These definitions lead to a powerful theorem:

Theorem (Roles assignments are flexible). *If $\bar{\rho} \models H$, where H is a data type or newtype, and $\bar{\rho}'$ is such that $\rho'_i \leq \rho_i$ (for $\rho_i \in \bar{\rho}$ and $\rho'_i \in \bar{\rho}'$), then $\bar{\rho}' \models H$.*

Proof. Straightforward induction on $\Omega \vdash \tau : \rho$. □

This theorem states that, given a sound role assignment for H , any more restrictive role assignment is also sound. This property of our system here is one of its distinguishing characteristics from our prior work on roles – see Section 10 for discussion.

4.4 Metatheory

The preceding discussion gave several non-obvious examples where admitting *too many* coercions would lead to unsoundness. However, we must have *enough* coercions to allow us to make progress when evaluating a program. (We do not have space to elaborate, but a key example is the use of `nth` in rule `S_KPUSH`, presented in Appendix C.3.) Happily, we can be confident that we have enough coercions, but not too many, because we prove the usual progress and preservation theorems for System FC. The structure of the proofs follows broadly that in previous work, such as [WVPZ11] or [YWC+12].

A key step in the proof of progress is to prove *consistency*; that is, that no coercion can exist between, say, `Int` and `Bool`. This is done by defining a non-deterministic, role-directed rewrite relation on types and showing that the rewrite system is confluent and preserves type constants (other than newtypes) appearing in the heads of types. We then prove that, if a coercion exists between two types τ_1 and τ_2 , these two types both rewrite to a type σ . We conclude then that τ_1 and τ_2 , if headed by a non-newtype type constant, must be headed by the same such constant.

Alas, the rewrite relation is *not* confluent! The non-linear patterns allowed in type families (that is, with a repeated variable on the left-hand side), combined with non-termination, break the confluence property (previous work gives full details [EVPW14]). However, losing confluence does not necessarily threaten consistency – it just threatens the particular proof technique we use. However, a more powerful proof appears to be an open problem in the term rewriting commu-

nity.⁵ For the purposes of our proof we dodge this difficulty by restricting type families to have only linear patterns, thus leading to confluence; consistency of the full system remains an open problem.

The full proof of type safety appears in the appendix; it exhibits no new proof techniques.

5. Roles on type constructors

In System FC we assume that, for every type constant H , the global environment specifies $roles(H)$, the roles of H 's parameters. However, there is some flexibility about this role assignment; the only requirement for type soundness is that $roles(H) \models H$.

In GHC, the roles of a type constructor are determined first by any role annotations provided by the programmer. If these are missing, the type checker calculates the default roles using the inference algorithm described below.

5.1 Role inference

A type constructor's roles are assigned depending on its nature:

- Primitive type constructors like (\rightarrow) and (\sim_{ρ}^X) have pre-defined roles (Figure 6).
- Type families (Section 2.2.3) and type classes (Section 3.2) have nominal roles for all parameters.
- For a **data** type or **newtype** T GHC *infers* the roles for T 's type parameters, possibly modified by role annotations (Section 3.1).

The role inference algorithm is quite straightforward. At a high level, it simply starts with the role information of the built-in constants (\rightarrow) , (\Rightarrow) , and (\sim_{ρ}) , and propagates the roles until it finds a fixpoint. In the description of the algorithm, we assume a mutable environment; $roles(H)$ pulls a list of roles from this environment. Only after the algorithm is complete will $roles(H) \models H$ hold.

1. Populate $roles(T)$ (for all T) with user-supplied annotations; omitted role annotations default to phantom. (See Section 5.2 for discussion about this choice of default.)
2. For every data type D , every constructor for that data type K , and every coercion type and term-level argument type σ to that constructor: run $walk(D, \sigma)$.
3. For every newtype N with representation type σ , run $walk(N, \sigma)$.
4. If the role of any parameter to any type constant changed in the previous steps, go to step 2.
5. For every T , check $roles(T)$ against a user-supplied annotation, if any. If these disagree, reject the program. Otherwise, $roles(T) \models T$ holds.

The procedure $walk(T, \sigma)$ is defined as follows, matching from top to bottom:

⁵Specifically, we believe that a positive answer to open problem #79 of the Rewriting Techniques and Applications (RTA) conference would lead to a proof of consistency; see <http://www.win.tue.nl/rtaloop/problems/79.html>.

$\text{walk}(T, \alpha) \quad := \text{mark the } \alpha \text{ parameter to } T \text{ as } R.$
 $\text{walk}(T, H \bar{\tau}) \quad := \text{let } \bar{\rho} = \text{roles}(H);$
 for every $i, 0 < i \leq \text{length}(\bar{\tau})$:
 if $\rho_i = N$, then
 mark all variables free in τ_i as N ;
 else if $\rho_i = R$, then $\text{walk}(T, \tau_i).$
 $\text{walk}(T, \tau_1 \tau_2) \quad := \text{walk}(T, \tau_1);$
 mark all variables free in τ_2 as $N.$
 $\text{walk}(T, F(\bar{\tau})) \quad := \text{mark all variables free in the } \bar{\tau} \text{ as } N.$
 $\text{walk}(T, \forall \beta:\kappa.\tau) := \text{walk}(T, \tau).$

When marking, we must follow these two rules:

1. If a variable to be marked does not appear as a type-level argument to the data type T in question, ignore it.
2. Never allow a variable previously marked N to be marked R . If such a mark is requested, ignore it.

The first rule above deals with existential and local (\forall -bound) type variables, and the second one deals with the case where a variable is used both in a nominal and in a representational context. In this case, we wish the variable to be marked N , not R .

Theorem. *The role inference algorithm always terminates.*

Theorem (Role inference is sound). *After running the role inference algorithm, $\text{roles}(H) \models H$ will hold for all H .*

Theorem (Role inference is optimal). *After running the role inference algorithm, any loosening of roles (a change from ρ to ρ' , where $\rho \leq \rho'$ and $\rho \neq \rho'$) would violate $\text{roles}(H) \models H$.*

Proofs of these theorems appear in Appendix I.

5.2 The role of role inference

According to the specification of sound role assignments in Figure 6, a type constructor H can potentially have several different sound role assignments. For example, assigning `Maybe`'s parameter to have a representational role is type-safe, but assigning a nominal role would be, too. Note that nominal roles are always sound for data types, according to the definition in Figure 6. However, as we saw in the description of the role inference algorithm, we choose default roles for data types to be as permissive as possible – in other words, the default role for a data type constructor parameter starts at phantom and only change when constrained by the algorithm. Here, we discuss this design decision and its consequences.

What if we had no role inference whatsoever and required programmers to annotate every data type? In this case, the burden on programmers seems drastic and migration to this system overwhelming, requiring all existing data type declarations to be annotated with roles.

Alternatively, we could specify that all unannotated roles default to nominal (thus removing the need for role inference). This choice would lead to greater abstraction safety by default – we would not have to worry that the implementor of `Map` is unaware of roles and forgets a critical role annotation.

However, we choose to use the most permissive roles by default for several reasons. First, for convenience: this choice increases the availability of `coerce` (as only those types with annotations would be `Coercible` otherwise), and it supports backward compatibility with the Generalized Newtype Deriving (GND) feature (see Section 7).

Furthermore, our choice of using phantom as the default also means that the majority of programmers do not need

to learn about roles. They will not need role annotations in their code. Users of `coerce` will need to consider roles, as will library implementors who use class-based invariants (see Section 3.1). Other users are unaffected by roles and will not be burdened by them.

Our choices in the design of the role system, and the default of phantom in particular, has generated vigorous debate.⁶ This discussion is healthy for the Haskell community. The difficulty with abstraction is not new: with GND, it has always been possible to lift coercions through data types, potentially violating their class-based invariants. The features described in this paper make this subversion both more convenient (through the use of `coerce`) and, more importantly, now preventable (through the use of role annotations).

6. Implementing Coercible

We have described the source-language view of `Coercible` (Sections 2, 3), and System FC, the intermediate language into which the source language is elaborated (Section 4). In this section we link the two by describing how the source-language use of `Coercible` is translated into Core.

6.1 Coercible and coerce

When the compiler transforms Haskell to Core, type classes become ordinary types and type class constraints turn into ordinary value arguments [WB89]. In particular, type classes typically become simple product types with one field per method.

The same holds for the type class `Coercible a b`, which has one method, namely the witness of representational equality $a \sim_R b$. As that type cannot be expressed in Haskell, the actual definition of `Coercible` is built in:

```
data Coercible a b = MkCoercible (a ~R b)
```

The definition of `coerce`, which is also only possible in Core, pattern-matches on `MkCoercible` to get hold of the equality witness, and then uses Core's primitive cast operation:

```
coerce :: forall α β. Coercible α β → α → β  
coerce = Λ α β. λ (c :: Coercible α β) (x :: α). case c of  
  MkCoercible eq → x ▷ eq
```

Since type applications are explicit in Core, `coerce` now takes four arguments: the types to cast from and to, the coercion witness, and finally the value to cast.

The data type `Coercible` also serves to `box` the primitive, unboxed type \sim_R , just as `Int` serves to `box` the primitive, unboxed type `Int#`:

```
data Int = Int# Int#
```

All boxed types are represented uniformly by a heap pointer. In GHC all constraints (such as `Eq a` or `Coercible a b`) are boxed, so that they can be treated uniformly, and even polymorphically [YWC⁺12]. In contrast, an unboxed type is represented by a non-pointer bit field, such as a 32 or 64-bit int in the case of `Int#` [PL91].

A witness of (unboxed) type \sim_R carries no information: we never actually inspect an equality proof at run-time. So the type \sim_R can be represented by a zero-width bit-field – that is, by nothing at all. This implementation trick, of boxing a zero-bit witness, is exactly analogous to the wrapping of boxed nominal equalities used to implement deferred type errors [VPMa12].

⁶To read some of this debate, see the thread beginning with this post: <http://www.haskell.org/pipermail/libraries/2014-March/022321.html>

Since Coercible is a regular data type, you might worry about bogus programs like this, which uses recursion to construct an unsound witness `co` whose value is bottom:

```
looksUnsound :: forall α β. α → β
looksUnsound = \α β x →
  let co :: Coercible α β = co in
  coerce α β co x
```

However, since `coerce` evaluates the `Coercible` argument (see the definition of `coerce` above), `looksUnsound` will simply diverge. Again, this follows the behaviour of deferred type errors [VPMa12].

In uses of `coerce`, the `Coercible` argument will be constructed from the instances which, as described below (Section 6.4), are guaranteed to be acyclic. The usual simplification machinery of GHC then ensures that these are inlined, causing the `case` to cancel with the `MkCoercible` constructor, leaving only the cast `x ▷ eq`, which is operationally free.

6.2 On-demand instance generation

The language of Section 2 suggests that we generate Haskell instance declarations for `Coercible`, based on type declarations. Although this is a useful way to explain the design to a programmer (who is already familiar with type classes and instance declarations), GHC's implementation is much simpler and more direct.

Rather than generate and compile instance declarations, the constraint solver treats `Coercible` constraints specially: to solve a `Coercible` constraint, the solver uses the rules of Section 2 directly to decompose the constraint into simpler sub-goals. This approach makes it easy to implement the non-standard visibility rules of `Coercible` instances (see Section 3.1), by simply not applying the newtype-unwrapping rule if the constructor is not in scope.

6.3 The higher rank instance

Consider this declaration, whose constructor uses a higher-rank type:

```
newtype Sel = MkSel (forall a. [a] → a)
```

We would expect its newtype-unwrapping instance to take the form

```
instance Coercible (forall a. [a] → a) b ⇒ Coercible Sel b
instance Coercible a (forall a. [a] → a) ⇒ Coercible a Sel
```

These declarations are illegal in source Haskell, even with all GHC extensions enabled. Nevertheless, we can generate internally and work with them in the solver just fine. This leads to constraints of the form

```
Coercible (forall a. s) (forall b. t)
```

which need special support in the solver. It already supports solving (nominal) type equalities of the form `(forall a. s) ~ (forall b. t)`, by generating a fresh type variable `c` and solving `s[c/a] ~ t[c/b]`. We generalised this functionality to handle representational type equalities as well.

6.4 Preventing circular reasoning and diverging instances

For most type classes, like `Show`, it is perfectly fine (and useful) to use a not-yet solved type class constraint to solve another, even though this can lead to cycles [LP05]. Consider the following code and execution:

```
newtype Fix a = MkFix (a (Fix a))
deriving instance Show (a (Fix a)) ⇒ Show (Fix a)
```

```
λ> show (MkFix (Just (MkFix (Just (MkFix Nothing))))))
"MkFix (Just (MkFix (Just (MkFix Nothing))))"
```

There are two `Show` instances at work: one for `Show (Maybe a)`, which uses the instance of `Show a`; and one for `Show (Fix a)`, which uses the the instance `Show (a (Fix a))`. Plugging them together to solve `Show (Fix Maybe)`, we see that this instance calls, by way of `Show (Maybe (Fix Maybe))`, itself. Nevertheless, the result is perfectly well-behaved and indeed terminates.

But with `Coercible`, such circular reasoning would be problematic; we could then seemingly write the bogus function `looksUnsoundH`:

```
newtype Id a = MkId a
c1 :: a → Fix Id
c1 = coerce
c2 :: Fix Id → b
c2 = coerce
looksUnsoundH :: a → b
looksUnsoundH = c2 ◦ c1
```

With the usual constraint solving, this code would type check: to solve the constraint `Coercible a (Fix Id)`, we need to solve `Coercible a (Id (Fix Id))`, which requires `Coercible a (Fix Id)`. This is a constraint we already looked at, so the constraint solver would normally consider all required constraints solved and accept the program.

Fortunately, there is no soundness problem here. Circular constraint-solving leads to a recursive definition of the `Coercible` constraints, exactly like the (Core) `looksUnsound` in Section 6.1, and `looksUnsoundH` will diverge just like `looksUnsound`. Nevertheless, unlike normal type classes, a recursive definition of `Coercible` is *never* useful, so it is more helpful to reject it statically. GHC therefore uses the existing depth-counter of the solver to spot and reject recursion of `Coercible` constraints.

6.5 Coercible and rewrite rules

What if a client of module `Html` writes this?

```
...(map unMk hs)...
```

She cannot use `coerce` because `HTML` is an abstract type, so the type system would (rightly) reject an attempt to use `coerce` (Section 3.1). However, since `HTML` is a newtype, one might hope that GHC's optimiser would transform `(map unMk)` to `coerce`. The optimiser must respect type soundness, but (by design) it does not respect abstraction boundaries: dissolving abstractions is one key to high performance.

The correctness of transforming `(map unMk)` to `coerce` depends on a theorem about `map`, which a compiler can hardly be expected to identify and prove all by itself. Fortunately GHC already comes with a mechanism that allows a library author to specify *rewrite rules* for their code [PTH01]. The author takes the proof obligation that the rewrite is semantics-preserving, while GHC simply applies the rewrite whenever possible. In this case the programmer could write

```
{-# RULES "map/co" map coerce = coerce #-}
```

In our example, the programmer wrote `(map unMk)`. The definition `unMk` in module `Html` does not mention `coerce`, but both produce the same System FC code (a cast). So via cross-module inlining (more dissolution of abstraction boundaries)

unMk will be inlined, transforming the call to the equivalent of (map coerce), and that in turn fires the rewrite rule. Indeed even a nested call like map (map unMk) will also be turned into a single call of coerce by this same process applied twice.

The bottom line is this: the author of a map-like function someMap can accompany someMap with a RULE, and thereby optimise calls of someMap that do nothing into a simple call to coerce.

Could we dispense with a user-visible coerce function altogether, instead using map-like functions and RULEs as above? No: doing so would replace the zero-cost guarantee with best-effort optimisation; it would burden the author of every map-like function with the obligation to write a suitable RULE; it would be much less convenient to use in deeply-nested cases; and there might simply *be* no suitable map-like function available.

7. Generalized Newtype Deriving done right

As mentioned before, **newtype** is a great tool to make programs more likely to be correct, by having the type checker enforce certain invariants or abstractions. But newtypes can also lead to tedious boilerplate. Assume the programmer needs an instance of the type class Monoid for her type HTML. The underlying type String already comes with a suitable instance for Monoid. Nevertheless, she has to write quite a bit of code to convert that instance into one for HTML:

```
instance Monoid HTML where
  mempty = Mk mempty
  mappend (Mk a) (Mk b) = Mk (mappend a b)
  mconcat xs = Mk (mconcat (map unMk xs))
```

Note that this definition is not only verbose, but also non-trivial, as invocations of Mk and unMk have to be put in the right places, possibly via some higher order functions like map – all just to say “just use the underlying instance”!

This task is greatly simplified with Coercible: Instead of wrapping and unwrapping arguments and results, she can directly coerce the method of the base type’s instance itself:

```
instance Monoid HTML where
  mempty = coerce (mempty :: String)
  mappend = coerce (mappend :: String → String → String)
  mconcat = coerce (mconcat :: [String] → String)
```

The code is pure boilerplate: apply coerce to the method, instantiated at the base type by a type signature. And because it is boilerplate, the compiler can do it for her; all she has to do is to declare which instances of the base type should be lifted to the new type by listing them in the **deriving** clause:

```
newtype HTML = Mk String deriving Monoid
```

This is not a new feature: GHC has provided this *Generalized Newtype Deriving* (GND) for many years. But, the implementation was “magic” – GND would produce code that a user could not write herself. Now, the feature can be explained easily and fully via coerce.

Furthermore, GND was previously unsound [WVPZ11]. When combined with other extensions of GHC, such as type families [CKP05, CKPM05] or GADTs [CH03], GND could be exploited to completely break the type system: Figure 7 shows how this notorious bug can allow any type to be coerced to any other. The clause “**deriving** (UnsafeCast b)” is the bogus use of GND, and now will generate the instance

```
instance UnsafeCast b c ⇒ UnsafeCast b (Id2 c) where
  unsafe = coerce (unsafe :: c → Discern c b)
```

```
newtype Id1 a = MkId1 a
newtype Id2 a = MkId2 (Id1 a) deriving (UnsafeCast b)
```

```
type family Discern a b
type instance Discern (Id1 a) b = a
type instance Discern (Id2 a) b = b
```

```
class UnsafeCast to from where
  unsafe :: from → Discern from to
```

```
instance UnsafeCast b (Id1 a) where
  unsafe (MkId1 x) = x
```

```
unsafeCoerce :: a → b
unsafeCoerce x = unsafe (MkId2 (MkId1 x))
```

Figure 7. The above implementation of unsafeCoerce compiles (with appropriate flags) in GHC 7.6.3 but does not in GHC 7.8.1.

which will rightly be rejected because Discern’s first parameter has a nominal role. Indeed, preventing abuse of GND was the entire subject of the previous work [WVPZ11] the current paper is based on.

Similarly, it was possible to use GND to break invariants of abstract data types. The addition of coerce makes it yet easier to break such abstractions. As discussed in Section 3.1, these abuses can now be prevented via role annotations.

8. Related work

Prior work discusses the relationship between roles in FC and languages with generativity and abstraction, type-indexed constructs, and universes in dependent type theory. We do not repeat that discussion here. Instead we use this section to clarify the relationship between this paper and [WVPZ11], as well as make connections to other systems.

8.1 Prior version of roles

The idea of *roles* was initially developed in [WVPZ11] as a solution to the Generalized Newtype Deriving problem. That work introduces the equality relations \sim_R and \sim_N (called “type equality” and “code equality” resp. in [WVPZ11]). However, the system presented in [WVPZ11] was quite invasive: it required annotating every sub-tree of every kind with a role. Kinds in GHC are already quite complicated because of kind polymorphism, and a new form of role-annotated kinds would be more complex still.

In this paper, we present a substantially simplified version of the roles system of [WVPZ11], requiring role information only on the parameters to data types. Our new design keeps roles and kinds modularly separate, so that roles can be handled almost entirely separately (both intellectually and in the implementation) from kinds. The key simplification is to “assume the worst” about higher-kinded parameters, by assuming that their arguments are all nominal. In exchange we give up some expressiveness; specifically, we give up the ability to abstract over type constructors with non-nominal argument roles (see Section 10).

Furthermore, the observation that it is sound to “assume the worst” and use parameterised types with less permissive roles opens the door to role annotations. In this work, programmers are allowed to deliberately specify less permissive roles, giving them the ability to preserve type abstractions.

Surprisingly, this flexibility means that our version of roles actually *increases* expressiveness compared to [WVPZ11] in some places. In [WVPZ11] a role is part of a type’s kind, so a type expecting a higher-kinded argument (such as `Monad`) would also have to specify the roles expected by its argument. Therefore if `Monad` is applicable to `Maybe`, it would not also be applicable to a type `T` whose parameter has a nominal role. In the current work, however, there is no problem because `Maybe` and `T` have the same kind.

Besides the simplification discussed above, this paper makes two other changes to the specification of roles presented in [WVPZ11].

- The treatment of the phantom role is entirely novel; the rule `CO_PHANTOM` has no analogue in prior work.
- The coercion formation rules (Figure 4) are refactored so that the role on the coercion is an *output* of the (syntax-directed) judgement instead of an input. This is motivated by the implementation (which does not know the role at which coercions should be checked) and requires the addition of the `CO_SUB` rule.

There are, of course, other minor differences between this system and [WVPZ11] in keeping with the evolution of System FC. The main significant change, unrelated to roles, is the re-introduction of **left** and **right** coercions; see Section 4.2.6.

One important non-difference relates to the linear-pattern requirement. Section 4.4 describes that our language is restricted to have only *linear* patterns in its type families. (GHC, on the other hand, allows non-linear patterns as well.) This restriction exists in the language in [WVPZ11] as well. Section 4.2.2 of [WVPZ11] defines so-called Good contexts as having certain properties. Condition 1 in this definition subtly implies that all type families have linear patterns – if a type family had a non-linear pattern, it would be impossible, in general, to establish this condition. The fact that the definition of Good implies linear patterns came as a surprise, further explored in [EVPW14]. The language described in the present paper clarifies this restriction, but it is not a new restriction.

Finally, because this system has been implemented in GHC, this paper discusses more details related to compilation from source Haskell. In particular, the role inference algorithm of Section 5 is a new contribution of this work.

8.2 OCaml and variance annotations

The interactions between sub-typing, type abstraction, and various type system extensions such as GADTs and parameter constraints also appear in the OCaml language. In that context, *variance annotations* act like roles; they ensure that subtype coercions between compatible types are safe. For example, the type α `list` of immutable lists is covariant in the parameter α : if $\sigma \leq \tau$ then σ `list` \leq τ `list`. Variances form a lattice, with *invariant*, the most restrictive, at the bottom; *covariant* and *contravariant* incomparable; and *bivariant* at the top, allowing sub-typing in both directions. It is tempting to identify invariant with nominal and bivariant with phantom, but the exact connection is unclear. Scherer and Rémy [SR13] show that GADT parameters are not always invariant.

Exploration of the interactions between type abstraction, GADTs, and other features have recently revealed a soundness issue in OCaml⁷ that has been confirmed to date back several years. Garrigue discusses these issues [Gar13]. His proposed solution is to “assume that nothing is known about

⁷<http://caml.inria.fr/mantis/view.php?id=5985>

abstract types when they are used in parameter constraints and GADT return types” – akin to assigning nominal roles. However, this solution is too conservative, and in practice the OCaml 4.01 compiler relies on no fewer than *six* flags to describe the variance of type parameters. However, lacking anything equivalent to Core and its tractable metatheory, the OCaml developers cannot demonstrate the soundness of their solution in the way that we have done here.

What is clear, however, is that generative type abstraction interacts in interesting and non-trivial ways with type equality and sub-typing. Roles and type-safe coercion solve an immediate practical problem in Haskell, but we believe that the ideas have broader applicability in advanced type systems.

9. Roles in Practice

We have described a mechanism to allow safe coercions among distinct types, and we have reimplemented GHC’s previously unsafe `GeneralizedNewtypeDeriving` extension in terms of these safe coercions. Naturally, this change causes some code that was previously accepted to be rejected. Given that Haskell has a large user base and a good deal of production code, how does this change affect the community?

Advance testing During the development of this feature, we tested it against several popular Haskell packages available through Hackage, an online Haskell open-source distribution site. These tests were all encouraging and did not find any instances of hard-to-repair code in the wild.

Compiling all of Hackage As of 30 September 2013, 3,234 packages on Hackage compiled with GHC 7.6.3, the last released version without roles. The development version of GHC at that time included roles. A total of only four packages failed to compile directly due to GND failure.⁸ Of these, three of the failures were legitimate – the use of GND was indeed unsafe. For example, one case involved coercing a type variable passed into a type family; the author implicitly assumed that a newtype and its representation type were always considered equivalent with respect to the type family. Only one package failed to compile because of the gap in expressiveness between the roles in [WVPZ11] and those here. No other Hackage package depends on this one, indicating it is not a key part of the Haskell open-source fabric. See Section 10 for discussion of the failure.

These data were gathered almost two months after the implementation of roles was pushed into the development version of GHC, so active maintainers may have made changes to their packages before the study took place. Indeed, we are aware of a few packages that needed manual updates. In these cases, instances previously derived using GND had to be written by hand, but quite straightforwardly.

10. Future directions

As of the date of writing (May 2014), roles seem not to have caused an undue burden to the community. The first release candidate for GHC 7.8 was released on 3 February 2014, followed by the full release on 9 April, and package authors have been updating their work to be compatible for some time. The authors of this paper are unaware of any major

⁸These data come from Bryan O’Sullivan’s work, described here: <http://www.haskell.org/pipermail/ghc-devs/2013-September/002693.html> That posting includes 3 additional GND failures; these were due to an implementation bug, since fixed.

problems that Haskellers have had in updating existing code, despite hundreds of packages being available for GHC 7.8.⁹

However, we are aware that some users wish to use roles in higher-order scenarios that are currently impossible. We focus on one such scenario, as it is representative of all examples we have seen, including the package that did not compile when testing all of Hackage (Section 9).

Imagine adding the `join` method to the `Monad` class, as follows:

```
class Monad m where
  ...
  join :: forall a. m (m a) -> m a
```

With this definition, GND would still work in many cases. For example, if we define

```
newtype M a = Mk (Maybe a)
  deriving Monad
```

GND will work without a problem. We would need to show `Coercible (Maybe (Maybe a) -> Maybe a) (M (M a) -> M a)`, which is straightforward.

More complicated constructions run into trouble, though. Take this definition, written to restrict a monad’s interface:

```
newtype Restr m a = Mk (m a)
  deriving Monad
```

To perform GND in this scenario, we must prove `Coercible (m (m a) -> m a) (Restr m (Restr m a) -> Restr m a)`. In solving for this constraint, we eventually simplify to `Coercible (m (m a)) (m (Restr m a))`. At this point, we are stuck, because we do not have any information about the role of `m`’s parameter, so we must assume it is nominal. The GND feature is thus not available here. Similar problems arise when trying to use GND on monad transformers, a relatively common idiom.

How would this scenario play out under the system proposed in [WVPZ11]? This particular problem wouldn’t exist – `m`’s kind could have the right roles – but a different problem would. A type’s kind also stores its roles in [WVPZ11]. This means that `Monad` instances could be defined only for types that expect a representational parameter. Yet, it is sometimes convenient to define a `Monad` instance for a data type whose parameter is properly assigned a nominal role. The fact that the system described in this paper can accept `Monad` instances both for types with representational parameters and nominal parameters is a direct consequence of the *Role assignments are flexible* theorem (Section 4.3), which does not hold of the system in [WVPZ11].

Looking forward, there is a proposal to indeed add `join` to `Monad`, and so we want to be able to allow the use of GND on this enhanced `Monad` class. We have started to formulate solutions to this problem and have hope that we can overcome this barrier without modifications to the core language.

11. Conclusion

Our focus has been on Haskell, for the sake of concreteness, but we believe that this work is important beyond the Haskell community. Any language that offers *both* generative type abstraction *and* type-level computation must deal with

their interaction, and those interactions are extremely subtle. We have described one sound and tractable way to combine the two, including the source language changes, type inference, core calculus, and metatheory. In doing so we have given a concrete foundation for others to build upon.

Acknowledgments

Thanks to Antal Spector-Zabusky for contributing to this version of FC; and to Edward Kmett and Dimitrios Vytiniotis for discussion and feedback. This material is based upon work supported by the National Science Foundation under grant nos. CCF-1116620 and CCF-1319880. The first author was supported by the Deutsche Telekom Stiftung.

References

- [BEPW14] Joachim Breitner, Richard A. Eisenberg, Simon Peyton Jones, and Stephanie Weirich, *Safe zero-cost coercions for Haskell*, ICFP, ACM, 2014.
- [CH03] James Cheney and Ralf Hinze, *First-class phantom types*, Tech. report, Cornell University, 2003.
- [CKP05] Manuel M. T. Chakravarty, Gabriele Keller, and Simon Peyton Jones, *Associated type synonyms*, ICFP, ACM, 2005, pp. 241–253.
- [CKPM05] Manuel M. T. Chakravarty, Gabriele Keller, Simon Peyton Jones, and Simon Marlow, *Associated types with class*, POPL, ACM, 2005, pp. 1–13.
- [EVPW14] Richard A. Eisenberg, Dimitrios Vytiniotis, Simon Peyton Jones, and Stephanie Weirich, *Closed type families with overlapping equations*, POPL, ACM, 2014, pp. 671–683.
- [Gar13] Jacques Garrigue, *On variance, injectivity, and abstraction*, OCaml Meeting, Boston., September 2013.
- [LP05] Ralf Lämmel and Simon Peyton Jones, *Scrap your boilerplate with class: Extensible generic functions*, ICFP, 2005.
- [Mar10] Simon Marlow (editor), *Haskell 2010 language report*, 2010.
- [MTHM97] Robin Milner, Mads Tofte, Robert Harper, and David MacQueen, *The definition of Standard ML (revised)*, 1997.
- [PL91] Simon Peyton Jones and J Launchbury, *Unboxed values as first class citizens*, FPCA, LNCS, vol. 523, 1991, pp. 636–666.
- [PTH01] Simon Peyton Jones, Andrew Tolmach, and Tony Hoare, *Playing by the rules: rewriting as a practical optimisation technique in GHC*, Haskell Workshop, 2001, pp. 203–233.
- [SR13] Gabriel Scherer and Didier Rémy, *GADTs meet subtyping*, ESOP, 2013, pp. 554–573.
- [VPMa12] Dimitrios Vytiniotis, Simon Peyton Jones, and José Pedro Magalhães, *Equality proofs and deferred type errors: A compiler pearl*, ICFP, ACM, 2012, pp. 341–352.
- [WB89] Philip Wadler and Stephen Blott, *How to make ad-hoc polymorphism less ad-hoc*, POPL, ACM, 1989, pp. 60–76.
- [WVPZ11] Stephanie Weirich, Dimitrios Vytiniotis, Simon Peyton Jones, and Steve Zdancewic, *Generative type abstraction and type-level computation*, POPL, ACM, 2011, pp. 227–240.
- [YWC⁺12] Brent A. Yorgey, Stephanie Weirich, Julien Cretin, Simon Peyton Jones, Dimitrios Vytiniotis, and José Pedro Magalhães, *Giving Haskell a promotion*, TLDI, ACM, 2012, pp. 53–66.

⁹Package authors have the option of specifying which compilers their package is known to work with. Of the 555 packages listed as working with one of the GHC 7.6 versions, 183 also are listed as compatible with GHC 7.8. These packages include 43 that use the GND extension.

A. Further discussion

There are many aspects of roles and Coercible that may be of interest, especially to those considering adopting a similar mechanism in another programming language. We collect these observations here.

A.1 Conservativity of roles

Roles are coarse-grained The system we describe has exactly three roles. However, by having only three roles, we have created a rather coarse-grained classification system. For example, consider the following definitions:

```
data Bar a = MkBar (F a)
type instance F Int = Char
type instance F Bool = Char
type instance F [a] = Double
```

It is safe to coerce a `Bar Int` to a `Bar Bool`. Unravelling definitions, we see that this is so. Yet, coercing `Bar Int` to `Bar [Double]` is clearly not safe. GHC assigns a nominal role to the parameter of `Bar`, but this choice of role eliminates the possibility of the `Bar Int` to `Bar Bool` coercion. If, instead, we had a *lattice* of roles, keyed by type families whose equality must be respected, we might be able to allow more safe coercions. We could similarly imagine a lattice keyed by classes whose instance definitions are to be respected; with such a lattice, we could allow the coercion of `Map Int v` to `Map Age v` precisely when `Int`'s and `Age`'s `Ord` instance correspond.

Equality does not propagate roles What role should be assigned to a parameter with an equality constraint involving a phantom? According to the rules in our formalism, such a parameter would get a nominal role. This case has come up in practice. Consider the following type from Edward Kmett's lens library:

```
data Magma i t b a where
  MagmaAp :: Magma i (x → y) b a
    → Magma i x b a → Magma i y b a
  MagmaPure :: x → Magma i x b a
  MagmaFmap :: (x → y)
    → Magma i x b a → Magma i y b a
  Magma :: i → a → Magma i b b a
```

Role inference gives these roles, respectively: representational, nominal, nominal, representational. Close inspection of the type definition shows us that the third parameter, `b`, is almost a phantom – it is never used outside phantom contexts except in one place: the return type of the `Magma` constructor. There, we see that the second and third parameters must be equal. Another way to write this last constructor is `Magma :: (t ~ b) ⇒ i → a → Magma i t b a`. Also, note that the second parameter, `t`, is used representationally everywhere but in that same spot.

What this all leads to is the fact that `Magma i x x a` has the same run-time representation as `Magma i y y a` whenever `x` has the same representation as `y`. Yet, the role mechanism is not expressive enough to prove this.

A.2 Type inference

Inferring polymorphic coercions Suppose we want

```
concatH :: [HTML] → HTML
```

It is tempting simply to write

```
concatH = coerce concat
```

However, this will not type-check, due to a failure of type inference. The use of `coerce` creates the constraint `Coercible ([a] → [a]) ([HTML] → HTML)`, for some `a`, but GHC cannot figure out that `a` should be `Char`. In this instance, it is quite obvious what the programmer wants, and it may be feasible to make the compiler smart enough to see that.

Explicit type application A different approach to the `concatH` problem immediately above is to be able to provide an explicit type argument to the `concat` function. With hypothetical syntax, we would want to write

```
concatH = coerce (concat @Char)
```

The use of `@` above denotes an explicitly-passed type parameter. With this new syntax, we are freed from the burden of instantiating all of `concat`'s type to specify a single type parameter. This desire for explicit type application came up numerous times in our experimentation with `coerce`, particularly in the implementation of `GND`.

A.3 Parametricity

There seems to be some relationship between roles, parametricity, and categorical structures. For example, consider the class

```
class Functor f where
  fmap :: (a → b) → f a → f b
```

with the functor law, stating that `fmap id` should be identically `id`. We conjecture that it is impossible to write a lawful instance of `Functor` over a type whose one parameter would be inferred to have a nominal role. This observation stems from the fact that `fmap` must be parametric in `a` and `b`, and thus cannot take any action which observes any nominal equality dealing with those type parameters. We have not proved this conjecture, however.

A.4 Eta-reduction

Consider the definition

```
newtype MyMaybe a = MkMMM (Maybe a)
```

What is the axiom created by this declaration? Naively, we could expect an axiom $C : [a:*\].\text{MyMaybe } a \sim_R \text{Maybe } a$. Now, suppose we are using the monad transformer `ListT`. As usual for a monad transformer, `ListT` has kind $(* \rightarrow *) \rightarrow (* \rightarrow *)$. That is, it takes a monad as its first argument and transforms it into a monad, now enhanced with non-determinism. Let's say we have a value `m` of type `ListT Maybe Int`. Is `m` coercible into `ListT MyMaybe Int`? We would like it to be – these would have the same representation. But, with the axiom C above, such a coercion would be impossible. There is no way to extract `MyMaybe ~R Maybe` from C .

Instead, we *η-reduce* axioms. Accordingly, the declaration above would yield $C_2 : \text{MyMaybe } \sim_R \text{Maybe}$. This axiom can easily be used to show that `ListT Maybe Int ~R ListT MyMaybe Int`.

This *η-reduction* is why the `CO_APP` rule works over representational coercions as well as nominal ones. An earlier version of the rule consumed and produced only nominal coercions; with that rule, it was impossible to derive `MyMaybe Int ~R Maybe Int` from C_2 .

A.5 Extending role to families

A.5.1 Roles on type and data families

In GHC 7.8, all type and data family parameters have nominal roles. This stands to good reason, as a type or data family can pattern-match on its parameters. For example:

```

type family TF a
type instance TF Int = Double
type instance TF Age = Char

```

Clearly, `TF Int` is not representationally equal to `TF Age`.

Yet, it would be sensible to extend the idea of roles to type and data families. A family with a non-nominal parameter would need extra checks on its instance declarations, to make sure that they are compatible with the choice of roles. For example:

```

type role If nominal representational representational
type family If True (a :: Bool) b c
type instance If True b c = b
type instance If False b c = c

```

The above definition, though not accepted in GHC 7.8, is perfectly type safe. Note that a representational parameter must not be matched on and must not be used in a nominal context on the right-hand side. The only barrier to implementing this is the extra complexity for the GHC maintainers and the extra complexity in the language. If a compelling use case for this comes up, we will likely add the feature.

A.5.2 Roles on data family instances

Roles on data families follow the same arguments as above. However, we can identify a separate issue involving roles on data family instances, which are, of course, data types. For example:

```

data family DF a
data instance DF (b, Int) = MkDF (Maybe b)

```

Data family instances are internally desugared into something resembling a type family instance and a fresh data type declaration, somewhat like this:¹⁰

```

type family DF a
type instance DF (b, Int) = DFPairIntInstance b
data DFPairIntInstance b = MkDF (Maybe b)

```

Here, it is apparent that `b` can be assigned a representational role, even while we require a nominal role for `a`.

Role inference for data family instances is not currently implemented. Instead, all type variables in a data family instance are assigned nominal roles. Why? Essentially because there is no way of writing a role annotation for data family instances. Without the ability to write role annotations, library writers would be unable to enforce abstraction on these, and so it is safer just to default these (somewhat uncommon) parameters to have nominal roles.

If you wish to request roles on either type/data families or on data family instances, you can comment on GHC bug #8177 here: <https://ghc.haskell.org/trac/ghc/ticket/8177>

B. Design decisions

Here, we collect some of the design decisions that we made while formulating roles concretely into GHC.

B.1 Concrete syntax

It turns out that designing the concrete syntax of role annotations was non-trivial. We identified several desired traits of the syntax:

¹⁰ Type inference is somewhat different between type families, which are not necessarily injective, and data families, which are. Along similar lines, data families can appear unsaturated, while type families cannot. This desugaring does not change these facts.

1. Role annotations must be optional. Otherwise, all existing code would be broken.
2. Role annotations should be succinct.
3. Role annotations will be a relatively obscure feature, and therefore should be searchable should a user come across one.
4. Code with role annotations should compile with older versions of GHC. This eases migration to GHC 7.8.
5. Role annotations should not be specified in a pragma; pragmas are meant to be reserved for implementation details (e.g., optimising), and roles are a type system feature.
6. Role annotations should be easy to refactor as a data type evolves.
7. Code is read much more often than it is written; favour readability over concision.

We will use `Map` as a running example to demonstrate the various alternatives we considered for the syntax. Note that all options satisfy desire (1).

1. Standalone role annotations:

```

type role Map nominal representational
data Map k v = ...

```

This is, of course, our final answer to the concrete syntax question. It satisfies (3), (5), and (7), at the cost of some others. In particular, this choice is not backward-compatible. A role annotation fails to parse in earlier versions of GHC. However, all is not lost, because GHC supports C preprocessor directives, and library authors can selectively include role annotations using preprocessor directives. The fact that the annotations are standalone means they can be grouped under one set of directives instead of sprinkled throughout the source file. Note that this syntax is very easy to search for, and the written-out nature of the roles makes them readable, if not so concise.

2. Inline, abbreviated role annotations:

```

data Map k@N v@R = ...

```

This version satisfies (2), (5), (6). It is not backwards-compatible, and the fact that the role annotations are inline means that each role-annotated definition would need its own preprocessor directives. Furthermore, now the type definition itself must be repeated in various places, so refactoring becomes more burdensome. This version is also not readable by non-experts and is nearly impossible to search for if a user is confused.

3. Inline pragmas:

```

data Map {—# ROLE nominal #—} k
         {—# ROLE representational #—} v = ...

```

This version satisfies (3), (4), (6), (7), but it clearly uses a pragma. We felt that, in several years, we would regret this decision. Backwards-compatibility would no longer be an issue, and we would be stuck with a pragma syntax for a core language feature.

4. Custom class constraints:

```

data Map k v = ...
instance NextParamNominal Map
instance NextParamRepresentational (Map k)

```

This version satisfies (3), (4), (5), (7). But, there is a mismatch between the class instance mechanism and role annotations. In particular, declarations such as these would make no sense if orphaned (that is, if put in a separate module from the parent data type declaration). Furthermore, what would it mean if one of these classes were used as a constraint on a function, like so:

```
fun :: (NextParamRepresentational f, Functor f) => ...
```

If we allowed the syntax, we would also have to implement the full complication of the roles system developed by [WVPZ11]. We didn't think the full system was necessary, so we would have to restrict the NextParamXXX classes to only instance declarations. In the end, there seemed to be too much of a mismatch to make this viable.

5. Custom Coercible instances:

```
data Map k v = ...
instance Coercible v1 v2 =>
  Coercible (Map k v1) (Map k v2)
```

This version satisfies (3), (4), (5), and (7). The roles are implicitly discovered by the usages of the variables in the declaration. This form has the advantage that it makes the Coercible instances obvious. However, it has many of the drawbacks of the previous version, with NextParamXXX. One important advantage this has over the previous version is that it doesn't lead to the full implementation of the system in [WVPZ11] – implication constraints are not allowed, so there is no way of writing the declaration above as a constraint. However, what would it mean if this annotation were omitted? It would seem strange that a Coercible instance would be automatically supplied if none were written. Also, what would it mean if the declaration used variables in a way inconsistent with roles? In the end, we felt that this version was in a bit of an uncanny valley: it's rather close to a "normal" instance declaration, but with some unexpected features and consequences. We thought it would lead to more confusion than other versions.

Do we staunchly defend our choice of the **type role** ... syntax? No, but we were unable to come up with a better one that might stand the test of time.

B.2 Safe Haskell

One of the consequences of the unsoundness of earlier versions of GND is that the feature was prohibited from the Safe Haskell subset. However, even after roles were implemented and GND written in terms of coerce, the feature *still* did not meet the Safe Haskell criteria. At issue were both abstraction and coherence:

Abstraction We describe in Section 3.1 that we allow coercions to happen even on data types for which the constructors are not available, such as Map. However, this violates Safe Haskell's promise that no abstraction barrier is broken through. To rectify this problem, GHC uses a more stringent check when satisfying a Coercible constraint when compiling in Safe mode: all constructors of all data types to be coerced under must be visible. This means, essentially, traversing the entire tree of data type definitions, making sure all constructors of all data types, recursively, are available. With this check in place, we can be sure not to break any abstraction boundaries.

Coherence Haskell classes are compiled into regular data-types in GHC Core. Accordingly, classes have roles assigned to their parameters. However, we tend to think of, for example, Ord String to be quite independent from Ord HTML. Thus, as discussed in Section 3.2, we default all class roles to be nominal.

However, it's possible that a user wishes to override this default. As an interesting example, it is quite sensible that Coercible's roles should be representational! This fact can be seen in the fact that the roles of the (\sim_R) operator are representational. Accordingly, we allow role annotations for classes, even though roles other than nominal can lead quickly to incoherence.

Safe Haskell claims to enforce class coherence. Thus, it is important that Safe Haskell restricts role annotations on classes. This is done by requiring the extension IncoherentInstances (which is disallowed in Safe Haskell) to have a non-trivial role annotation on a class.

C. System FC, in full

Throughout this entire proof of type safety, any omitted proof is by (perhaps mutual) straightforward induction on the relevant derivations.

As usual, all definitions and proofs are only up to α -equivalence. If there is a name clash, assume a variable renaming to a fresh variable.

We assume the regularity of typing judgements throughout the proof. That is, if $\Gamma \vdash \tau : \kappa$, $\Gamma \vdash \gamma : \phi$, or $\Gamma \vdash e : \tau$, we can conclude $\vdash \Gamma$. Accordingly, whenever proving one of the judgements above, we must also prove $\vdash \Gamma$. In practice, tracking these context consistency judgements presents no problems and is elided throughout.

C.1 The remainder of the grammar

Φ	::=	$[\bar{\alpha}:\bar{\kappa}].\tau \sim_\rho \sigma$	axiom types
e	::=	v	value
		x	variable
		$e_1 e_2$	application
		$e \tau$	type application
		$e \gamma$	coercion application
		case $_\tau e$ of \overline{alt}	pattern match
		$e \triangleright \gamma$	cast
		contra $\gamma \tau$	absurdity
v	::=	$\lambda x:\tau.e$	expression values
		$\Lambda \alpha:\kappa.e$	value abstraction
		$\lambda c:\phi.e$	type abstraction
		$K \bar{\tau} \bar{\gamma} \bar{e}$	coercion abstraction
			applied data constructor
alt	::=	$K \bar{\alpha} \bar{c} \bar{x} \rightarrow e$	alternative in pattern match
ψ	::=	D	value types
		(\rightarrow)	data type (<i>not newtypes!</i>)
		(\Rightarrow)	arrow
		(\Rightarrow^K)	prop. arrow
		(\sim_ρ^K)	equality
		$\forall \alpha:\kappa.\tau$	polymorphism
		$\psi \tau$	application

C.2 Typing judgements

Note that the statement, for example, $\alpha \# \Gamma$ means that the variable α is fresh in the context Γ .

$\boxed{\vdash \Gamma}$ Context validity

$$\frac{}{\vdash \emptyset} \text{CTX_EMPTY}$$

$$\frac{\alpha \# \Gamma}{\vdash \Gamma, \alpha : \kappa} \text{CTX_TYVAR}$$

$$\frac{\Gamma \vdash \tau \sim_{\rho} \sigma : \star \quad c \# \Gamma}{\vdash \Gamma, c : \phi} \text{CTX_COVAR}$$

$$\frac{\Gamma \vdash \tau : \star \quad x \# \Gamma}{\vdash \Gamma, x : \tau} \text{CTX_VAR}$$

$\boxed{\Gamma \vdash \tau : \kappa}$ Type kinding

$$\frac{\alpha : \kappa \in \Gamma}{\Gamma \vdash \alpha : \kappa} \text{TY_VAR}$$

$$\frac{\Gamma \vdash \tau_1 : \kappa_1 \rightarrow \kappa_2 \quad \Gamma \vdash \tau_2 : \kappa_1}{\Gamma \vdash \tau_1 \tau_2 : \kappa_2} \text{TY_APP}$$

$$\frac{T : \kappa}{\Gamma \vdash T : \kappa} \text{TY_ADT}$$

$$\frac{}{\Gamma \vdash (\rightarrow) : \star \rightarrow \star \rightarrow \star} \text{TY_ARROW}$$

$$\frac{}{\Gamma \vdash (\Rightarrow) : \star \rightarrow \star \rightarrow \star} \text{TY_PROPARGROW}$$

$$\frac{}{\Gamma \vdash (\sim_{\rho}^{\kappa}) : \kappa \rightarrow \kappa \rightarrow \star} \text{TY_EQUALITY}$$

$$\frac{\Gamma, \alpha : \kappa \vdash \tau : \star}{\Gamma \vdash \forall \alpha : \kappa. \tau : \star} \text{TY_FORALL}$$

$$\frac{F : [\overline{\alpha} : \kappa'] . \kappa}{\Gamma \vdash \tau : \kappa'} \text{TY_TYFUN}$$

$$\frac{}{\Gamma \vdash F(\overline{\tau}) : \kappa} \text{TY_TYFUN}$$

$\boxed{\Gamma \vdash e : \tau}$ Expression typing

$$\frac{x : \tau \in \Gamma}{\Gamma \vdash x : \tau} \text{TM_VAR}$$

$$\frac{\Gamma, x : \tau \vdash e : \sigma}{\Gamma \vdash \lambda x : \tau. e : \tau \rightarrow \sigma} \text{TM_ABS}$$

$$\frac{\Gamma \vdash e_1 : \tau \rightarrow \sigma \quad \Gamma \vdash e_2 : \tau}{\Gamma \vdash e_1 e_2 : \sigma} \text{TM_APP}$$

$$\frac{\Gamma, \alpha : \kappa \vdash e : \tau}{\Gamma \vdash \Lambda \alpha : \kappa. e : \forall \alpha : \kappa. \tau} \text{TM_TABS}$$

$$\frac{\Gamma \vdash e : \forall \alpha : \kappa. \sigma \quad \Gamma \vdash \tau : \kappa}{\Gamma \vdash e \tau : \sigma[\tau/\alpha]} \text{TM_TAPP}$$

$$\frac{\Gamma, c : \sigma_1 \sim_{\rho} \sigma_2 \vdash e : \tau}{\Gamma \vdash \lambda c : \sigma_1 \sim_{\rho} \sigma_2. e : \phi \Rightarrow \tau} \text{TM_CABS}$$

$$\frac{\Gamma \vdash e : (\sigma_1 \sim_{\rho} \sigma_2) \Rightarrow \tau \quad \Gamma \vdash \gamma : \sigma_1 \sim_{\rho} \sigma_2}{\Gamma \vdash e \gamma : \tau} \text{TM_CAPP}$$

$$\frac{K : \tau}{\Gamma \vdash K : \tau} \text{TM_DATACON}$$

$$\Gamma \vdash e : D \overline{\sigma}$$

$$\Gamma \vdash \tau : \star$$

$$\forall \text{alt}_i \text{ s.t. } \text{alt}_i \in \overline{\text{alt}} :$$

$$\text{alt}_i = K_i \overline{\alpha}_i \overline{c}_i \overline{x}_i \rightarrow e_i$$

$$K_i : \forall \overline{\alpha}'_i : \kappa_i. \forall \overline{\beta}'_i : \kappa'_i. \overline{\phi}_i \Rightarrow \overline{\tau}_i \rightarrow D \overline{\alpha}'_i$$

$$\Gamma, \overline{\alpha}_i : \kappa'_i, (\overline{c}_i : \phi_i, \overline{x}_i : \tau_i) [\overline{\sigma} / \overline{\alpha}'_i] [\overline{\alpha}_i / \overline{\beta}'_i] \vdash e_i : \tau$$

$$\overline{\text{alt}} \text{ is exhaustive}$$

$$\frac{}{\Gamma \vdash \text{case}_{\tau} e \text{ of } \overline{\text{alt}} : \tau} \text{TM_CASE}$$

$$\frac{\Gamma \vdash e : \tau_1 \quad \Gamma \vdash \gamma : \tau_1 \sim_{\mathbb{R}} \tau_2}{\Gamma \vdash e \triangleright \gamma : \tau_2} \text{TM_CAST}$$

$$\frac{\emptyset \vdash \gamma : H_1 \sim_{\mathbb{N}} H_2 \quad H_1 \neq H_2}{\Gamma \vdash \tau : \star} \text{TM_CONTRA}$$

C.3 Small-step operational semantics

$\boxed{e_1 \longrightarrow e_2}$ Small-step operational semantics

$$\frac{}{\overline{(\lambda x : \tau. e_1) e_2} \longrightarrow e_1[e_2/x]} \text{S_BETA}$$

$$\frac{}{\overline{(\Lambda \alpha : \kappa. e) \tau} \longrightarrow e[\tau/\alpha]} \text{S_TBETA}$$

$$\frac{}{\overline{(\lambda c : \phi. e) \gamma} \longrightarrow e[\gamma/c]} \text{S_CBETA}$$

$$\frac{\text{alt}_i = K \overline{\alpha} \overline{c} \overline{x} \rightarrow e'}{\text{case}_{\tau_0} K \overline{\tau} \overline{\sigma} \overline{\gamma} \overline{e} \text{ of } \overline{\text{alt}} \longrightarrow e'[\overline{\sigma}/\overline{\alpha}][\overline{\gamma}/\overline{c}][\overline{e}/\overline{x}]} \text{S_IOTA}$$

$$\frac{}{\overline{(v \triangleright \gamma_1) \triangleright \gamma_2} \longrightarrow v \triangleright (\gamma_1 \mathbin{\&}; \gamma_2)} \text{S_TRANS}$$

$$\frac{e_1 \longrightarrow e'_1}{e_1 e_2 \longrightarrow e'_1 e_2} \text{S_APP_CONG}$$

$$\frac{e \longrightarrow e'}{e \tau \longrightarrow e' \tau} \text{S_TAPP_CONG}$$

$$\frac{e \longrightarrow e'}{e \gamma \longrightarrow e' \gamma} \text{S_CAPP_CONG}$$

$$\frac{e \longrightarrow e'}{\text{case}_{\tau} e \text{ of } \overline{\text{alt}} \longrightarrow \text{case}_{\tau} e' \text{ of } \overline{\text{alt}}} \text{S_CASE_CONG}$$

$$\frac{e \longrightarrow e'}{e \triangleright \gamma \longrightarrow e' \triangleright \gamma} \text{S_CAST_CONG}$$

$$\frac{\eta_1 = \text{sym}(\text{nth}^1 \eta_0) \quad \eta_2 = \text{nth}^2 \eta_0}{\emptyset \vdash v : \sigma_1 \rightarrow \sigma_2 \quad \emptyset \vdash v : \sigma_1 \rightarrow \sigma_2}{(v \triangleright \eta_0) e' \longrightarrow v (e' \triangleright \eta_1) \triangleright \eta_2} \text{S_PUSH}$$

$$\frac{\emptyset \vdash v : \forall \alpha : \kappa. \sigma' \quad \emptyset \vdash \tau : \kappa}{(v \triangleright \gamma) \tau \longrightarrow v \tau \triangleright \gamma @ \tau} \text{S_TPUSH}$$

$$\frac{\eta_{11} = \text{nth}^1(\text{nth}^1 \eta_0) \quad \eta_{12} = \text{nth}^2(\text{nth}^1 \eta_0) \quad \eta_2 = \text{nth}^2 \eta \quad \gamma'' = \eta_{11} \mathbin{\&}; \gamma' \mathbin{\&}; \text{sym} \eta_{12}}{\emptyset \vdash v : \sigma_1 \sim_{\rho} \sigma_2 \Rightarrow \sigma_3 \quad \emptyset \vdash \gamma' : \sigma_4 \sim_{\rho} \sigma_5}{(v \triangleright \eta_0) \gamma' \longrightarrow v \gamma'' \triangleright \eta_2} \text{S_CPUSH}$$

$$\frac{\emptyset \vdash \eta : D \overline{\tau} \sim_{\mathbb{R}} D \overline{\tau}' \quad K : \forall \overline{\alpha} : \kappa. \forall \overline{\beta} : \kappa'. (\sigma' \sim_{\rho} \sigma'') \Rightarrow \overline{\tau}'' \rightarrow D \overline{\alpha}}{\emptyset \vdash \gamma : (\sigma' \sim_{\rho} \sigma'') [\overline{\tau} / \overline{\alpha}] [\overline{\sigma} / \overline{\beta}]} \text{S_KPPUSH}$$

$$\frac{\gamma' = \text{sym}(\sigma' [\text{nth} \eta / \overline{\alpha}]_{\rho}) \mathbin{\&}; \gamma \mathbin{\&}; \sigma'' [\text{nth} \eta / \overline{\alpha}]_{\rho}}{e' = e \triangleright \tau'' [\text{nth} \eta / \overline{\alpha}]_{\mathbb{R}}} \text{S_KPPUSH}$$

$$\frac{}{\text{case}_{\tau_0} (K \overline{\tau} \overline{\sigma} \overline{\gamma} \overline{e}) \triangleright \eta \text{ of } \overline{\text{alt}} \longrightarrow \text{case}_{\tau_0} K \overline{\tau}' \overline{\sigma}' \overline{\gamma}' \overline{e}' \text{ of } \overline{\text{alt}}} \text{S_KPPUSH}$$

D. Global context well-formedness

We assume throughout the paper and this appendix that the global context is well formed. Here, we explain precisely what can appear in the global context and what restrictions there are:

1. The global context may contain $C : [\bar{\alpha}:\bar{\kappa}].\tau \sim_{\rho} \sigma$:
There are two forms of axiom, for which different rules apply:
 - (a) Newtype axioms: All of the following must hold.
 - i. $\tau = N\bar{\alpha}$
 - ii. $\rho = R$
 - iii. There must not be two axioms mentioning the same newtype N .
 - iv. The length of $\text{roles}(N)$ must match the arity of the axiom C .
 - (b) Type family axioms: All of the following must hold.
 - i. $\tau = F(\bar{\tau}')$
 - ii. $\rho = N$
 - iii. The types $\bar{\tau}'$ must not mention type families.
 - iv. Each $\beta \in \bar{\alpha}$ must appear exactly once in the list $\bar{\tau}'$.
 Regardless of the form of axiom, the following must hold:
 - (c) $\bar{\alpha}:\bar{\kappa} \vdash \tau : \kappa_0$
 - (d) $\bar{\alpha}:\bar{\kappa} \vdash \sigma : \kappa_0$
 - (e) Consider two axioms $C_1 : [\bar{\alpha}:\bar{\kappa}].\tau_1 \sim_{\rho} \sigma_1$ and $C_2 : [\bar{\beta}:\bar{\kappa}'].\tau_2 \sim_{\rho} \sigma_2$ (where variables are renamed so that $\bar{\alpha} \cap \bar{\beta} = \emptyset$). Then, if there exists some θ with $\theta(\tau_1) = \theta(\tau_2)$, it must be that $\theta(\sigma_1) = \theta(\sigma_2)$.
2. The global context may contain $T : \kappa$.
3. The global context may contain $K : \tau$:
 - (a) $\tau = \forall \bar{\alpha}:\bar{\kappa}.\forall \bar{\beta}:\bar{\kappa}'.\bar{\phi} \Rightarrow \bar{\sigma} \rightarrow D\bar{\alpha}$
 - (b) $\emptyset \vdash \tau : \star$
4. The global context may contain $F : [\bar{\alpha}:\bar{\kappa}].\kappa_0$.
5. For all H , $\text{roles}(H) \models H$.

E. Properties of roles

Lemma 1 (Permutation of role checking). *If $\Omega \vdash \tau : \rho$ and Ω' is a permutation of Ω , then $\Omega' \vdash \tau : \rho$.*

Lemma 2 (Weakening of role checking). *If $\Omega \vdash \tau : \rho$, then $\Omega, \alpha:\rho' \vdash \tau : \rho$.*

Lemma 3 (Strengthening of role checking). *If $\Omega, \alpha:\rho' \vdash \tau : \rho$ and α does not appear free in τ , then $\Omega \vdash \tau : \rho$.*

Lemma 4 (Nominal roles are infectious). *Let $\bar{\alpha}$ be the free variables in σ . We have $\Omega \vdash \sigma : N$ if and only if every $\alpha_i \in \bar{\alpha}$ is at role N in Ω .*

Lemma 5 (Sub-roling). *If $\Omega \vdash \tau : \rho$ and $\rho \leq \rho'$, then $\Omega \vdash \tau : \rho'$.*

F. Structural properties

F.1 Weakening

Let bnd be a metavariable for a context binding. That is:

$$bnd ::= \begin{array}{l} \alpha:\kappa \\ | \\ c:\phi \\ | \\ x:\tau \end{array}$$

Lemma 6 (Type kinding weakening). *If $\Gamma, \Gamma' \vdash \tau : \kappa$ and $\vdash \Gamma, bnd, \Gamma'$, then $\Gamma, bnd, \Gamma' \vdash \tau : \kappa$.*

Lemma 7 (Coercion typing weakening). *If $\Gamma, \Gamma' \vdash \gamma : \phi$ and $\vdash \Gamma, bnd, \Gamma'$, then $\Gamma, bnd, \Gamma' \vdash \gamma : \phi$.*

Lemma 8 (Term typing weakening). *If $\Gamma, \Gamma' \vdash e : \tau$ and $\vdash \Gamma, bnd, \Gamma'$, then $\Gamma, bnd, \Gamma' \vdash e : \tau$.*

F.2 Substitution

Lemma 9 (Type variable substitution). *Suppose $\Gamma \vdash \sigma : \kappa_1$. Then:*

1. *If $\vdash \Gamma, \alpha:\kappa_1, \Gamma'$, then $\vdash \Gamma, \Gamma'[\sigma/\alpha]$;*
2. *If $\Gamma, \alpha:\kappa_1, \Gamma' \vdash \tau : \kappa_2$, then $\Gamma, \Gamma'[\sigma/\alpha] \vdash \tau[\sigma/\alpha] : \kappa_2$.*

Lemma 10 (Type variable substitution in coercions). *If $\Gamma, \alpha:\kappa, \Gamma' \vdash \gamma : \phi$ and $\Gamma \vdash \sigma : \kappa$, then $\Gamma, \Gamma'[\sigma/\alpha] \vdash \gamma[\sigma/\alpha] : \phi[\sigma/\alpha]$.*

Lemma 11 (Type variable substitution in terms). *If $\Gamma, \alpha:\kappa, \Gamma' \vdash e : \tau$ and $\Gamma \vdash \sigma : \kappa$, then $\Gamma, \Gamma'[\sigma/\alpha] \vdash e[\sigma/\alpha] : \tau[\sigma/\alpha]$.*

Lemma 12 (Coercion strengthening).

1. *If $\vdash \Gamma, c:\phi, \Gamma'$, then $\vdash \Gamma, \Gamma'$;*
2. *If $\Gamma, c:\phi, \Gamma' \vdash \tau : \kappa$, then $\Gamma, \Gamma' \vdash \tau : \kappa$.*

Lemma 13 (Coercion substitution). *If $\Gamma, c:\phi_1, \Gamma' \vdash \gamma : \phi_2$ and $\Gamma \vdash \eta : \phi_1$, then $\Gamma, \Gamma' \vdash \gamma[\eta/c] : \phi_2$.*

Lemma 14 (Coercion substitution in terms). *If $\Gamma, c:\phi, \Gamma' \vdash e : \tau$ and $\Gamma \vdash \eta : \phi$, then $\Gamma, \Gamma' \vdash e[\eta/c] : \tau$.*

Lemma 15 (Term strengthening).

1. *If $\vdash \Gamma, x:\tau, \Gamma'$, then $\vdash \Gamma, \Gamma'$;*
2. *If $\Gamma, x:\tau, \Gamma' \vdash \sigma : \kappa$, then $\Gamma, \Gamma' \vdash \sigma : \kappa$.*

Lemma 16 (Term strengthening in coercions). *If $\Gamma, x:\tau, \Gamma' \vdash \gamma : \phi$, then $\Gamma, \Gamma' \vdash \gamma : \phi$.*

Lemma 17 (Term substitution). *If $\Gamma, x:\sigma, \Gamma' \vdash e : \tau$ and $\Gamma \vdash e' : \sigma$, then $\Gamma, \Gamma' \vdash e[e'/x] : \tau$.*

F.3 Classifier regularity

Lemma 18 (Coercion typing regularity). *If $\Gamma \vdash \gamma : \tau \sim_{\rho} \sigma$, then $\Gamma \vdash \tau \sim_{\rho} \sigma : \star$.*

Lemma 19 (Coercion homogeneity). *If $\Gamma \vdash \gamma : \tau \sim_{\rho} \sigma$, then $\Gamma \vdash \tau : \kappa$ and $\Gamma \vdash \sigma : \kappa$.*

Proof. Direct from Lemma 18. □

Lemma 20 (Term typing regularity). *If $\Gamma \vdash e : \tau$, then $\Gamma \vdash \tau : \star$.*

G. Preservation

G.1 Lifting

Lifting is defined by the following algorithm, with patterns to be tried in order from top to bottom.

$$\begin{array}{lcl}
\tau[\overline{\gamma/\beta}]_{\text{P}} & = & \langle \tau[\overline{\sigma/\beta}], \tau[\overline{\sigma'/\beta}] \rangle_{\text{P}} \quad (\overline{\Gamma} \vdash \gamma : \sigma \sim_{\rho} \sigma') \\
\alpha[\overline{\gamma/\beta}]_{\rho} & = & \gamma_i \quad (\alpha = \beta_i \wedge \overline{\Gamma} \vdash \gamma_i : \sigma \sim_{\rho} \sigma') \\
\alpha[\overline{\gamma/\beta}]_{\text{R}} & = & \mathbf{sub} \gamma_i \quad (\alpha = \beta_i) \\
\alpha[\overline{\gamma/\beta}]_{\text{N}} & = & \langle \alpha \rangle \quad (\alpha \notin \beta) \\
\alpha[\overline{\gamma/\beta}]_{\text{R}} & = & \mathbf{sub} \langle \alpha \rangle \quad (\alpha \notin \beta) \\
(H\overline{\tau})[\overline{\gamma/\beta}]_{\text{R}} & = & H(\tau[\overline{\gamma/\beta}]_{\rho}) \quad (\overline{\rho} \text{ is a prefix of } \mathit{roles}(H)) \\
H[\overline{\gamma/\beta}]_{\text{N}} & = & \langle H \rangle \\
(\tau_1 \tau_2)[\overline{\gamma/\beta}]_{\rho} & = & \tau_1[\overline{\gamma/\beta}]_{\rho} \tau_2[\overline{\gamma/\beta}]_{\text{N}} \\
(\forall \alpha : \kappa. \tau)[\overline{\gamma/\beta}]_{\rho} & = & \forall \alpha : \kappa. \tau[\overline{\gamma/\beta}]_{\rho} \\
(F(\overline{\tau}))[\overline{\gamma/\beta}]_{\text{N}} & = & F(\tau[\overline{\gamma/\beta}]_{\text{N}}) \\
(F(\overline{\tau}))[\overline{\gamma/\beta}]_{\text{R}} & = & \mathbf{sub} F(\tau[\overline{\gamma/\beta}]_{\text{N}})
\end{array}$$

Lemma 21 (Lifting). *If:*

1. $\overline{\Gamma} \vdash \gamma : H\overline{\tau} \sim_{\text{R}} H\overline{\sigma}$;
2. $\overline{\Gamma} \vdash \tau : \kappa$;
3. $\overline{\Gamma} \vdash \sigma : \kappa$;
4. H is not a **newtype**;
5. $\Omega \vdash \sigma_0 : \rho_0$, where $\overline{\beta}$ is the type variables in $\overline{\Gamma}, \Gamma'$
 $\Omega = \overline{\beta}' : \text{N}, \overline{\beta} : \mathit{roles}(H)$;
6. $\overline{\Gamma}, \overline{\beta} : \kappa, \Gamma' \vdash \sigma_0 : \kappa'$; and
7. Γ' contains only type variable bindings.

then:

$$\overline{\Gamma}, \Gamma' \vdash \sigma_0[\overline{\mathbf{nth}} \gamma/\beta]_{\rho_0} : \sigma_0[\overline{\tau/\beta}] \sim_{\rho_0} \sigma_0[\overline{\sigma/\beta}]$$

Proof. First, because Γ' contains only type variable bindings, then a type variable substitution has no effect on Γ' (which can contain only *kinds*).

If $\rho_0 = \text{P}$, then the first equation of the algorithm matches, and we have $\sigma_0[\overline{\mathbf{nth}} \gamma/\beta]_{\text{P}} = \langle \sigma_0[\overline{\tau/\beta}], \sigma_0[\overline{\sigma/\beta}] \rangle_{\text{P}}$, and we are done, applying Lemma 9.

So, we assume now that $\rho_0 \neq \text{P}$.

Let $\overline{\rho} = \mathit{roles}(H)$.

We proceed by induction on the derivation of $\overline{\Gamma}, \overline{\beta} : \kappa, \Gamma' \vdash \sigma_0 : \kappa'$. Each case concludes by the application of the appropriate substitution lemma(s).

Case TY_VAR: We know $\sigma_0 = \alpha$.

Case ($\alpha = \beta_i$):

Case ($\rho_0 = \rho_i$): Here, we have $\sigma_0[\overline{\mathbf{nth}} \gamma/\beta]_{\rho_0} = \mathbf{nth}^i \gamma$, $\sigma_0[\overline{\tau/\beta}] = \tau_i$, and $\sigma_0[\overline{\sigma/\beta}] = \sigma_i$. Thus, we are done, by CO_NTH.

Case ($\rho_0 = \text{R}, \rho_i = \text{N}$): Similar to the last case, fixing the roles with a use of **sub**.

Case ($\rho_0 = \text{N}, \rho_i \neq \text{N}$): This case is impossible. We know $\Omega \vdash \alpha : \text{N}$. By inversion then, we know $\alpha : \text{N} \in \Omega$. Yet, we know that ρ_i is the i th role in $\mathit{roles}(H)$, and by the definition of $\Omega, \alpha : \rho_i \in \Omega$. This contradicts $\rho_i \neq \text{N}$, and we are done.

Case ($\alpha \notin \overline{\beta}$):

Case ($\rho_0 = \text{N}$): Here, $\sigma_0[\overline{\mathbf{nth}} \gamma/\beta]_{\text{N}} = \langle \sigma_0 \rangle$, $\sigma_0[\overline{\tau/\beta}] = \sigma_0$, and $\sigma_0[\overline{\sigma/\beta}] = \sigma_0$, so we are done, by CO_REFL.

Case ($\rho_0 = \text{R}$): Similar to last case, fixing the output role with **sub**.

Case TY_APP:

Case ($\sigma_0 = H' \overline{\sigma'}, \rho_0 = \text{R}$): Here $(H' \overline{\sigma'})[\overline{\mathbf{nth}} \gamma/\beta]_{\text{R}} = H'(\overline{\sigma'}[\overline{\mathbf{nth}} \gamma/\beta]_{\rho'})$, where $\overline{\rho}'$ is a prefix of $\mathit{roles}(H')$.

Let $\eta = H'(\overline{\sigma'}[\overline{\mathbf{nth}} \gamma/\beta]_{\rho'})$. Then, we must show

$\overline{\Gamma}, \Gamma' \vdash \eta : H' \overline{\sigma'}[\overline{\tau/\beta}] \sim_{\text{R}} H' \overline{\sigma'}[\overline{\sigma/\beta}]$. We will use CO_TYCONAPP. We must show

$$\overline{\Gamma}, \Gamma' \vdash \overline{\sigma'}[\overline{\mathbf{nth}} \gamma/\beta]_{\rho'} : \overline{\sigma'}[\overline{\tau/\beta}] \sim_{\rho'} \overline{\sigma'}[\overline{\sigma/\beta}].$$

We do this by induction, for each $\sigma'_i \in \overline{\sigma}'$. All of the premises of the lifting lemma are satisfied automatically, except for premise 5. Fix i . We must show $\Omega \vdash \sigma'_i : \rho'_i$. We know $\Omega \vdash H' \overline{\sigma}' : \text{R}$. This can be proved by either RTY_TYCONAPP or RTY_APP. If it is by the former, we are done by inversion. If it is by the latter, then we know $\Omega \vdash \sigma'_i : \text{N}$. We apply Lemma 5, and we are done.

Other applications: Apply the induction hypothesis. Premise 5 of the lifting lemma is satisfied by correspondence between RTY_APP and CO_APP.

Case TY_ADT:

Case ($\rho_0 = \text{N}$): Here $H[\overline{\mathbf{nth}} \gamma/\beta]_{\text{N}} = \langle H \rangle$, and we are done by CO_REFL.

Case ($\rho_0 = \text{R}$): Here $H[\overline{\mathbf{nth}} \gamma/\beta]_{\text{R}} = H(\emptyset)$ and we are done by CO_TYCONAPP.

Cases TY_ARROW, TY_EQUALITY: Similar to TY_ADT.

Case TY_FORALL: By the induction hypothesis. Note that the roles in RTY_FORALL and CO_FORALL match up, and that the new binding in RTY_FORALL is given a nominal role, echoed in the definition of Ω in this lemma's premises.

Case TY_TYFUN: By the induction hypothesis, once again noting the correspondence between RTY_TYFAM and CO_TYFAM. □

G.2 Preservation

Theorem 22 (Preservation). *If $\overline{\Gamma} \vdash e : \tau$ and $e \longrightarrow e'$, then $\overline{\Gamma} \vdash e' : \tau$.*

Proof. By induction on the derivation of $e \longrightarrow e'$.

Beta rules: By substitution.

Case S_IOTA: We know $\overline{\Gamma} \vdash \mathbf{case}_{\tau_0} K \overline{\tau} \overline{\sigma} \overline{\gamma} \overline{e} \mathbf{of} \overline{alt} : \tau_0$, where $\overline{alt}_i = K \overline{\alpha} \overline{c} \overline{x} \rightarrow e'$. We must show $\overline{\Gamma} \vdash e'[\overline{\sigma/\alpha}][\overline{\gamma/c}][\overline{e/x}] : \tau_0$. By inversion on TM_CASE, we see

$$\overline{\Gamma} \vdash K \overline{\tau} \overline{\sigma} \overline{\gamma} \overline{e} : D \overline{\tau}$$

$$K : \forall \overline{\alpha}' : \kappa. \forall \overline{\beta}' : \kappa'. \overline{\phi} \Rightarrow \overline{\tau}' \rightarrow D \overline{\alpha}'$$

$$\overline{\Gamma}, \overline{\alpha} : \kappa', c : \overline{\phi}[\overline{\tau/\alpha'}][\overline{\alpha/\beta'}], x : \overline{\tau}'[\overline{\tau/\alpha'}][\overline{\alpha/\beta'}] \vdash e' : \tau_0$$

We also know that $\overline{\Gamma} \vdash \tau_0 : \star$, which implies that none of the variables $\overline{\alpha}$ are mentioned in τ_0 . We can do induction on the length of $\overline{\tau}$ to see that

$$\overline{\Gamma} \vdash K \overline{\tau} : \forall \overline{\beta}' : \kappa'. \overline{\phi}[\overline{\tau/\alpha'}] \Rightarrow \overline{\tau}'[\overline{\tau/\alpha'}] \rightarrow D \overline{\alpha}'[\overline{\tau/\alpha'}]$$

This simplifies to

$$\overline{\Gamma} \vdash K \overline{\tau} : \forall \overline{\beta}' : \kappa'. \overline{\phi}[\overline{\tau/\alpha'}] \Rightarrow \overline{\tau}'[\overline{\tau/\alpha'}] \rightarrow D \overline{\tau}$$

Now, we do induction on the length of $\overline{\sigma}$ to see that

$$\overline{\Gamma} \vdash K \overline{\sigma} : \overline{\phi}[\overline{\tau/\alpha'}][\overline{\sigma/\beta'}] \Rightarrow \overline{\tau}'[\overline{\tau/\alpha'}][\overline{\sigma/\beta'}] \rightarrow D \overline{\tau}$$

and

$$\overline{\Gamma} \vdash \sigma : \kappa'$$

We can then use repeated application of the type variable substitution lemma to get

$$\overline{\Gamma}, c : \overline{\phi}[\overline{\tau/\alpha'}][\overline{\sigma/\beta'}], x : \overline{\tau}'[\overline{\tau/\alpha'}][\overline{\sigma/\beta'}] \vdash e'[\overline{\sigma/\alpha}] : \tau_0$$

using the following facts

$$\begin{aligned}\tau_0[\overline{\sigma/\alpha}] &= \tau_0 \\ \phi[\overline{\tau/\alpha'}][\overline{\alpha/\beta'}][\overline{\sigma/\alpha}] &= \phi[\overline{\tau/\alpha'}][\overline{\sigma/\beta'}] \\ \tau'[\overline{\tau/\alpha'}][\overline{\alpha/\beta'}][\overline{\sigma/\alpha}] &= \tau'[\overline{\tau/\alpha'}][\overline{\sigma/\beta'}]\end{aligned}$$

So, we have

$$\Gamma, c: \overline{\phi[\overline{\tau/\alpha'}][\overline{\sigma/\beta'}]}, x: \overline{\tau'[\overline{\tau/\alpha'}][\overline{\sigma/\beta'}]} \vdash e'[\overline{\sigma/\alpha}] : \tau_0$$

Starting from the type of $K\bar{\tau}\bar{\sigma}$, we do induction on the length of $\bar{\gamma}$ to get

$$\Gamma \vdash K\bar{\tau}\bar{\sigma}\bar{\gamma} : \overline{\tau'[\overline{\tau/\alpha'}][\overline{\sigma/\beta'}]} \rightarrow D\bar{\tau}$$

and

$$\overline{\Gamma \vdash \gamma : \phi[\overline{\tau/\alpha'}][\overline{\sigma/\beta'}]}$$

Thus, we can use the coercion variable substitution lemma to get

$$\overline{\Gamma, x: \tau'[\overline{\tau/\alpha'}][\overline{\sigma/\beta'}]} \vdash e'[\overline{\sigma/\alpha}][\overline{\gamma/c}] : \tau_0$$

Finally we use analogous reasoning for term arguments \bar{e} to conclude

$$\Gamma \vdash e'[\overline{\sigma/\alpha}][\overline{\gamma/c}][\overline{e/x}] : \tau_0$$

as desired.

Case S_TRANS: We know that $\Gamma \vdash (v \triangleright \gamma_1) \triangleright \gamma_2 : \tau$ and need to show that $\Gamma \vdash v \triangleright (\gamma_1 \mathbin{\&\#} \gamma_2) : \tau$. Inversion gives us $\Gamma \vdash v : \sigma_1$, $\Gamma \vdash \gamma_1 : \sigma_1 \sim_R \sigma_2$, and $\Gamma \vdash \gamma_2 : \sigma_2 \sim_R \tau$. Straightforward use of typing rules shows that $\Gamma \vdash v \triangleright (\gamma_1 \mathbin{\&\#} \gamma_2) : \tau$, as desired.

Congruence rules: By induction.

Case S_PUSH: We adopt the variable names from the statement of the rule:

$$\frac{\eta_1 = \mathbf{sym}(\mathbf{nth}^1 \eta_0) \quad \eta_2 = \mathbf{nth}^2 \eta_0 \quad \emptyset \vdash v : \sigma_1 \rightarrow \sigma_2}{(v \triangleright \eta_0) e' \rightarrow v(e' \triangleright \eta_1) \triangleright \eta_2} \text{ S_PUSH}$$

We know that $\Gamma \vdash (v \triangleright \eta_0) e' : \sigma_4$ and we must show $\Gamma \vdash (v(e' \triangleright \eta_1)) \triangleright \eta_2 : \sigma_4$. Inversion tells us that $\Gamma \vdash \eta_0 : (\sigma_1 \rightarrow \sigma_2) \sim_R (\sigma_3 \rightarrow \sigma_4)$ and $\Gamma \vdash e' : \sigma_3$. We can now see that $\Gamma \vdash \eta_1 : \sigma_3 \sim_R \sigma_1$ and $\Gamma \vdash \eta_2 : \sigma_2 \sim_R \sigma_4$. Thus, $\Gamma \vdash e' \triangleright \eta_1 : \sigma_1$ and $\Gamma \vdash v(e' \triangleright \eta_1) \triangleright \eta_2 : \sigma_4$ as desired.

Case S_TPUSH: We adopt the variable names from the statement of the rule:

$$\frac{\emptyset \vdash v : \forall \alpha: \kappa. \sigma' \quad \emptyset \vdash \tau : \kappa}{(v \triangleright \gamma) \tau \rightarrow v \tau \triangleright \gamma @ \tau} \text{ S_TPUSH}$$

We know that $\Gamma \vdash (v \triangleright \gamma) \tau : \tau'$ and we must show that $\Gamma \vdash v \tau \triangleright \gamma @ \tau : \tau'$. Inversion tells us that $\Gamma \vdash \gamma : (\forall \alpha: \kappa. \sigma') \sim_R (\forall \alpha: \kappa. \sigma'')$ where $\tau' = \sigma''[\tau/\alpha]$. We can see that $\Gamma \vdash \gamma @ \tau : \sigma'[\tau/\alpha] \sim_R \sigma''[\tau/\alpha]$ and thus that $\Gamma \vdash v \tau \triangleright \gamma @ \tau : \tau'$ as desired.

Case S_CPUSH: We adopt the variables names from the statement of the rule:

$$\frac{\eta_{11} = \mathbf{nth}^1(\mathbf{nth}^1 \eta_0) \quad \eta_{12} = \mathbf{nth}^2(\mathbf{nth}^1 \eta_0) \quad \eta_2 = \mathbf{nth}^2 \eta \quad \gamma'' = \eta_{11} \mathbin{\&\#} \gamma' \mathbin{\&\#} \mathbf{sym} \eta_{12} \quad \emptyset \vdash v : \sigma_1 \sim_\rho \sigma_2 \Rightarrow \sigma_3 \quad \emptyset \vdash \gamma' : \sigma_4 \sim_\rho \sigma_5}{(v \triangleright \eta_0) \gamma' \rightarrow v \gamma'' \triangleright \eta_2} \text{ S_CPUSH}$$

We know that $\Gamma \vdash (v \triangleright \eta_0) \gamma' : \sigma_6$ and we must show that $\Gamma \vdash v \gamma'' \triangleright \eta_2 : \sigma_6$. Inversion tells us that $\Gamma \vdash \eta_0 :$

$(\sigma_1 \sim_\rho \sigma_2 \Rightarrow \sigma_3) \sim_R (\sigma_4 \sim_\rho \sigma_5 \Rightarrow \sigma_6)$. We can now see the following:

$$\begin{aligned}\Gamma \vdash \eta_{11} : \sigma_1 \sim_\rho \sigma_4 \\ \Gamma \vdash \eta_{12} : \sigma_2 \sim_\rho \sigma_5 \\ \Gamma \vdash \eta_2 : \sigma_3 \sim_\rho \sigma_6 \\ \Gamma \vdash \gamma'' : \sigma_1 \sim_\rho \sigma_2\end{aligned}$$

Thus $\Gamma \vdash v \gamma'' \triangleright \eta_2 : \sigma_6$ as desired.

Case S_KPUSH: We adopt the variable names from the statement of S_KPUSH:

$$\frac{\emptyset \vdash \eta : D\bar{\tau} \sim_R D\bar{\tau}' \quad K : \forall \bar{\alpha}: \bar{\kappa}. \forall \bar{\beta}: \bar{\kappa}'. (\sigma' \sim_\rho \sigma'') \Rightarrow \bar{\tau}'' \rightarrow D\bar{\alpha}}{\emptyset \vdash \gamma : (\sigma' \sim_\rho \sigma'')[\overline{\tau/\alpha}][\overline{\sigma/\beta}]} \quad \frac{\gamma' = \mathbf{sym}(\sigma'[\mathbf{nth} \eta/\bar{\alpha}]) \mathbin{\&\#} \gamma \mathbin{\&\#} \sigma''[\mathbf{nth} \eta/\bar{\alpha}]_\rho}{e' = e \triangleright \tau''[\mathbf{nth} \eta/\bar{\alpha}]_R} \quad \text{S_KPUSH}$$

Inversion gives us the premises of this rule. We also know $\Gamma \vdash (K\bar{\tau}\bar{\sigma}\bar{\gamma}\bar{e}) \triangleright \eta : D\bar{\tau}'$. We must show $\Gamma \vdash (K\bar{\tau}'\bar{\sigma}\bar{\gamma}'\bar{e}') : D\bar{\tau}'$. Note that τ_0 and the $\bar{\alpha}$ do not change, so we need not worry about them here.

Let $\bar{\phi} = (\sigma' \sim_\rho \sigma'')$. From repeated inversion (and induction on the length of $\bar{\tau}$), we can derive

$$\overline{\Gamma \vdash \tau : \kappa}$$

Then, from homogeneity of coercions (Lemma 19) (and more induction on $\bar{\tau}'$), we see that

$$\overline{\Gamma \vdash \tau' : \kappa}$$

Putting this together, we get

$$\Gamma \vdash K\bar{\tau}' : (\forall \bar{\beta}: \bar{\kappa}'. \bar{\phi} \Rightarrow \bar{\tau}'' \rightarrow D\bar{\alpha})[\overline{\tau'/\alpha}]$$

or

$$\Gamma \vdash K\bar{\tau}' : \forall \bar{\beta}: \bar{\kappa}'. \bar{\phi}[\overline{\tau'/\alpha}] \Rightarrow \bar{\tau}''[\overline{\tau'/\alpha}] \rightarrow D\bar{\tau}'$$

Taking $K\bar{\tau}\bar{\sigma}\bar{\gamma}\bar{e}$ apart further (and induction on $\bar{\sigma}$) tells us

$$\overline{\Gamma \vdash \sigma : \kappa'}$$

and thus that

$$\Gamma \vdash K\bar{\tau}'\bar{\sigma} : \bar{\phi}[\overline{\tau'/\alpha}][\overline{\sigma/\beta}] \Rightarrow \bar{\tau}''[\overline{\tau'/\alpha}][\overline{\sigma/\beta}] \rightarrow D\bar{\tau}'[\overline{\sigma/\beta}]$$

But, from $\overline{\Gamma \vdash \tau' : \kappa}$, we see that $\bar{\beta}$ do not appear in $\bar{\tau}'$. So, we have

$$\Gamma \vdash K\bar{\tau}'\bar{\sigma} : \bar{\phi}[\overline{\tau'/\alpha}][\overline{\sigma/\beta}] \Rightarrow \bar{\tau}''[\overline{\tau'/\alpha}][\overline{\sigma/\beta}] \rightarrow D\bar{\tau}'$$

Using techniques similar to that for $\bar{\tau}$ and $\bar{\sigma}$, we can derive the following:

$$\begin{aligned}\overline{\Gamma \vdash \gamma : \phi[\overline{\tau/\alpha}][\overline{\sigma/\beta}]} \\ \overline{\Gamma \vdash e : \tau''[\overline{\tau/\alpha}][\overline{\sigma/\beta}]}\end{aligned}$$

We need to conclude the following:

$$\begin{aligned}\overline{\Gamma \vdash \gamma' : \phi[\overline{\tau'/\alpha}][\overline{\sigma/\beta}]} \\ \overline{\Gamma \vdash e' : \tau''[\overline{\tau'/\alpha}][\overline{\sigma/\beta}]}\end{aligned}$$

We wish to use the lifting lemma (Lemma 21) to get types for $\sigma'[\mathbf{nth} \eta/\bar{\alpha}]_\rho$ and $\sigma''[\mathbf{nth} \eta/\bar{\alpha}]_\rho$. So, we must first establish the premises of the lifting lemma.

1. $\Gamma \vdash \eta : D\bar{\tau} \sim_R D\bar{\tau}'$, from the inversion on S_KPUSH (and weakening to change the context);

2. $\overline{\Gamma \vdash \tau : \kappa}$, as above;
3. $\overline{\Gamma \vdash \tau' : \kappa}$, as above;
4. D is not a **newtype**: by choice of metavariable.
5. $\overline{\Omega \vdash \sigma' : \rho}$ and $\overline{\Omega \vdash \sigma'' : \rho}$: Here, $\Omega = \overline{\beta' : \mathbb{N}, \bar{\alpha} : \text{roles}(D)}$ where β' are the type variables bound in Γ , along with the existential variables $\bar{\beta}$. (That is, the Γ' in the statement of the lifting lemma is $\overline{\beta : \kappa'}$.) By ROLES_DATA, we can see that $\overline{\Omega \vdash (\sigma' \rightsquigarrow_{\rho} \sigma'') : \mathbb{R}}$. This can be established by either RTY_TYCONAPP or by RTY_APP. In the former case, we get the desired outcome by looking at ROLES_EQUALITY. In the latter case, we see that $\Omega \vdash \sigma'_i : \mathbb{N}$ or $\Omega \vdash \sigma''_i : \mathbb{N}$ and then use role subsumption (Lemma 5).
6. $\Gamma, \bar{\alpha} : \bar{\kappa}, \bar{\beta} : \kappa' \vdash \sigma' : \kappa''$ and the same for σ'' : This comes from the well-formedness of the global context, including the type of K .
7. $\bar{\beta} : \kappa'$ must contain only type variable bindings: It sure does.

Now, we can conclude

$$\frac{\overline{\Gamma, \beta : \kappa' \vdash \sigma' [\mathbf{nth} \eta / \bar{\alpha}]_{\rho} : \sigma' [\tau / \bar{\alpha}] \rightsquigarrow_{\rho} \sigma' [\tau' / \bar{\alpha}]}}{\overline{\Gamma, \beta : \kappa' \vdash \sigma'' [\mathbf{nth} \eta / \bar{\alpha}]_{\rho} : \sigma'' [\tau / \bar{\alpha}] \rightsquigarrow_{\rho} \sigma'' [\tau' / \bar{\alpha}]}}$$

We then do type variable substitution to get

$$\frac{\overline{\Gamma \vdash \sigma' [\mathbf{nth} \eta / \bar{\alpha}]_{\rho} [\sigma / \beta] : \sigma' [\tau / \bar{\alpha}] [\sigma / \beta] \rightsquigarrow_{\rho} \sigma' [\tau' / \bar{\alpha}] [\sigma / \beta]}}{\overline{\Gamma \vdash \sigma'' [\mathbf{nth} \eta / \bar{\alpha}]_{\rho} [\sigma / \beta] : \sigma'' [\tau / \bar{\alpha}] [\sigma / \beta] \rightsquigarrow_{\rho} \sigma'' [\tau' / \bar{\alpha}] [\sigma / \beta]}}$$

Now, by CO_TRANS, we can conclude

$$\overline{\Gamma \vdash \gamma' : \phi [\tau' / \bar{\alpha}] [\sigma / \beta]}$$

as desired.

To type the e' , we need to apply the lifting lemma once again, this time to $\tau'' [\mathbf{nth} \eta / \bar{\alpha}]_{\mathbb{R}}$. Much of our work at establishing premises carries over, except for these:

5. $\overline{\Omega \vdash \tau'' : \mathbb{R}}$ (with Ω as above): This comes directly from the premises of ROLES_DATA, noting that τ'' appears in as an argument type to K .
6. $\Gamma, \bar{\alpha} : \bar{\kappa}, \bar{\beta} : \kappa' \vdash \tau'' : \kappa''$: This comes from the well-formedness of the global context, including the type of K .

We then apply the lifting lemma to conclude that

$$\overline{\Gamma, \beta : \kappa' \vdash \tau'' [\mathbf{nth} \gamma / \bar{\alpha}]_{\mathbb{R}} : \tau'' [\tau / \bar{\alpha}] \rightsquigarrow_{\mathbb{R}} \tau'' [\tau' / \bar{\alpha}]}$$

We use type variable substitution to get

$$\overline{\Gamma \vdash \tau'' [\mathbf{nth} \gamma / \bar{\alpha}]_{\mathbb{R}} [\sigma / \beta] : \tau'' [\tau / \bar{\alpha}] [\sigma / \beta] \rightsquigarrow_{\mathbb{R}} \tau'' [\tau' / \bar{\alpha}] [\sigma / \beta]}$$

We can then conclude

$$\overline{\Gamma \vdash e' : \tau'' [\tau' / \bar{\alpha}] [\sigma / \beta]}$$

as desired.

Putting this all together, we see that $\Gamma \vdash K \bar{\tau} \bar{\sigma} \bar{\gamma}' \bar{e}' : D \bar{\tau}$ as originally desired, and we are done. \square

H. Progress

H.1 Consistency

Definition 23 (Type consistency). *Two types τ_1 and τ_2 are consistent if, whenever they are both value types:*

1. If $\tau_1 = H \bar{\sigma}$, then $\tau_2 = H \bar{\sigma}'$;

2. If $\tau_1 = \forall \alpha : \kappa. \sigma$ then $\tau_2 = \forall \alpha : \kappa. \sigma'$.

Note that if either τ_1 or τ_2 is *not* a value type, then they are vacuously consistent. Also, recall that a type headed by a **newtype** is not a value type.

Definition 24 (Context consistency). *The global context is consistent if, whenever $\emptyset \vdash \gamma : \tau_1 \rightsquigarrow_{\mathbb{R}} \tau_2$, τ_1 and τ_2 are consistent.*

In order to prove consistency, we define a type reduction relation $\tau \rightsquigarrow_{\rho} \sigma$, show that the relation preserves value type heads, and then show that any well-typed coercion corresponds to a path in the rewrite relation.

Here is the type rewrite relation:

$$\boxed{\tau \rightsquigarrow_{\rho} \sigma} \quad \text{Type reduction}$$

$$\frac{}{\tau \rightsquigarrow_{\rho} \tau} \quad \text{RED_REFL}$$

$$\frac{\tau_1 \rightsquigarrow_{\rho} \sigma_1 \quad \tau_2 \rightsquigarrow_{\mathbb{N}} \sigma_2}{\tau_1 \tau_2 \rightsquigarrow_{\rho} \sigma_1 \sigma_2} \quad \text{RED_APP}$$

$$\frac{\overline{\tau \rightsquigarrow_{\rho} \bar{\sigma}} \quad \overline{\bar{\rho} \text{ is a prefix of roles}(H)}}{H \bar{\tau} \rightsquigarrow_{\mathbb{R}} H \bar{\sigma}} \quad \text{RED_TYCONAPP}$$

$$\frac{\tau \rightsquigarrow_{\rho} \sigma}{\forall \alpha : \kappa. \tau \rightsquigarrow_{\rho} \forall \alpha : \kappa. \sigma} \quad \text{RED_FORALL}$$

$$\frac{\overline{\tau \rightsquigarrow_{\mathbb{N}} \bar{\sigma}}}{F(\bar{\tau}) \rightsquigarrow_{\rho} F(\bar{\sigma})} \quad \text{RED_TYFAM}$$

$$\frac{C : [\bar{\alpha} : \bar{\kappa}]. \tau_1 \rightsquigarrow_{\rho} \tau_2 \quad \rho \leq \rho'}{\tau_1 [\sigma / \alpha] \rightsquigarrow_{\rho'} \tau_2 [\sigma / \alpha]} \quad \text{RED_AXIOM}$$

$$\frac{}{\tau \rightsquigarrow_{\mathbb{P}} \sigma} \quad \text{RED_PHANTOM}$$

Lemma 25 (Simple rewrite substitution). *If $\tau_1 \rightsquigarrow_{\rho} \tau_2$, then $\tau_1 [\sigma / \alpha] \rightsquigarrow_{\rho} \tau_2 [\sigma / \alpha]$.*

Proof. By straightforward induction, noting that axioms have no free variables. \square

Lemma 26 (Rewrite substitution). *Let $\bar{\alpha}$ be the free variables in a type σ . If $\bar{\alpha} : \bar{\rho} \vdash \sigma : \mathbb{R}$:*

1. If $\overline{\tau \rightsquigarrow_{\rho} \tau'}$, then $\sigma [\tau / \bar{\alpha}] \rightsquigarrow_{\mathbb{R}} \sigma [\tau' / \bar{\alpha}]$;
2. If $\overline{\tau \rightsquigarrow_{\mathbb{N}} \tau'}$, then $\sigma [\tau / \bar{\alpha}] \rightsquigarrow_{\mathbb{N}} \sigma [\tau' / \bar{\alpha}]$.

Proof. Let $\Omega = \bar{\alpha} : \bar{\rho}$. Proceed by induction on the structure of σ .

Case $\sigma = \alpha$: There is thus only one free variable, α in σ . The one role ρ is \mathbb{R} . For clause (1), we know $\tau \rightsquigarrow_{\mathbb{R}} \tau'$, so we are done. For clause (2), we know $\tau \rightsquigarrow_{\mathbb{N}} \tau'$, so we are done.

Case $\sigma = \sigma_1 \sigma_2$:

Case (σ can be written as $H \bar{\sigma}$): Here, we assume that the length of $\bar{\sigma}$ is at most the length of $\text{roles}(H)$. If this is not the case, fall through to the “otherwise” case.

Clause (1): We know $\overline{\tau \rightsquigarrow_{\rho} \tau'}$. We must show that $H \bar{\sigma} [\tau / \bar{\alpha}] \rightsquigarrow_{\mathbb{R}} H \bar{\sigma} [\tau' / \bar{\alpha}]$. We will use RED_TYCONAPP. Let $\bar{\rho}'$ be a prefix of $\text{roles}(H)$ of the same length as $\bar{\sigma}$. We must show $\sigma [\tau / \bar{\alpha}] \rightsquigarrow_{\rho'} \sigma [\tau' / \bar{\alpha}]$.

Fix i . We will show that $\sigma_i [\tau / \bar{\alpha}] \rightsquigarrow_{\rho'_i} \sigma_i [\tau' / \bar{\alpha}]$.

Case ($\rho'_i = N$): In order to use the induction hypothesis, we must show that for every j such that α_j appears free in σ_i , $\rho_j = N$. To use Lemma 4, we must establish that $\Omega \vdash \sigma_i : N$. We can get this by inversion on $\Omega \vdash H\bar{\sigma} : R$ – whether by `RTY_TYCONAPP` or by `RTY_APP`, we get $\Omega \vdash \sigma_i : N$. So, we can use the induction hypothesis and we are done.

Case ($\rho'_i = R$): Inverting $\bar{\alpha};\bar{\rho} \vdash H\bar{\sigma} : R$ gives us two possibilities:

Case `RTY_TYCONAPP`: Here, we see $\overline{\Omega \vdash \sigma : \rho'}$, and thus, that $\Omega \vdash \sigma_i : R$ (because $\rho'_i = R$). We can then use the induction hypothesis (and using Lemma 3 to make the contexts line up) and we are done.

Case `RTY_APP`: We invert repeatedly, and we either get $\Omega \vdash \sigma_i : N$ or $\Omega \vdash \sigma_i : \rho'_i$, depending on whether we hit a `RTY_TYCONAPP` during the inversions. In the second case, we proceed as above (the `RTY_TYCONAPP` case). In the first case, we use Lemma 5 to conclude $\Omega \vdash \sigma_i : R$ and use the induction hypothesis.

Case ($\rho'_i = P$): We are done by `RED_PHANTOM`.

Clause (2): We know that $\bar{\tau} \rightsquigarrow_N \bar{\tau}'$. We must show that $H\bar{\sigma}[\bar{\tau}/\bar{\alpha}] \rightsquigarrow_N H\bar{\sigma}[\bar{\tau}'/\bar{\alpha}]$. It is easier to consider the original type σ just as $\sigma_1 \sigma_2$, not as $H\bar{\sigma}$; fall through to the next case.

Otherwise:

Clause (1): We know $\bar{\tau} \rightsquigarrow_\rho \bar{\tau}'$ and need to show that $(\sigma_1 \sigma_2)[\bar{\tau}/\bar{\alpha}] \rightsquigarrow_R (\sigma_1 \sigma_2)[\bar{\tau}'/\bar{\alpha}]$. The fact $\Omega \vdash \sigma_1 \sigma_2 : R$ must be by `RTY_APP`. So, we can conclude $\Omega \vdash \sigma_1 : R$ and $\Omega \vdash \sigma_2 : N$. Then, we can use the induction hypothesis to get $\sigma_1[\bar{\tau}/\bar{\alpha}] \rightsquigarrow_R \sigma_1[\bar{\tau}'/\bar{\alpha}]$. To use the induction hypothesis for σ_2 , we must first establish that, for every j such that α_j appears free in σ_2 , $\tau_j \rightsquigarrow_N \tau'_j$. Lemma 4 provides exactly this information, so we get $\sigma_2[\bar{\tau}/\bar{\alpha}] \rightsquigarrow_N \sigma_2[\bar{\tau}'/\bar{\alpha}]$. We are done by `RED_APP`.

Clause (2): We know $\bar{\tau} \rightsquigarrow_N \bar{\tau}'$ and need to show that $(\sigma_1 \sigma_2)[\bar{\tau}/\bar{\alpha}] \rightsquigarrow_N (\sigma_1 \sigma_2)[\bar{\tau}'/\bar{\alpha}]$. We simply use induction to get:

$$\begin{aligned} \sigma_1[\bar{\tau}/\bar{\alpha}] &\rightsquigarrow_N \sigma_1[\bar{\tau}'/\bar{\alpha}] \\ \sigma_2[\bar{\tau}/\bar{\alpha}] &\rightsquigarrow_N \sigma_2[\bar{\tau}'/\bar{\alpha}] \end{aligned}$$

We are done by `RED_APP`.

Case $\sigma = H$: We are done by `RED_REFL`.

Case $\sigma = \forall \beta:\kappa.\sigma'$: We assume that we have renamed variables so that $\beta \notin \bar{\alpha}$. We see that inverting $\Omega \vdash \forall \beta:\kappa.\sigma' : R$ gives us $\Omega, \beta:N \vdash \sigma' : R$, where $\bar{\alpha}, \beta$ are the free variables in σ' . We can then use the induction hypothesis and we are done by `RED_FORALL`.

Case $\sigma = F(\bar{\sigma})$: Inversion on $\Omega \vdash F(\bar{\sigma}) : R$ gives us $\overline{\Omega \vdash \bar{\sigma} : N}$. We can then apply Lemma 4 to see that $\bar{\rho} = N$. We then use the induction hypothesis repeatedly to get

$$\overline{\sigma[\bar{\tau}/\bar{\alpha}] \rightsquigarrow_N \sigma[\bar{\tau}'/\bar{\alpha}]}$$

We are now done by `RED_TYFAM`. \square

Lemma 27 (Sub-roling in the rewrite relation). *If $\tau_1 \rightsquigarrow_N \tau_2$, then $\tau_1 \rightsquigarrow_\rho \tau_2$.*

Proof. By straightforward induction on $\tau_1 \rightsquigarrow_N \tau_2$. \square

Lemma 28 (`RED_APP/RED_TYCONAPP`). *If $H\bar{\tau} \tau' \rightsquigarrow_R H\bar{\sigma} \sigma'$ by `RED_APP`, the length of $\bar{\tau}$ is less than the length of $\text{roles}(H)$, then $H\bar{\tau} \tau' \rightsquigarrow_R H\bar{\sigma} \sigma'$ also by `RED_TYCONAPP`.*

Proof. Fix H . We then proceed by induction on the length of $\bar{\tau}$.

Base case ($H\bar{\tau} \rightsquigarrow_R H\sigma'$): The premises of `RED_APP` give us $H \rightsquigarrow_R H$ and $\tau' \rightsquigarrow_N \sigma'$. Regardless of $\text{roles}(H)$, we can use the sub-roling lemma (Lemma 27) to show $\tau' \rightsquigarrow_\rho \sigma'$ and we are done. (In the case where $\text{roles}(H)$ is empty, an assumption is violated, and we are done anyway.)

Inductive case: Our inductive hypothesis says: if $H\bar{\tau} \rightsquigarrow_R H\bar{\sigma}$ and $\tau' \rightsquigarrow_N \sigma'$ (and the length of $\text{roles}(H)$ is sufficient), then $\bar{\tau} \rightsquigarrow_\rho \bar{\sigma}$ and $\tau' \rightsquigarrow_{\rho_i} \sigma'$, where $i = (\text{length of } \bar{\tau}) + 1$. We must show that, if $H\bar{\tau} \tau' \rightsquigarrow_R H\bar{\sigma} \sigma'$ and $\tau'' \rightsquigarrow_N \sigma''$ (and the length of $\text{roles}(H)$ is sufficient) then $\bar{\tau} \rightsquigarrow_\rho \bar{\sigma}$, $\tau' \rightsquigarrow_{\rho_i} \sigma'$, and $\tau'' \rightsquigarrow_{\rho_j} \sigma''$ (where $j = i + 1$).

Inverting $H\bar{\tau} \tau' \rightsquigarrow_R H\bar{\sigma} \sigma'$ gives us several possibilities:

Case `RED_REFL`: We get $\bar{\tau} \rightsquigarrow_\rho \bar{\sigma}$ and $\tau' \rightsquigarrow_{\rho_i} \sigma'$ by `RED_REFL`. We get $\tau'' \rightsquigarrow_{\rho_j} \sigma''$ by Lemma 27.

Case `RED_APP`: We get our first two desiderata from use of the induction hypothesis and our last from Lemma 27.

Case `RED_TYCONAPP`: Our first two desiderata come from the premises of `RED_TYCONAPP`, and the last one comes from Lemma 27.

Case `RED_AXIOM`: This case is impossible, because there can be only one newtype axiom for a newtype, and its arity is greater than $(\text{length of } \bar{\tau}) + 1$. \square

Lemma 29 (Pattern). *Let $\bar{\alpha}$ be the free variables in a type τ . We require that each variable α is mentioned exactly once in τ and that no type families appear in τ . Then, if, for some $\bar{\sigma}$, $\tau[\bar{\sigma}/\bar{\alpha}] \rightsquigarrow_N \tau'$, then there exist $\bar{\sigma}'$ such that $\tau' = \tau[\bar{\sigma}'/\bar{\alpha}]$ and $\bar{\sigma} \rightsquigarrow_N \bar{\sigma}'$.*

Proof. We proceed by induction on the structure of τ .

Case $\tau = \alpha$: There is just one free variable (α), and thus just one type σ . We have $\sigma \rightsquigarrow_N \tau'$. Let $\sigma' = \tau'$ and we are done.

Case $\tau = \tau_1 \tau_2$: Partition the free variables into a list $\bar{\beta}_1$ that appear in τ_1 and $\bar{\beta}_2$ that appear in τ_2 . This partition must be possible by assumption. Similarly, partition $\bar{\sigma}$ into $\bar{\sigma}_1$ and $\bar{\sigma}_2$. We can see that $\tau_1[\bar{\sigma}_1/\bar{\beta}_1] \tau_2[\bar{\sigma}_2/\bar{\beta}_2] \rightsquigarrow_N \tau'$. Thus must be by `RED_APP` (noting that all newtype axioms are at role R). Thus, $\tau' = \tau'_1 \tau'_2$ and $\tau_1[\bar{\sigma}_1/\bar{\beta}_1] \rightsquigarrow_N \tau'_1$ and $\tau_2[\bar{\sigma}_2/\bar{\beta}_2] \rightsquigarrow_N \tau'_2$. We then use the induction hypothesis to get $\bar{\sigma}'_1$ and $\bar{\sigma}'_2$ such that $\tau'_1 = \tau_1[\bar{\sigma}'_1/\bar{\beta}_1]$ and $\tau'_2 = \tau_2[\bar{\sigma}'_2/\bar{\beta}_2]$. We conclude that $\bar{\sigma}'$ is the combination of $\bar{\sigma}'_1$ and $\bar{\sigma}'_2$, undoing the partition done earlier.

Case $\tau = H$: Trivial.

Case $\tau = \forall \beta:\kappa.\tau_0$: We first note that, according to the definition of $\bar{\alpha}$, $\beta \notin \bar{\alpha}$. We wish to use the induction hypothesis, but we must be careful because τ_0 may mention β multiple times. So, we linearise τ_0 into $\bar{\tau}'_0$, replacing every occurrence of β with fresh variables $\bar{\beta}'$. (Note that $\bar{\beta}'$ can be empty.) We know that $(\forall \beta:\kappa.\tau_0)[\bar{\sigma}/\bar{\alpha}] \rightsquigarrow_N \tau'$. We note that $(\forall \beta:\kappa.\tau_0)[\bar{\sigma}/\bar{\alpha}] = \forall \beta:\kappa.(\tau_0[\bar{\sigma}/\bar{\alpha}]) = \forall \beta:\kappa.(\tau'_0[\bar{\sigma}/\bar{\alpha}][\bar{\beta}/\bar{\beta}'])$.

(We have abused notation somewhat in the second substitution. There is only one β ; it is substituted for every variable in $\bar{\beta}'$.) Let $\bar{\sigma}''$ be $\bar{\sigma}$ appended with the right number of copies of β . Let $\bar{\alpha}'$ be $\bar{\alpha}$ appended with β' . Then, we can say $\forall \beta:\kappa.(\tau_0'[\bar{\sigma}''/\bar{\alpha}']) \rightsquigarrow_{\mathbb{N}} \tau'$. We invert to get that $\tau' = \forall \beta:\kappa.\tau''$ and $\tau_0'[\bar{\sigma}''/\bar{\alpha}'] \rightsquigarrow_{\mathbb{N}} \tau''$. We can now use the induction hypothesis to get $\bar{\sigma}'''$ such that $\tau' = \tau[\bar{\sigma}'''/\bar{\alpha}']$ and $\bar{\sigma}'' \rightsquigarrow_{\mathbb{N}} \bar{\sigma}'''$. But, we can see that, β steps only to itself. Thus, the last entries in $\bar{\sigma}'''$ must be the same list of β s that $\bar{\sigma}''$ has. We let σ' be the prefix of $\bar{\sigma}'''$ without the β s, and we are done.

Case $\tau = F(\bar{\tau})$: Impossible, by assumption. \square

Lemma 30 (Patterns). *Let $\bar{\alpha}$ be the free variables in a list of types $\bar{\tau}$. Assume each variable α is mentioned exactly once in $\bar{\tau}$ and that no type families appear in $\bar{\tau}$. If, for some $\bar{\sigma}, \tau[\bar{\sigma}/\bar{\alpha}] \rightsquigarrow_{\mathbb{N}} \tau'$, then there exist $\bar{\sigma}'$ such that $\bar{\tau}' = \tau[\bar{\sigma}'/\bar{\alpha}]$ and $\bar{\sigma} \rightsquigarrow_{\mathbb{N}} \bar{\sigma}'$.*

Proof. By induction on the length of $\bar{\tau}$.

Base case: Trivial.

Inductive case: We partition and recombine variables as in the $\tau_1 \tau_2$ case in the previous proof and proceed by induction. \square

Lemma 31 (Local diamond). *If $\tau \rightsquigarrow_{\rho} \sigma_1$ and $\tau \rightsquigarrow_{\rho} \sigma_2$, then there exists σ_3 such that $\sigma_1 \rightsquigarrow_{\rho} \sigma_3$ and $\sigma_2 \rightsquigarrow_{\rho} \sigma_3$.*

Proof. If $\rho = \text{P}$, then the result is trivial, by RED_PHANTOM. So, we assume $\rho \neq \text{P}$.

If $\sigma_1 = \tau$ or $\sigma_2 = \tau$, the result is trivial. So, we assume that neither reduction is by RED_REFL.

By induction on the structure of τ :

Case $\tau = \alpha$: We note that the left-hand side of an axiom can never be a bare variable, and so the only possibility of stepping is by RED_REFL. We are done.

Case $\tau = \tau_1 \tau_2$: Suppose $\rho = \text{N}$. All axioms at nominal role have a type family application on their left-hand side, so RED_AXIOM cannot apply. Thus, only RED_APP can be used, and we are done by induction.

Now, we can assume $\rho = \text{R}$. If $\tau_1 \tau_2$ cannot be rewritten as $H\bar{\tau}$ (for some H and some $\bar{\tau}$), then the only applicable rule is RED_APP (noting that relevant axiom left-hand sides can indeed be written as $H\bar{\tau}$) and we are done by induction.

So, we now rewrite τ as $H\bar{\tau}_0$. There are six possible choices of the two reductions, among RED_APP, RED_TYCONAPP, and RED_AXIOM. We handle each case separately:

Case RED_APP/RED_APP: We are done by induction.

Case RED_APP/RED_TYCONAPP: We apply Lemma 28 and finish by induction.

Case RED_APP/RED_AXIOM: Rewrite $\sigma_1 = \sigma_{11} \sigma_{12}$. We know then that $\tau_1 \rightsquigarrow_{\mathbb{R}} \sigma_{11}$ and $\tau_2 \rightsquigarrow_{\mathbb{N}} \sigma_{12}$. (Recall that $\tau_1 \tau_2 = \tau = H\bar{\tau}_0$.) We also know that $H\bar{\tau}_0 \rightsquigarrow_{\mathbb{R}} \sigma_2$ by a newtype axiom $C : [\bar{\alpha}:\bar{\kappa}].H\bar{\alpha} \sim_{\mathbb{R}} \sigma_0$, where $\sigma_2 = \sigma_0[\bar{\tau}_0/\bar{\alpha}]$.

By induction we can discover that σ_{11} has the form $H\bar{\sigma}$ – we know that τ_1 cannot reduce by RED_AXIOM because the well-formedness of the global context says

that newtype axioms are unique, and the axiom used on τ has a higher arity than any axiom that could be used on τ_1 . Thus, $\sigma_1 = H\bar{\sigma} \sigma_{12}$. The same axiom C applies here. Let $\bar{\sigma}' = \bar{\sigma}, \sigma_{12}$. So, we can step σ_1 to $\sigma_3 = \sigma_0[\bar{\sigma}'/\bar{\alpha}]$ by RED_AXIOM.

Now, we must show $\sigma_2 \rightsquigarrow_{\mathbb{R}} \sigma_3$. We wish to apply the rewrite-substitution lemma (Lemma 26). We must show that $\bar{\tau}_0 \rightsquigarrow_{\rho} \bar{\sigma}'$, where $\bar{\alpha}:\bar{\rho} \vdash \sigma_0 : \text{R}$. This last fact is exactly what appears in the premise to ROLES_NEWTYPE (which, in turn, is guaranteed by the well-formedness of the global context). Now, we know $\tau = H\bar{\tau}_0$ and $\sigma_1 = H\bar{\sigma}'$, and that $\tau \rightsquigarrow_{\mathbb{R}} \sigma_1$ by RED_APP. We also know that an axiom is applicable to τ . Thus, the length of $\bar{\tau}$ must be the length of $\text{roles}(H)$, by context well-formedness. So, we can use Lemma 28 to get $\bar{\tau}_0 \rightsquigarrow_{\rho} \bar{\sigma}'$, as desired. We then apply Lemma 26 to conclude $\sigma_2 \rightsquigarrow_{\mathbb{R}} \sigma_3$, and we are done.

Case RED_TYCONAPP/RED_TYCONAPP: We are done by induction.

Case RED_TYCONAPP/RED_AXIOM: We see that $\sigma_1 = H\bar{\sigma}'$ where $\bar{\rho}$ is a prefix of $\text{roles}(H)$ and $\bar{\tau}_0 \rightsquigarrow_{\rho} \bar{\sigma}'$. We also see that $C : [\bar{\alpha}:\bar{\kappa}].H\bar{\alpha} \sim_{\mathbb{R}} \sigma_0$ and that $\sigma_2 = \sigma_0[\bar{\tau}_0/\bar{\alpha}]$.

Let $\sigma_3 = \sigma_0[\bar{\sigma}'/\bar{\alpha}]$. We can see that $\sigma_1 \rightsquigarrow_{\mathbb{R}} \sigma_3$ by RED_AXIOM. And, by Lemma 26 (the rewrite-substitution lemma), we see that $\sigma_2 \rightsquigarrow_{\mathbb{R}} \sigma_3$. So, we are done.

Case RED_AXIOM/RED_AXIOM: Consider the possibility that the two reductions are by different axioms. This would violate context well-formedness, so it is impossible. Thus, we can assume that the axiom used in both reductions is the same: $C : [\bar{\alpha}:\bar{\kappa}].H\bar{\alpha} \sim_{\mathbb{R}} \sigma_0$. The only way that σ_1 and σ_2 can be different is if the types substituted in the rule conclusion ($\bar{\sigma}$) are different in the two different reductions. Suppose then that we have $\bar{\sigma}$ and $\bar{\sigma}'$ so that $\sigma_1 = \sigma_0[\bar{\sigma}/\bar{\alpha}]$ and $\sigma_2 = \sigma_0[\bar{\sigma}'/\bar{\alpha}]$. It must be that $\tau = H\bar{\sigma}$ and that $\tau = H\bar{\sigma}'$. But, this tells us that $\bar{\sigma} = \bar{\sigma}'$ and thus that $\sigma_1 = \sigma_2$. We are done.

Case $\tau = H$: The only non-trivial step H can make is by RED_AXIOM. However, given that only one axiom for a newtype can exist, both steps must step to the same type, so we are done.

Case $\tau = \forall \alpha:\kappa.\tau'$: We are done by induction.

Case $\tau = F(\bar{\tau})$: Here, two rules may apply. We handle the different possibilities separately:

Case RED_TYFAM/RED_TYFAM: We are done by induction.

Case RED_TYFAM/RED_AXIOM: Here, we know that $\sigma_1 = F(\bar{\sigma})$ where $\bar{\tau} \rightsquigarrow_{\mathbb{N}} \bar{\sigma}$, and that $\sigma_2 = \sigma_0[\bar{\sigma}'/\bar{\alpha}]$ where $C : [\bar{\alpha}:\bar{\kappa}].F(\bar{\tau}') \sim_{\mathbb{N}} \sigma_0$ and $\bar{\tau} = \tau'[\bar{\sigma}'/\bar{\alpha}]$.

We wish to use RED_AXIOM to reduce $F(\bar{\sigma})$. We apply Lemma 30 to get $\bar{\sigma}''$ such that $\bar{\sigma} = \tau'[\bar{\sigma}''/\bar{\alpha}]$ and $\bar{\sigma}' \rightsquigarrow_{\mathbb{N}} \bar{\sigma}''$. We then use RED_AXIOM to get $\sigma_1 \rightsquigarrow_{\mathbb{N}} \sigma_3$, where $\sigma_3 = \sigma_0[\bar{\sigma}''/\bar{\alpha}]$. Now, we must show that $\sigma_2 \rightsquigarrow_{\mathbb{N}} \sigma_3$. This comes directly from Lemma 26, and we are done.

Case RED_AXIOM/RED_AXIOM:

We have $C_1 : [\bar{\alpha}:\bar{\kappa}].F(\bar{\tau}_1) \sim_{\mathbb{N}} \sigma_1'$ and $C_2 : [\bar{\beta}:\bar{\kappa}'].F(\bar{\tau}_2) \sim_{\mathbb{N}} \sigma_2'$.

We also know that $\tau = F(\bar{\tau}_1)[\bar{\sigma}'/\bar{\alpha}]$ and $\tau = F(\bar{\tau}_2)[\bar{\sigma}''/\bar{\beta}']$.

Thus, $F(\bar{\tau}_1)[\bar{\sigma}'/\bar{\alpha}] = F(\bar{\tau}_2)[\bar{\sigma}''/\bar{\beta}']$. Thus, $[\bar{\sigma}', \bar{\sigma}''/\bar{\alpha}, \bar{\beta}']$ is a unifier for $F(\bar{\tau}_1)$ and $F(\bar{\tau}_2)$. Thus, by context well-formedness, we have $\sigma_1'[\bar{\sigma}'/\bar{\alpha}] = \sigma_2'[\bar{\sigma}''/\bar{\beta}']$. But,

$\sigma_1 = \sigma'_1[\overline{\sigma'/\alpha}]$ and $\sigma_2 = \sigma'_2[\overline{\sigma''/\beta}]$, and so $\sigma_1 = \sigma_2$ and we are done. \square

Let the notation $\tau_1 \Leftrightarrow_\rho \tau_2$ mean that there exists a σ such that $\tau_1 \rightsquigarrow_\rho^* \sigma$ and $\tau_2 \rightsquigarrow_\rho^* \sigma$.

Lemma 32 (Confluence). *The rewrite relation \rightsquigarrow_ρ is confluent. That is, if $\tau \rightsquigarrow_\rho^* \sigma_1$ and $\tau \rightsquigarrow_\rho^* \sigma_2$, then $\sigma_1 \Leftrightarrow_\rho \sigma_2$.*

Proof. Confluence is a consequence of the local diamond property, Lemma 31. \square

Lemma 33 (Stepping preserves value type heads). *If τ_1 is a value type and $\tau_1 \rightsquigarrow_R \tau_2$, then τ_2 has the same head as τ_1 .*

Proof. By induction, noting that the left-hand side of well-formed axioms are never value types. \square

Lemma 34 (Rewrite relation consistency). *If $\tau_1 \Leftrightarrow_R \tau_2$, then τ_1 and τ_2 are consistent.*

Proof. If either τ_1 or τ_2 is not a value type, then we are trivially done. So, we assume τ_1 and τ_2 are value types. By assumption, there exists σ such that $\tau_1 \rightsquigarrow_R^* \sigma$ and $\tau_2 \rightsquigarrow_R^* \sigma$. By induction over the length of these reductions and the use of Lemma 33, we can see that σ must have the same head as both τ_1 and τ_2 . Thus, τ_1 and τ_2 have the same head, and are thus consistent. \square

Lemma 35 (Completeness of the rewrite relation). *If $\emptyset \vdash \gamma : \tau_1 \sim_\rho \tau_2$, then $\tau_1 \Leftrightarrow_\rho \tau_2$.*

Proof. By induction on $\emptyset \vdash \gamma : \tau_1 \sim_\rho \tau_2$.

Case CO_REFL: Trivial, as \Leftrightarrow_ρ is manifestly reflexive.

Case CO_SYM: By induction, as \Leftrightarrow_ρ is manifestly symmetric.

Case CO_TRANS: We adopt the variable names in the statement of the rule:

$$\frac{\Gamma \vdash \gamma_1 : \tau_1 \sim_\rho \tau_2 \quad \Gamma \vdash \gamma_2 : \tau_2 \sim_\rho \tau_3}{\Gamma \vdash \gamma_1 \circ \gamma_2 : \tau_1 \sim_\rho \tau_3} \text{CO_TRANS}$$

By induction, we know $\tau_1 \Leftrightarrow_\rho \tau_2$ and $\tau_2 \Leftrightarrow_\rho \tau_3$. Thus, we must find σ_{13} such that $\tau_1 \rightsquigarrow_\rho^* \sigma_{13}$ and $\tau_3 \rightsquigarrow_\rho^* \sigma_{13}$. Note that there must be σ_{12} with $\tau_1 \rightsquigarrow_\rho^* \sigma_{12}$ and $\tau_2 \rightsquigarrow_\rho^* \sigma_{12}$, and there must be σ_{23} with $\tau_2 \rightsquigarrow_\rho^* \sigma_{23}$ and $\tau_3 \rightsquigarrow_\rho^* \sigma_{23}$. Thus, we can use Lemma 32 (confluence) to find a σ_{13} such that $\sigma_{12} \rightsquigarrow_\rho^* \sigma_{13}$ and $\sigma_{23} \rightsquigarrow_\rho^* \sigma_{13}$. By transitivity of \rightsquigarrow_ρ^* , we are done. \square

Case CO_TYCONAPP: We know by induction that $\overline{\tau} \Leftrightarrow_\rho \overline{\sigma}$.

Let the list of common reducts be $\overline{\tau}'$. We can see that $H\overline{\tau} \rightsquigarrow_R^* H\overline{\tau}'$ by repeated use of RED_TYCONAPP, and similarly for $H\overline{\sigma} \rightsquigarrow_R^* H\overline{\tau}'$. Thus $H\overline{\tau}'$ is our common reduct and we are done.

Case CO_TYFAM: We are done by induction and repeated use of RED_TYFAM.

Case CO_APP: We are done by induction and repeated use of RED_APP.

Case CO_FORALL: We are done by induction and repeated use of RED_FORALL.

Case CO_PHANTOM: We are done by RED_PHANTOM.

Case CO_VAR: Not possible, as the context is empty. \square

Case CO_AXIOM: We are done by RED_AXIOM.

Case CO_NTH: We adopt the variable names in the rule:

$$\frac{\Gamma \vdash \gamma : H\overline{\tau} \sim_R H\overline{\sigma} \quad \overline{\rho} \text{ is a prefix of } \text{roles}(H) \quad H \text{ is not a newtype}}{\Gamma \vdash \text{nth}^i \gamma : \tau_i \sim_{\rho_i} \sigma_i} \text{CO_NTH}$$

We know by induction that $H\overline{\tau} \Leftrightarrow_R H\overline{\sigma}$. In other words, there exists some τ_0 such that $H\overline{\tau} \rightsquigarrow_R^* \tau_0$ and $H\overline{\sigma} \rightsquigarrow_R^* \tau_0$. We can see by induction on the number of steps in the derivation (and a nested induction in the RED_APP case) that τ_0 must have the form $H\overline{\tau}'$ for some $\overline{\tau}'$. In particular, note that no axioms can apply because H is not a newtype. Thus, each step is from either RED_APP or from RED_TYCONAPP. However, by Lemma 28, we can consider just the RED_TYCONAPP case. This says that $\tau_i \rightsquigarrow_{\rho_i}^* \tau'_i$ and $\sigma_i \rightsquigarrow_{\rho_i}^* \tau'_i$, as desired, so we are done.

Case CO_LEFT: We adopt the variable names from the rule:

$$\frac{\Gamma \vdash \gamma : \tau_1 \tau_2 \sim_N \sigma_1 \sigma_2 \quad \Gamma \vdash \tau_1 : \kappa \quad \Gamma \vdash \sigma_1 : \kappa}{\Gamma \vdash \text{left } \gamma : \tau_1 \sim_N \sigma_1} \text{CO_LEFT}$$

We know by induction that $\tau_1 \tau_2 \Leftrightarrow_N \sigma_1 \sigma_2$. The steps to reach the common reduct must all be RED_APP, because newtype axioms are all at role R. Thus, the common reduct must be $\tau'_1 \tau'_2$ where $\tau_1 \rightsquigarrow_N^* \tau'_1$, and $\sigma_1 \rightsquigarrow_N^* \tau'_1$, so we are done.

Case CO_RIGHT: Similar to previous case.

Case CO_INST: We adopt the variable names from the rule:

$$\frac{\Gamma \vdash \gamma : \forall \alpha : \kappa. \tau_1 \sim_\rho \forall \alpha : \kappa. \sigma_1 \quad \Gamma \vdash \tau : \kappa}{\Gamma \vdash \gamma @ \tau : \tau_1[\tau/\alpha] \sim_\rho \sigma_1[\tau/\alpha]} \text{CO_INST}$$

We know by induction that $\forall \alpha : \kappa. \tau_1 \Leftrightarrow_\rho \forall \alpha : \kappa. \sigma_1$. We can easily see by inspection of the rewrite relation that the common reduct must have the form $\forall \alpha : \kappa. \tau_0$ for some τ_0 . We can also see by a straightforward induction that $\tau_1 \rightsquigarrow_\rho^* \tau_0$ and $\sigma_1 \rightsquigarrow_\rho^* \tau_0$. We must show that $\tau_1[\tau/\alpha] \rightsquigarrow_\rho^* \tau_0[\tau/\alpha]$ and $\sigma_1[\tau/\alpha] \rightsquigarrow_\rho^* \tau_0[\tau/\alpha]$. These facts come from an induction over the lengths of the derivations and the use of the simple rewrite substitution lemma, Lemma 25.

Case CO_SUB: We adopt the variable names in the rule:

$$\frac{\Gamma \vdash \gamma : \tau \sim_N \sigma}{\Gamma \vdash \text{sub } \gamma : \tau \sim_R \sigma} \text{CO_SUB}$$

We know that $\tau \Leftrightarrow_N \sigma$ and we need $\tau \Leftrightarrow_R \sigma$. This follows by induction over the lengths of the reduction and the use of Lemma 27. \square

Lemma 36 (Consistency). *The global context is consistent.*

Proof. Take a γ such that $\emptyset \vdash \gamma : \tau_1 \sim_R \tau_2$. By the completeness of the rewrite relation (Lemma 35), we see that $\tau_1 \Leftrightarrow_R \tau_2$. But, the rewrite relation consistency lemma (Lemma 34) tells us that τ_1 and τ_2 are consistent. Thus, the context admits only consistent coercions and is itself consistent. \square

H.2 Progress

Lemma 37 (Canonical forms).

1. If $\emptyset \vdash v : \tau_1 \rightarrow \tau_2$, then v is either $\lambda x : \tau. e'$ or $K\overline{\tau}\overline{\gamma}\overline{e}$.
2. If $\emptyset \vdash v : \forall \alpha : \kappa. \tau$, then v is either $\Lambda \alpha : \kappa. e'$ or $K\overline{\tau}$.

3. If $\emptyset \vdash v : \phi \Rightarrow \tau$, then v is either $\lambda c:\phi.e'$ or $K\bar{\tau}\bar{\gamma}$.
4. If $\emptyset \vdash v : D\bar{\sigma}$, then v is $K\bar{\tau}\bar{\gamma}\bar{e}$.

Lemma 38 (Value types). *If $\Gamma \vdash v : \tau$, then τ is a value type.*

Proof. If v is an abstraction, then the result is trivial. So, we assume that $v = K\bar{\tau}\bar{\gamma}\bar{e}$. Induction on the lengths of the lists of arguments yields

$$K : \forall \bar{\alpha}:\bar{\kappa}.\forall \bar{\beta}:\bar{\kappa}'.\bar{\phi} \Rightarrow \bar{\sigma} \rightarrow D\bar{\alpha}$$

We can see (again, by induction on the argument lists) that no matter what K is applied to, its type will always be a value type, headed by one of \forall , \Rightarrow , \rightarrow or D , all of which form value types. \square

Theorem 39 (Progress). *If $\emptyset \vdash e : \tau$, then either e is a value or a coerced value, or $e \rightarrow e'$ for some e' .*

Proof. We proceed by induction on the typing judgement $\emptyset \vdash e : \tau$.

Case TM_VAR: Cannot happen in an empty context.

Abstraction forms: Trivial.

Case TM_APP: We know $e = e_1 e_2$. By induction, we know that e_1 is either a value, a coerced value, or steps to e'_1 . If e_1 steps, we are done by S_APP_CONG. If e_1 is a value, the canonical forms lemma now gives us several cases:

Case $e_1 = \lambda x:\tau.e_3$: We are done by S_BETA.

Case $e_1 = K\bar{\tau}\bar{\gamma}\bar{e}$: Then, $e_1 e_2$ is a value.

If e_1 is a coerced value $v \triangleright \gamma$, then by the value types lemma (Lemma 38) and the consistency lemma (Lemma 36), the type of v must be headed by (\rightarrow) . We are done by S_PUSH.

Case TM_TAPP: Similar to previous case.

Case TM_CAPP: Similar to previous cases.

Case TM_DATACON: e is a value.

Case TM_CASE: We adopt the variable names from the rule:

$$\frac{\begin{array}{l} \Gamma \vdash e : D\bar{\sigma} \\ \Gamma \vdash \tau : * \\ \forall alt_i \text{ s.t. } alt_i \in \overline{alt} : \\ alt_i = K_i \bar{\alpha}_i \bar{c}_i \bar{x}_i \rightarrow e_i \\ K_i : \forall \bar{\alpha}'_i:\bar{\kappa}_i.\forall \bar{\beta}'_i:\bar{\kappa}'_i.\bar{\phi}_i \Rightarrow \bar{\tau}_i \rightarrow D\bar{\alpha}'_i \\ \Gamma, \bar{\alpha}_i:\bar{\kappa}'_i, (\bar{c}_i:\bar{\phi}_i, \bar{x}_i:\bar{\tau}_i)[\sigma/\alpha'_i][\alpha_i/\beta'_i] \vdash e_i : \tau \\ \overline{alt} \text{ is exhaustive} \end{array}}{\Gamma \vdash \text{case}_{\tau} e \text{ of } \overline{alt} : \tau} \quad \text{TM_CASE}$$

We know by induction that e is a value, a coerced value, or $e \rightarrow e'$ for some e' . If e steps, then we are done by S_CASE_CONG.

We know that T must actually be a data type (not a newtype), because it has a constructor. Thus, e has a value type. Therefore, if it has the form $v \triangleright \gamma$, the value v has a type headed by T as well. Thus $v = K\bar{\tau}\bar{\gamma}\bar{e}$ and we apply S_KPUSH, noting that the premises are all satisfied by straightforward use of typing judgements.

The final case is that e is a value. By the canonical forms lemma, we see that $e = K\bar{\tau}\bar{\gamma}\bar{e}$. Thus, S_IOTA applies, noting that the match must be exhaustive.

Case TM_CAST: We adopt the variable names from the rule:

$$\frac{\begin{array}{l} \Gamma \vdash e : \tau_1 \\ \Gamma \vdash \gamma : \tau_1 \sim_{\mathbf{R}} \tau_2 \end{array}}{\Gamma \vdash e \triangleright \gamma : \tau_2} \quad \text{TM_CAST}$$

By induction, we know that e is a value, a coerced value, or $e \rightarrow e'$.

If e steps, we are done by S_CAST_CONG.

If e is a value, then $e \triangleright \gamma$ is a coerced value, and we are done.

If e is a coerced value, then we are done by S_TRANS.

Case TM_CONTRA: We adopt the variable names from the rule:

$$\frac{\begin{array}{l} \emptyset \vdash \gamma : H_1 \sim_{\mathbf{N}} H_2 \quad H_1 \neq H_2 \\ \Gamma \vdash \tau : * \end{array}}{\Gamma \vdash \text{contra } \gamma \tau : \tau} \quad \text{TM_CONTRA}$$

By completeness of the rewrite relation (Lemma 35), we know that $H_1 \Leftrightarrow_{\mathbf{N}} H_2$. But, if $H \rightsquigarrow_{\mathbf{N}} H'$, then $H = H'$ (by induction on $H \rightsquigarrow_{\mathbf{N}} H'$, noting that all newtype axioms are at role R). So $H_1 = H_2$, contradicting a premise to this rule. Thus, this case cannot happen. \square

I. Role inference

Lemma 40 (Walking). *Let $\bar{\alpha}$ be the parameters to some type constant T . For some type σ , let $\bar{\beta}$ be the free variables in σ that are not in $\bar{\alpha}$. Let $\bar{\rho}$ be a list of roles of the same length as $\bar{\alpha}$. Let $\Omega = \bar{\alpha}:\bar{\rho}, \bar{\beta}:\mathbf{N}$.*

If $\text{walk}(T, \sigma)$ makes no change to the role of any of the $\bar{\alpha}$, then $\Omega \vdash \sigma : \mathbf{R}$.

Proof. By induction on the structure of σ :

Case $\sigma = \alpha'$: By assumption, it must be that $\alpha' : \mathbf{R} \in \Omega$ or $\alpha' : \mathbf{N} \in \Omega$. In either case, we can derive $\Omega \vdash \alpha' : \mathbf{R}$, so we are done.

Case $\sigma = \sigma_1 \sigma_2$: We check if σ can also be written as $H' \bar{\tau}$.

Case $\sigma = H' \bar{\tau}$: Let $\bar{\rho}' = \text{roles}(H')$. In order to conclude $\Omega \vdash H' \bar{\tau} : \mathbf{R}$, we will show that $\Omega \vdash \tau : \rho'$. Fix i ; we will show $\Omega \vdash \tau_i : \rho'_i$. Here, we have three cases:

Case $\rho'_i = \mathbf{N}$: By assumption, it must be that all the free variables in τ_i are assigned to N in Ω . Thus, by Lemma 4, we have $\Omega \vdash \tau_i : \mathbf{N}$ and we are done.

Case $\rho'_i = \mathbf{R}$: By assumption, it must be that $\text{walk}(T, \tau_i)$ makes no change. We then use the induction hypothesis to say that $\Omega \vdash \tau_i : \mathbf{R}$, and we are done.

Case $\rho'_i = \mathbf{P}$: We are done by RTY_PHANTOM.

Other applications: We wish to use RTY_APP. Thus, we must show that $\Omega \vdash \sigma_1 : \mathbf{R}$ and $\Omega \vdash \sigma_2 : \mathbf{N}$. For the former, we see that $\text{walk}(T, \sigma_1)$ must make no change, and we are done by induction. For the latter, we see that all the free variables in σ_2 must be assigned to N, and we are done by Lemma 4.

Case $\sigma = H$: We are done by immediate application of RTY_TY-CONAPP.

Case $\sigma = \forall \alpha':\kappa.\sigma_1$: We are done by induction, noting that in RTY_FORALL, α' gets assigned role N when checking σ_1 . This matches our expectations that the type variables $\bar{\beta}$ are at role N in the inductive hypothesis.

Case $\sigma = F(\bar{\tau})$: Repeated use of Lemma 4 tells us that $\Omega \vdash \tau : \mathbf{N}$. We are done by RTY_TYFAM. \square

Theorem 41. *The role inference algorithm always terminates.*

Proof. First, we observe that the walk procedure always terminates, as it is structurally recursive.

For the algorithm to loop in step 4, a role assigned to a variable must have changed. Yet, there are a finite number of such variables, and each variable may be updated only at most twice (from P to R and from R to N). Thus, at some point no more updates will happen and the algorithm will terminate. \square

Theorem 42 (Role inference is sound). *After running the role inference algorithm, $\text{roles}(H) \models H$ will hold for all H .*

Proof. We handle the data type case first. Fix a D . We will show that $\text{roles}(D) \models D$. Because the role inference algorithm has terminated, we know that $\text{walk}(D, \sigma)$ has caused no change for every σ that appears as a coercion type or term-level argument type in a constructor for D . Choose a constructor K , such that

$$K : \forall \bar{\alpha}:\bar{\kappa}.\forall \bar{\beta}:\bar{\kappa}'.\bar{\phi} \Rightarrow \bar{\sigma} \rightarrow D \bar{\alpha}$$

Let $\bar{\rho} = \text{roles}(D)$ and $\Omega = \bar{\alpha}:\bar{\rho}, \bar{\beta}:\bar{N}$. We have satisfied the premises of the walking lemma (Lemma 40), and thus we can conclude that $\Omega \vdash \sigma : R$. We have shown $\text{roles}(D) \models D$ by `ROLES_DATA`.

The newtype case is similar, using the right-hand side of the newtype definition in place of σ . \square

Lemma 43 (Stumbling). *Let $\bar{\alpha}$ be the parameters to some type constant T . For some type σ , let $\bar{\beta}$ be the free variables in σ that are not in $\bar{\alpha}$. Let $\bar{\rho}$ be a list of roles of the same length as $\bar{\alpha}$. Let $\Omega = \bar{\alpha}:\bar{\rho}, \bar{\beta}:\bar{N}$.*

If $\text{walk}(T, \sigma)$ were modified to skip one of its attempts to mark a variable, then it is not possible to conclude $\Omega \vdash \sigma : R$.

Proof. By induction on the structure of σ :

Case $\sigma = \alpha'$: If that mark were not done, then Ω would contain $\alpha':P$; this clearly violates $\Omega \vdash \alpha' : R$.

Case $\sigma = \sigma_1 \sigma_2$: We check if σ can also be written as $H' \bar{\tau}$.

Case $\sigma = H' \bar{\tau}$: Let $\bar{\rho}' = \text{roles}(H')$. Fix i .

Case $\rho'_i = N$: If we do not mark every free variable in τ_i as N, then it would be impossible to conclude $\Omega \vdash \tau_i : N$, by Lemma 4. Thus, we would not be able to conclude $\Omega \vdash H' \bar{\tau} : R$ by `RTY_TYCONAPP`. What about by `RTY_APP`? This, too, would require $\Omega \vdash \tau_i : N$, which we are unable to do.

Case $\rho'_i = R$: By induction, it is not possible to conclude $\Omega \vdash \tau_i : R$, and thus impossible to use `RTY_TYCONAPP`. What about `RTY_APP`? This would require $\Omega \vdash \tau_i : N$, which is not possible via the contrapositive of Lemma 5.

Case $\rho'_i = P$: There is no marking to be done here, so the assumption that walk is modified is false.

Other applications: Suppose the skipped marking were in the recursive call. Then, by induction, it is not possible to conclude $\Omega \vdash \sigma_1 : R$. Thus, it is not possible to conclude $\Omega \vdash \sigma_1 \sigma_2 : R$ by `RTY_APP`.

Now, suppose the skipped marking is when marking all free variables in σ_2 as N. In this case, we know that $\Omega \vdash \sigma_2 : N$ is impossible (by Lemma 4) and thus we cannot use `RTY_APP`.

Case $\sigma = H$: No mark was skipped, so the assumption that walk is modified is false.

Case $\sigma = \forall \alpha':\kappa.\sigma_1$: We are done by induction, noting that in `RTY_FORALL`, α' gets assigned role N when checking σ_1 . This matches our expectations that the type variables $\bar{\beta}$ are at role N in the inductive hypothesis.

Case $\sigma = F(\bar{\tau})$: If one of the variables free in the $\bar{\tau}$ were not marked as N, then it would be impossible to conclude $\Omega \vdash \tau_i : N$ for that τ_i (by Lemma 4). Thus, we would be unable to use `RTY_TYFAM`. \square

Theorem 44 (Role inference is optimal). *After running the role inference algorithm, any loosening of roles (a change from ρ to ρ' , where $\rho \leq \rho'$ and $\rho \neq \rho'$) would violate $\text{roles}(H) \models H$.*

Proof. Every time the role inference algorithm changes an assigned role from ρ' to ρ , it is the case that $\rho \leq \rho'$ and $\rho \neq \rho'$. Thus, all we must show is that every change the algorithm makes is necessary – that is, not making the change would then violate $\text{roles}(H) \models H$.

Role inference runs only on algebraic data types, so we need only concern ourselves with T s, not general H s. In both the data type and newtype cases, showing $\text{roles}(T) \models T$ requires showing $\Omega \vdash \sigma : R$, where $\Omega = \bar{\alpha}:\bar{\rho}, \bar{\beta}:\bar{N}$ and $\bar{\alpha}$ are the parameters to T and $\bar{\beta}$ are the remaining free variables of σ . (In the newtype case, $\bar{\beta}$ is empty.) The list of roles $\bar{\rho}$ is $\text{roles}(T)$. So, we must show that skipping any change in the $\text{walk}(T, \sigma)$ algorithm means that $\Omega \vdash \sigma : R$ would not be derivable. This is precisely what Lemma 43 shows and so we are done. \square