

Spores, Formally

Heather Miller and Philipp Haller

December 2013

1 Overview

Spores are designed to avoid problems of closures. This is done using two mechanisms: the spore shape and context bounds for the spore's environment.

A spore is a closure with a specific shape that dictates how the environment of a spore is declared. In general, a spore has the following shape:

```
spore {  
  val y1: S1 = <expr1>  
  ...  
  val yn: Sn = <exprn>  
  (x: T) => {  
    // ...  
  }  
}
```

A spore consists of two parts: the header and the body. The list of value definitions at the beginning is called the spore header. The header is followed by a regular closure, the spore's body. The characteristic property of a spore is that the body of its closure is only allowed to access its parameter, values in the spore header, as well as top-level singleton objects (public, global state). In particular, the spore closure is not allowed to capture variables in the environment. Only an expression on the right-hand side of a value definition in the spore header is allowed to capture variables.

By enforcing this shape, the environment of a spore is always declared explicitly in the spore header which avoids accidentally capturing problematic references. Moreover, and that's important for OO languages, it's no longer possible to accidentally capture the "this" reference.

Note that the evaluation semantics of a spore is equivalent to a closure obtained by leaving out the "spore" marker:

```
{  
  val y1: S1 = <expr1>  
  ...  
  val yn: Sn = <exprn>  
  (x: T) => {  
    // ...  
  }  
}
```

In Scala, the above block first initializes all value definitions in order and then evaluates to a closure that captures the introduced local variables `y1`, ..., `yn`. The corresponding spore has the exact same evaluation semantics. What's interesting is that this closure shape is already used in production systems such as Spark to avoid problems with accidentally captured "this" references. However, in these systems the above shape is not enforced, whereas with spores it is.

The result type of the "spore" constructor is not a regular function type, but a subtype of one of Scala's function types. This is possible, because in Scala functions are instances of classes that mix in one of the function traits. For example, the trait for functions of arity one looks like this:¹

```
trait Function1[-A, +B] {  
  def apply(x: A): B  
}
```

The `apply` method is abstract; a concrete implementation applies the body of the function that's being defined to the argument `x`. Functions are contravariant in their argument type `A`, indicated using the "-" symbol, and covariant in their result type `B`, indicated using the "+" symbol.

The type of a spore of arity one is a subtype of `Function1`:

```
trait Spore[-A, +B] extends Function1[A, B]
```

Using the `Spore` trait methods can require argument closures to be spores:

```
def sendOverWire(s: Spore[Int, Int]): Unit = ...
```

This way, libraries and frameworks can enforce the use of spores instead of plain closures, thereby reducing the risk for common programming errors.

1.1 Context bounds

The fact that for spores a certain shape is enforced is very useful. However, in some situations this is not enough. For example, using closures in a concurrent setting is very error-prone, because of the fact that it's possible to capture mutable objects which leads to race conditions. Thus, closures should only capture immutable objects to avoid interference. However, such constraints cannot be enforced using the spore shape alone (captured objects are stored in constant values in the spore header, but such a constant might still refer to a mutable object).

In this section we introduce a form of type-based constraints called "context bounds" that can be attached to a spore which enforce certain type-based properties for all captured variables of a spore.

Taking another example, it might be necessary for a spore to require the availability of instances of a certain type class for the types of all its captured variables. A typical example for such a type class is `Pickler`: types with an instance of the `Pickler` type class can be pickled using a new pickling framework

¹For simplicity we omit definitions of the 'andThen' and 'compose' methods in the definition of 'Function1'.

for Scala. To be able to pickle a spore, it's necessary that all its captured types have an instance of `Pickler`.²

Spores allow expressing such a requirement using implicit properties. The idea is that if there is an implicit of type `Property[Pickler]` in scope at the point where a spore is created, then it is enforced that all captured types in the spore header have an instance of the `Pickler` type class:

```
import spores.withPickler

spore {
  val name: String = <expr1>
  val age: Int = <expr2>
  (x: String) => {
    // ...
  }
}
```

While an imported property does not have an impact on how a spore is constructed (besides the property import), it has an impact on the result type of the spore macro. In the above example, the result type would be a refinement of the `Spore` type:

```
Spore[String, Int] {
  type Captured = (String, Int)
  val captured: Captured
  implicit val p$1 = implicitly[Pickler[(String, Int)]]
  (x: String) => {
    // ...
  }
}
```

The refinement type contains a type member `Captured` which is defined to be a tuple of all the captured types. The values of the actual captured variables are accessible using the captured value member. What's more, the refinement type contains for each type class that's required an implicit value with a type class instance for type `Captured`.

Such implicit values allow retrieving a type class instance for the captured types of a given spore using Scala's `implicitly` function as follows:

```
val s = spore { ... }

implicitly[Pickler[s.Captured]]
```

Note that `s.Captured` is defined to be the type of the environment of spore `s`: a tuple with all types of captured variables.

²A spore can be pickled by pickling its environment and the fully-qualified class name of its corresponding function class.

2 Formalization

$t ::= x$	variable
$ (x : T) \Rightarrow t$	abstraction
$ t t$	application
$ \text{let } x = t \text{ in } t$	let binding
$ \{\overline{l = t}\}$	record construction
$ t.l$	selection
$ \text{spore } \{ \overline{x : T = t} ; \overline{pn}; (x : T) \Rightarrow t \}$	spore
$ \text{import } pn \text{ in } t$	property import
$ t \text{ compose } t$	spore composition
$v ::= (x : T) \Rightarrow t$	abstraction
$ \{\overline{l = v}\}$	record value
$ \text{spore } \{ \overline{x : T = v} ; \overline{pn}; (x : T) \Rightarrow t \}$	spore value
$T ::= T \Rightarrow T$	function type
$ \{\overline{l : T}\}$	record type
$ \mathcal{S}$	
$\mathcal{S} ::= T \Rightarrow T \{ \text{type } \mathcal{C} = \overline{T} ; \overline{pn} \}$	spore type
$ T \Rightarrow T \{ \text{type } \mathcal{C} ; \overline{pn} \}$	abstract spore type
$P \in pn \rightarrow \mathcal{T}$	property map
$\mathcal{T} \in \mathcal{P}(T)$	type family
$\Gamma ::= \overline{x : T}$	type environment
$\Delta ::= \overline{pn}$	property environment

Figure 1: Core language syntax

We formalize spores in the context of a standard, typed lambda calculus with records. Apart from novel language and type-systematic features, our formal development follows a well-known methodology ?. Figure 1 shows the syntax of our core language. Terms are standard except for the **spore**, **import**, and **compose** terms. A **spore** term creates a new spore. It contains a list of variable definitions (the spore header), a list of property names, and the spore’s closure. A property name refers to a type family (a set of types) that all captured types must belong to.

An illustrative example of a property name and its associated type family, but in the context of Scala, is a type class: a spore satisfies such a property if there is a type class instance for all its captured types.

An **import** term imports a property name into the property environment within a lexical scope (a term); the property environment contains properties that are registered as requirements whenever a spore is created. This is explained in more detail in Section 2.2. A **compose** term is used to compose two spores. The core language provides spore composition as a built-in feature, because type checking spore composition is markedly different from type checking regular function composition (see Section 2.2).

The grammar of values is standard except for spore values; in a spore value each term on the right-hand side of a definition in the spore header is a value.

The grammar of types is standard except for spore types. Spore types are refinements of function types. They additionally contain a (possibly-empty) sequence of captured types, which can be left abstract, and a sequence of property names.

2.1 Subtyping

Figure 2 shows the subtyping rules. Record (S-REC) and function (S-FUN) subtyping are standard.

The subtyping rule for spores (S-SPORE) is analogous to the subtyping rule for functions with respect to the argument and result types. Additionally, for two spore types to be in a subtyping relationship either their captured types have to be the same ($M_1 = M_2$) or the supertype must be an abstract spore type ($M_2 = \text{type } \mathcal{C}$). The subtype must guarantee at least the properties of its supertype, or a superset thereof. Taken together, this rule expresses the fact that a spore type whose type member \mathcal{C} is not abstract is compatible with an abstract spore type as long as it has a superset of the supertype’s properties. This is important for spores used as first-class values: functions operating on spores with arbitrary environments can simply demand an abstract spore type. The way both the captured types and the properties are modeled corresponds to (but simplifies) the subtyping rule for refinement types in Scala (see Section ??).

Rule S-SPOREFUN expresses the fact that spore types are refinements of their corresponding function types, giving rise to a subtyping relationship.

$$\begin{array}{c}
 \text{S-REC} \\
 \frac{\bar{l} \subseteq \bar{l} \quad l_i = l'_i \rightarrow T_i <: T'_i \wedge T'_i <: T_i}{\{\bar{l} : \bar{T}\} <: \{\bar{l}' : \bar{T}'\}} \\
 \\
 \text{S-FUN} \\
 \frac{T_2 <: T_1 \quad R_1 <: R_2}{T_1 \Rightarrow R_1 <: T_2 \Rightarrow R_2} \\
 \\
 \text{S-SPORE} \\
 \frac{T_2 <: T_1 \quad R_1 <: R_2 \quad \overline{pn'} \subseteq \overline{pn} \quad M_1 = M_2 \vee M_2 = \text{type } \mathcal{C}}{T_1 \Rightarrow R_1 \{ M_1 ; \overline{pn} \} <: T_2 \Rightarrow R_2 \{ M_2 ; \overline{pn'} \}} \\
 \\
 \text{S-SPOREFUN} \\
 T_1 \Rightarrow R_1 \{ M ; \overline{pn} \} <: T_1 \Rightarrow R_1
 \end{array}$$

Figure 2: Subtyping

2.2 Typing rules

Typing derivations use a judgement of the form $\Gamma; \Delta \vdash t : T$. Besides the standard variable environment Γ we use a property environment Δ which is a sequence of property names that are “active” while deriving the type T of term t . The property environment is reminiscent of the implicit parameter context used in the original work on implicit parameters ?; it is an environment for names whose definition sites “just happen to be far removed from their usages.”

In the typing rules we assume the existence of a global property mapping P from property names pn to type families \mathcal{T} . This technique is reminiscent of the way some object-oriented core languages provide a global class table for

$$\begin{array}{c}
\text{T-VAR} \quad \frac{x : T \in \Gamma}{\Gamma; \Delta \vdash x : T} \quad \text{T-SUB} \quad \frac{\Gamma; \Delta \vdash t : T' \quad T' <: T}{\Gamma; \Delta \vdash t : T} \quad \text{T-ABS} \quad \frac{\Gamma, x : T_1; \Delta \vdash t : T_2}{\Gamma; \Delta \vdash (x : T_1) \Rightarrow t : T_1 \Rightarrow T_2} \\
\\
\text{T-APP} \quad \frac{\Gamma; \Delta \vdash t_1 : T_1 \Rightarrow T_2 \quad \Gamma; \Delta \vdash t_2 : T_1}{\Gamma; \Delta \vdash (t_1 t_2) : T_2} \quad \text{T-LET} \quad \frac{\Gamma; \Delta \vdash t_1 : T_1 \quad \Gamma, x : T_1; \Delta \vdash t_2 : T_2}{\Gamma; \Delta \vdash \text{let } x = t_1 \text{ in } t_2 : T_2} \\
\\
\text{T-REC} \quad \frac{\Gamma; \Delta \vdash \overline{t} : \overline{T}}{\Gamma; \Delta \vdash \{\overline{l} = \overline{t}\} : \{\overline{l} : \overline{T}\}} \quad \text{T-SEL} \quad \frac{\Gamma; \Delta \vdash t : \{\overline{l} : \overline{T}\}}{\Gamma; \Delta \vdash t.l_i : T_i} \quad \text{T-IMP} \quad \frac{\Gamma; \Delta, pn \vdash t : T}{\Gamma; \Delta \vdash \text{import } pn \text{ in } t : T} \\
\\
\text{T-SPORE} \quad \frac{\forall s_i \in \overline{s}. \Gamma; \Delta \vdash s_i : S_i \quad \overline{y} : \overline{S}, x : T_1; \Delta \vdash t_2 : T_2 \quad \forall pn \in \Delta, \Delta'. \overline{S} \subseteq P(pn)}{\Gamma; \Delta \vdash \text{spore } \{ \overline{y} : \overline{S} = s ; \Delta' ; (x : T_1) \Rightarrow t_2 \} : T_1 \Rightarrow T_2 \{ \text{type } \mathcal{C} = \overline{S} ; \Delta, \Delta' \}} \\
\\
\text{T-COMP} \quad \frac{\Gamma; \Delta \vdash t_1 : T_1 \Rightarrow T_2 \{ \text{type } \mathcal{C} = \overline{S} ; \Delta_1 \} \quad \Gamma; \Delta \vdash t_2 : U_1 \Rightarrow T_1 \{ \text{type } \mathcal{C} = \overline{R} ; \Delta_2 \} \quad \Delta' = \{ pn \in \Delta_1 \cup \Delta_2 \mid \overline{S} \subseteq P(pn) \wedge \overline{R} \subseteq P(pn) \}}{\Gamma; \Delta \vdash t_1 \text{ compose } t_2 : U_1 \Rightarrow T_2 \{ \text{type } \mathcal{C} = \overline{S}, \overline{R} ; \Delta' \}}
\end{array}$$

Figure 3: Typing rules

type-checking. The main difference is that our core language does not include constructs to extend the global property map; such constructs are left out of the core language for simplicity, since the creation of properties is not essential to our model.

The typing rules are standard except for rules T-IMP, T-SPORE, and T-COMP, which are new. Only these three type rules inspect or modify the property environment Δ . Note that there is no rule for spore application, since there is a subtyping relationship between spores and functions (see Section 2.1). Using the subsumption rule T-SUB spore application is expressed using the standard rule for function application (T-APP).

Rule T-IMP imports a property pn into the property environment within the scope defined by term t .

Rule T-SPORE derives a type for a spore term. In the spore, all terms on right-hand sides of variable definitions in the spore header must be well-typed in the same environment $\Gamma; \Delta$ according to their declared type. The body of the spore's closure, t_2 , must be well-typed in an environment containing only the variables in the spore header and the closure's parameter, one of the central properties of spores. The last premise requires all captured types to satisfy both the properties in the current property environment, Δ , as well as the properties listed in the spore term, Δ' . Finally, the resulting spore type contains the argument and result types of the spore's closure, the sequence of captured types according to the spore header, and the concatenation of properties Δ and Δ' . The intuition here is that properties in the environment have been

$$\begin{array}{c}
\text{E-LET1} \\
\frac{t_1 \rightarrow t'_1}{\text{let } x = t_1 \text{ in } t_2 \rightarrow \text{let } x = t'_1 \text{ in } t_2} \\
\\
\text{E-LET2} \\
\text{let } x = v_1 \text{ in } t_2 \rightarrow [x \mapsto v_1]t_2 \\
\\
\text{E-REC} \\
\frac{t_k \rightarrow t'_k}{\overline{\{\bar{l} = v, l_k = t_k, \bar{l}' = t'\} \rightarrow \{\bar{l} = v, l_k = t'_k, \bar{l}' = t'\}}} \\
\\
\text{E-SEL1} \\
\frac{t \rightarrow t'}{t.l \rightarrow t'.l} \\
\\
\text{E-SEL2} \\
\{\bar{l} = v\}.l_i \rightarrow v_i \\
\\
\text{E-APP1} \\
\frac{t_1 \rightarrow t'_1}{t_1 t_2 \rightarrow t'_1 t_2} \\
\\
\text{E-APP2} \\
\frac{t_2 \rightarrow t'_2}{v_1 t_2 \rightarrow v_1 t'_2} \\
\\
\text{E-APPABS} \\
((x : T) \Rightarrow t)v \rightarrow [x \mapsto v]t \\
\\
\text{E-APPSPORE} \\
\frac{\forall pn \in \overline{pn}. \bar{T} \subseteq P(pn)}{\text{spore } \{ x : T = v; \overline{pn}; (x' : T) \Rightarrow t \} v' \rightarrow [x \mapsto v][x' \mapsto v']t} \\
\\
\text{E-SPORE} \\
\frac{t_k \rightarrow t'_k}{\text{spore } \{ x : T = v, x_k : T_k = t_k, \overline{x' : T' = t'}; (x : T) \Rightarrow t \} \rightarrow \text{spore } \{ x : T = v, x_k : T_k = t'_k, \overline{x' : T' = t'}; (x : T) \Rightarrow t \}} \\
\\
\text{E-IMP} \\
\text{import } pn \text{ in } t \rightarrow \text{insert}(pn, t) \\
\\
\text{E-COMP1} \\
\frac{t_1 \rightarrow t'_1}{t_1 \text{ compose } t_2 \rightarrow t'_1 \text{ compose } t_2} \\
\\
\text{E-COMP2} \\
\frac{t_2 \rightarrow t'_2}{v_1 \text{ compose } t_2 \rightarrow v_1 \text{ compose } t'_2} \\
\\
\text{E-COMP3} \\
\frac{\Delta = \{p \mid p \in \overline{pn}, \overline{qn}. \bar{T} \subseteq P(p) \wedge \bar{S} \subseteq P(p)\}}{\text{spore } \{ x : T = v; \overline{pn}; (x' : T') \Rightarrow t \} \text{ compose spore } \{ y : S = w; \overline{qn}; (y' : S') \Rightarrow t' \} \rightarrow \text{spore } \{ x : T = v, y : S = w; \Delta; (y' : S') \Rightarrow \text{let } z' = t' \text{ in } [x' \mapsto z']t \}}
\end{array}$$

Figure 4: Operational Semantics³

explicitly imported by the user, thus indicating that all spores in the scope of the corresponding import should satisfy them.

Rule T-COMP derives a result type for the composition of two spores. It inspects the captured types of both spores (\bar{S} and \bar{R}) to ensure that the properties of the resulting spore, Δ , are satisfied by the captured variables of both spores. Otherwise, the argument and result types are analogous to regular function composition. Note that it's always possible to weaken the properties of a spore through spore subtyping and subsumption (T-SUB).

2.3 Operational semantics

Figure 4 shows the evaluation rules of a small-step operational semantics for our core language. The only non-standard rules are E-APPSPORE, E-SPORE, E-IMP, and E-COMP3. Rule E-APPSPORE applies a spore literal to an argument value. The differences to regular function application (E-APPABS) are (a) that the types in the spore header must satisfy the properties of the spore dynamically, and (b) that the variables in the spore header must be replaced by their values in the body of the spore’s closure. Rule E-SPORE is a simple congruence rule. Rule E-IMP is a computation rule that is always enabled. It adds property name pn to all spore terms within the body t . The *insert* helper function is defined in Figure 5 (we omit rules for `compose` and `let`, since they are analogous to rules H-INSAPP and H-INSSEL).

Rule E-COMP3 is the computation rule for spore composition. Besides computing the composition in a way analogous to regular function composition, it defines the spore header of the result spore, as well as its properties. The properties of the result spore are restricted to those that are satisfied by the captured variables of both argument spores.

$$\begin{array}{c}
 \text{H-INSPORE1} \\
 \frac{\forall t_i \in \bar{l}. \text{insert}(pn, t_i) = t'_i \quad \text{insert}(pn, t) = t'}{\text{insert}(pn, \text{spore } \{ \overline{x : T = t; \bar{pn}}; (x' : T) \Rightarrow t \}) = \text{spore } \{ \overline{x : T = t'; \bar{pn}, pn}; (x' : T) \Rightarrow t' \}} \\
 \\
 \text{H-INSPORE2} \\
 \text{insert}(pn, \text{spore } \{ \overline{x : T = v; \bar{pn}}; (x' : T) \Rightarrow t \}) = \text{spore } \{ \overline{x : T = v; \bar{pn}, pn}; (x' : T) \Rightarrow t \} \\
 \\
 \text{H-INSAPP} \\
 \text{insert}(pn, t_1 t_2) = \text{insert}(pn, t_1) \text{insert}(pn, t_2) \\
 \\
 \text{H-INSSEL} \\
 \text{insert}(pn, t.l) = \text{insert}(pn, t).l
 \end{array}$$

Figure 5: Helper function *insert*

2.4 Soundness

This section presents a soundness proof of the spore type system. The proof is based on a pair of progress and preservation theorems[?]. A complete proof of soundness appears in the companion technical report[?]. In addition to standard lemmas, such as Lemma 2.4 and Lemma 2.5, we also prove a lemma specific to our type system, namely Lemma 2.3, which ensures types are preserved under property import. Soundness of the type system follows from Theorem 2.2 and Theorem 2.6.

Lemma 2.1. (Canonical forms)

³For the sake of brevity, here we omit the standard evaluation rules. The complete set of evaluation rules can be found in the accompanying technical report[?]

1. If v is a value of type $\{\overline{l : T}\}$, then v is $\{\overline{l = v}\}$ where \overline{v} is a sequence of values.
2. If v is a value of type $T \Rightarrow R$, then v is either $(x : T_1) \Rightarrow t$ or $\mathbf{spore} \{ \overline{y : S = v} ; \overline{pn}; (x : T_1) \Rightarrow t \}$ where $T <: T_1$ and \overline{v} is a sequence of values.
3. If v is a value of type $T \Rightarrow R \{ \mathbf{type} \mathcal{C} = \overline{S} ; \overline{pn} \}$, then v is $\mathbf{spore} \{ \overline{y : S = v} ; \overline{pn}; (x : T_1) \Rightarrow t \}$ where $T <: T_1$ and \overline{v} is a sequence of values.

Proof. According to the grammar in Figure 1, values in the core language can have three forms: $(x : T) \Rightarrow t$, $\{\overline{l = v}\}$, and $\mathbf{spore} \{ \overline{x : T = v} ; \overline{pn}; (x : T) \Rightarrow t \}$ where \overline{v} is a sequence of values.

For the first part, according to (T-REC) and the subtyping rules, v is $\{\overline{l = v}\}$ where \overline{v} is a sequence of values of types \overline{T} .

For the second part, according to the subtyping rules v can have either type $T_1 \Rightarrow R_1$, $T_1 \Rightarrow R_1 \{ \mathbf{type} \mathcal{C} = \overline{S} ; \overline{pn} \}$, or $T_1 \Rightarrow R_1 \{ \mathbf{type} \mathcal{C} ; \overline{pn} \}$ where $T <: T_1$ and $R_1 <: R$. If v has type $T_1 \Rightarrow R_1$, then according to the grammar and (T-ABS) v must be $(x : T) \Rightarrow t$. If v has either type $T_1 \Rightarrow R_1 \{ \mathbf{type} \mathcal{C} = \overline{S} ; \overline{pn} \}$ or type $T_1 \Rightarrow R_1 \{ \mathbf{type} \mathcal{C} ; \overline{pn} \}$, then according to the grammar and (T-SPORE) v must be $\mathbf{spore} \{ \overline{x : T = v} ; \overline{pn}; (x : T_1) \Rightarrow t \}$ where \overline{v} is a sequence of values.

Part three is similar. \square

Theorem 2.2. (Progress) *Suppose t is a closed, well-typed term (that is, $\vdash t : T$ for some T). Then either t is a value or else there is some t' with $t \rightarrow t'$.*

Proof. By induction on a derivation of $t : T$. The only three interesting cases are the ones for spore creation, application (where we might apply a spore to some argument), and spore composition.

Case T-SPORE: $t = \mathbf{spore} \{ \overline{x : S = t} ; \Delta'; (x : T_1) \Rightarrow t_2 \}$, $\forall t_i \in \overline{t}. \vdash t_i : S_i$, and $\overline{x : S}, x : T_1 \vdash t_2 : T_2$. By the induction hypothesis, either all \overline{t} are values, in which case t is a value; or there is a term t_i such that $t_i \rightarrow t'_i$ (since $\vdash t_i : S_i$). Thus, by (E-SPORE), $t \rightarrow t'$ for some term t' .

Case T-APP: $t = t_1 t_2$ and $\vdash t_1 : T_1 \Rightarrow T_2$ and $\vdash t_2 : T_1$. By the induction hypothesis, either t_1 is a value v_1 , or $t_1 \rightarrow t'_1$. In the latter case it follows from (E-APP1) that $t \rightarrow t'$ for some t' . In the former case, by the induction hypothesis t_2 is either a value v_2 or $t_2 \rightarrow t'_2$. In the former case by the canonical forms lemma we have that v_2 is either $(x : T_1) \Rightarrow t$ or $\mathbf{spore} \{ \overline{x : T = v} ; \overline{pn}; (x : T_1) \Rightarrow t \}$ where $T <: T_1$ and \overline{v} is a sequence of values; thus, either (E-APPABS) or (E-APPSPORE) apply. In the latter case, the result follows from (E-APP2).

Case T-COMP: $t = t_1 \mathbf{compose} t_2$ and $\vdash t_1 : T_1 \Rightarrow T_2 \{ \mathbf{type} \mathcal{C} = \overline{S} ; \Delta_1 \}$ and $\vdash t_2 : U_1 \Rightarrow T_1 \{ \mathbf{type} \mathcal{C} = \overline{R} ; \Delta_2 \}$. If either t_1 or t_2 is not a value, the result follows from the induction hypothesis and (E-COMP1) or (E-COMP2). If t_1 is a value v_1 and t_2 is a value v_2 , then by the canonical forms lemma, $v_1 = \mathbf{spore} \{ \overline{y : S = v} ; \Delta_1; (x : T_1) \Rightarrow s_1 \}$ and $v_2 = \mathbf{spore} \{ \overline{z : R = w} ; \Delta_2; (u : U_1) \Rightarrow s_2 \}$. Thus, by (E-COMP3), $t \rightarrow t'$ for some t' . \square

Lemma 2.3. (Preservation of types under import) *If $\Gamma; \Delta, pn \vdash t : T$ then $\Gamma; \Delta \vdash \mathit{insert}(pn, t) : T$*

Proof. By induction on a derivation of $t : T$. The only three interesting cases are the ones for spore creation, application (where we might apply a spore to some argument), and spore composition.

Case T-SPORE: $t = \mathbf{spore} \{ \overline{x : \overline{S = t}} ; \Delta' ; (x : T_1) \Rightarrow t_2 \}$, $\forall t_i \in \overline{t}. \vdash t_i : S_i$, and $\overline{x : \overline{S}}, x : T_1 \vdash t_2 : T_2$. By the induction hypothesis, either all \overline{t} are values, in which case t is a value; or there is a term t_i such that $t_i \rightarrow t'_i$ (since $\vdash t_i : S_i$). Thus, by (E-SPORE), $t \rightarrow t'$ for some term t' .

Case T-APP: $t = t_1 t_2$ and $\vdash t_1 : T_1 \Rightarrow T_2$ and $\vdash t_2 : T_1$. By the induction hypothesis, either t_1 is a value v_1 , or $t_1 \rightarrow t'_1$. In the latter case it follows from (E-APP1) that $t \rightarrow t'$ for some t' . In the former case, by the induction hypothesis t_2 is either a value v_2 or $t_2 \rightarrow t'_2$. In the former case by the canonical forms lemma we have that v_2 is either $(x : T_1) \Rightarrow t$ or $\mathbf{spore} \{ \overline{x : \overline{T = v}} ; \overline{pn}; (x : T_1) \Rightarrow t \}$ where $T <: T_1$ and \overline{v} is a sequence of values; thus, either (E-APPABS) or (E-APPSPORE) apply. In the latter case, the result follows from (E-APP2).

Case T-COMP: $t = t_1 \mathbf{compose} t_2$ and $\vdash t_1 : T_1 \Rightarrow T_2 \{ \mathbf{type} \mathcal{C} = \overline{S} ; \Delta_1 \}$ and $\vdash t_2 : U_1 \Rightarrow T_1 \{ \mathbf{type} \mathcal{C} = \overline{R} ; \Delta_2 \}$. If either t_1 or t_2 is not a value, the result follows from the induction hypothesis and (E-COMP1) or (E-COMP2). If t_1 is a value v_1 and t_2 is a value v_2 , then by the canonical forms lemma, $v_1 = \mathbf{spore} \{ \overline{y : \overline{S = v}} ; \Delta_1 ; (x : T_1) \Rightarrow s_1 \}$ and $v_2 = \mathbf{spore} \{ \overline{z : \overline{R = w}} ; \Delta_2 ; (u : U_1) \Rightarrow s_2 \}$. Thus, by (E-COMP3), $t \rightarrow t'$ for some t' . □

Lemma 2.4. (Preservation of types under substitution) *If $\Gamma, x : S ; \Delta \vdash t : T$ and $\Gamma ; \Delta \vdash s : S$, then $\Gamma ; \Delta \vdash [x \mapsto s]t : T$*

Proof. By induction on a derivation of $\Gamma, x : S ; \Delta \vdash t : T$. □

Lemma 2.5. (Weakening) *If $\Gamma ; \Delta \vdash t : T$ and $x \notin \text{dom}(\Gamma)$, then $\Gamma, x : S ; \Delta \vdash t : T$.*

Proof. By induction on a derivation of $\Gamma ; \Delta \vdash t : T$. □

Theorem 2.6. (Preservation) *If $\Gamma ; \Delta \vdash t : T$ and $t \rightarrow t'$, then $\Gamma ; \Delta \vdash t' : T$.*

Proof. By induction on a derivation of $t : T$.

- Case T-SEL: $t = s.l_i$ and $\Gamma ; \Delta \vdash s : \{ \overline{l : \overline{S}} \}$. Since $t \rightarrow t'$ we have either by (E-SEL1) $s \rightarrow s'$ and $t' = s'.l_i$, or we have by (E-SEL2) $s = \{ \overline{l = v} \}$ and $t' = v_i$. In the former case, by the induction hypothesis, $\Gamma ; \Delta \vdash s' : \{ \overline{l : \overline{S}} \}$ and thus by (T-SEL), $\Gamma ; \Delta \vdash s'.l_i : S_i$. In the latter case, by (T-REC), $\Gamma ; \Delta \vdash v_i : S_i$.
- Case T-IMP: $t = \mathbf{import} \textit{pn} \textit{ in } s$ and $\Gamma ; \Delta, \textit{pn} \vdash s : T$. Since $t \rightarrow t'$, we have by (E-IMP) $t' = \mathbf{insert}(\textit{pn}, s)$. By Lemma 2.3, $\Gamma ; \Delta \vdash \mathbf{insert}(\textit{pn}, s) : T$.
- Case T-APP: $t = s_1 s_2$ and $T = S_2$. By (T-APP), $\Gamma ; \Delta \vdash s_1 : S_1 \Rightarrow S_2$ and $\Gamma ; \Delta \vdash s_2 : S_1$. Since $t \rightarrow t'$, either (E-APP1), (E-APP2), (E-APPABS), or (E-APPSPORE) applies. If (E-APP1) applies, then $s_1 \rightarrow s'_1$ and $t' = s'_1 s_2$. By the induction hypothesis, $\Gamma ; \Delta \vdash s'_1 : S_1 \Rightarrow S_2$. By (T-APP), $\Gamma ; \Delta \vdash t' : S_2$. The case where (E-APP2) applies is similar. If (E-APPABS) applies, then $s_1 = (x : S_1) \Rightarrow t_2$ and $s_2 = v$ and $t' = [x \mapsto v]t_2$. By (T-ABS),

$\Gamma, x : S_1; \Delta \vdash t_2 : S_2$. By (T-APP), $\Gamma; \Delta \vdash v : S_1$. By Lemma 2.4, $\Gamma; \Delta \vdash [x \mapsto v]t_2 : S_2$.

If (E-APPSPORE) applies, then $s_1 = \mathbf{spore} \{ \overline{x : T = v}; \overline{pn}; (y : S_1) \Rightarrow t_2 \}$ and $s_2 = v'$ and $\forall pn \in \overline{pn}. \overline{S} \subseteq P(pn)$ and $t' = \overline{[x \mapsto v][y \mapsto v']}t_2$. By (T-SPORE), $\overline{x : T}, y : S_1; \Delta \vdash t_2 : S_2$. By (T-APP), $\Gamma; \Delta \vdash v' : S_1$. By Lemma 2.5, $\Gamma, x : T, y : S_1; \Delta \vdash t_2 : S_2$. By Lemma 2.5, $\Gamma, x : T; \Delta \vdash v' : S_1$. By Lemma 2.4, $\Gamma, x : T; \Delta \vdash [y \mapsto v']t_2 : S_2$. By (T-SPORE), we also have $\forall v_i \in \overline{v}. \Gamma; \Delta \vdash v_i : T_i$. By Lemma 2.4, $\Gamma; \Delta \vdash [x \mapsto v][y \mapsto v']t_2 : S_2$.

- Case T-SPORE: $t = \mathbf{spore} \{ \overline{y : S = s}; \Delta'; (x : T_1) \Rightarrow t_2 \}$ and $T = T_1 \Rightarrow \overline{T_2} \{ \mathbf{type} \mathcal{C} = \overline{S}; \Delta, \Delta' \}$. By (T-SPORE), $\forall s_i \in \overline{s}. \Gamma; \Delta \vdash s_i : S_i$ and $\overline{y : S}, x : T_1; \Delta \vdash t_2 : T_2$ and $\forall pn \in \Delta, \Delta'. \overline{S} \subseteq P(pn)$. Since $t \rightarrow t'$, rule (E-SPORE) must apply, and thus $s_i \rightarrow s'_i$ for some s_i . By the induction hypothesis, $\Gamma; \Delta \vdash s'_i : S_i$. Thus, by (T-SPORE), $\Gamma; \Delta \vdash t' : T$.
- Case T-COMP: $t = s_1 \mathbf{compose} s_2$ and $T = T_1 \Rightarrow T_2 \{ \mathbf{type} \mathcal{C} = \overline{S}, \overline{R}; \Delta_3 \}$. By (T-COMP), $\Gamma \vdash s_1 : U_1 \Rightarrow T_2 \{ \mathbf{type} \mathcal{C} = \overline{S}; \Delta_1 \}$ and $\Gamma \vdash s_2 : T_1 \Rightarrow U_1 \{ \mathbf{type} \mathcal{C} = \overline{R}; \Delta_2 \}$ and $\Delta_3 = \{pn \in \Delta_1 \cup \Delta_2 \mid \overline{S} \subseteq P(pn) \wedge \overline{R} \subseteq P(pn)\}$. Since $t \rightarrow t'$, either (E-COMP1), (E-COMP2), or (E-COMP3) applies.

If (E-COMP1) applies, then $s_1 \rightarrow s'_1$, and by (T-COMP), $\Gamma; \Delta \vdash s_1 : U_1 \Rightarrow T_2 \{ \mathbf{type} \mathcal{C} = \overline{S}; \Delta_1 \}$, and $t' = s'_1 \mathbf{compose} s_2$. By the induction hypothesis, $\Gamma; \Delta \vdash s'_1 : U_1 \Rightarrow T_2 \{ \mathbf{type} \mathcal{C} = \overline{S}; \Delta_1 \}$. By (T-COMP), we know that $\Gamma; \Delta \vdash s_2 : T_1 \Rightarrow U_1 \{ \mathbf{type} \mathcal{C} = \overline{R}; \Delta_2 \}$ and $\Delta_3 = \{pn \in \Delta_1 \cup \Delta_2 \mid \overline{S} \subseteq P(pn) \wedge \overline{R} \subseteq P(pn)\}$. By (T-COMP), $\Gamma; \Delta \vdash t' : T$.

If (E-COMP2) applies, then $s_2 \rightarrow s'_2$, and by (T-COMP), $\Gamma; \Delta \vdash s_2 : T_1 \Rightarrow U_1 \{ \mathbf{type} \mathcal{C} = \overline{R}; \Delta_2 \}$, and $t' = v_1 \mathbf{compose} s'_2$. By the induction hypothesis, $\Gamma; \Delta \vdash s'_2 : T_1 \Rightarrow U_1 \{ \mathbf{type} \mathcal{C} = \overline{R}; \Delta_2 \}$. Since (E-COMP2) applies, $s_1 = v_1$, so by (T-COMP), we know that $\Gamma; \Delta \vdash v_1 : U_1 \Rightarrow T_2 \{ \mathbf{type} \mathcal{C} = \overline{S}; \Delta_1 \}$ and $\Delta_3 = \{pn \in \Delta_1 \cup \Delta_2 \mid \overline{S} \subseteq P(pn) \wedge \overline{R} \subseteq P(pn)\}$. By (T-COMP), $\Gamma; \Delta \vdash t' : T$.

If (E-COMP3) applies, then $s_1 = \mathbf{spore} \{ \overline{x : S = v}; \Delta_1; (y : U_1) \Rightarrow t_2 \}$ and $s_2 = \mathbf{spore} \{ \overline{y : R = w}; \Delta_2; (z : T_1) \Rightarrow u_1 \}$ and $\Delta_3 = \{p \mid p \in \Delta_1, \Delta_2. \overline{S} \subseteq P(p) \wedge \overline{R} \subseteq P(p)\}$. By (E-COMP3), $t' = \mathbf{spore} \{ \overline{x : S = v}, \overline{y : R = w}; \Delta_3; (z : T_1) \Rightarrow \mathbf{let} x = u_1 \mathbf{in} [y \mapsto x]t_2 \}$.

First, we show that $\forall v_i \in \overline{v}. \Gamma; \Delta \vdash v_i : S_i$ and $\forall w_i \in \overline{w}. \Gamma; \Delta \vdash w_i : R_i$. This follows from the fact that s_1 and s_2 are well-typed spores and (T-SPORE).

Second, we show that $\overline{x : S}, \overline{y : R}, z : T_1; \Delta \vdash \mathbf{let} x = u_1 \mathbf{in} [y \mapsto x]t_2 : T_2$. By (T-LET), we need to show that $\overline{x : S}, \overline{y : R}, z : T_1; \Delta \vdash u_1 : U_1$ and $\overline{x : S}, \overline{y : R}, z : T_1, x : U_1; \Delta \vdash [y \mapsto x]t_2 : T_2$. The former follows from (T-SPORE) and Lemma 2.5. To prove the latter: given that s_1 is well-typed, by (T-SPORE) we have that $\overline{x : S}, y : U_1 \vdash t_2 : T_2$. By Lemma 2.5, $\overline{x : S}, y : U_1, x : U_1 \vdash t_2 : T_2$. By Lemma 2.4, $\overline{x : S}, x : U_1 \vdash [y \mapsto x]t_2 : T_2$. By Lemma 2.5, $\overline{x : S}, \overline{y : R}, z : T_1, x : U_1; \Delta \vdash [y \mapsto x]t_2 : T_2$.

Third, we show that $\forall pn \in \Delta, \Delta_3. \overline{S} \subseteq P(pn) \wedge \overline{R} \subseteq P(pn)$. Since s_1 is well-typed, we have $\forall pn \in \Delta, \Delta_1. \overline{S} \subseteq P(pn)$. Since s_2 is well-typed, we

have $\forall pn \in \Delta, \Delta_2. \overline{R} \subseteq P(pn)$. Moreover, we have that $\Delta_3 = \{p \mid \overline{p} \in \Delta_1, \Delta_2. \overline{S} \subseteq P(p) \wedge \overline{R} \subseteq P(p)\}$. Thus, $\forall pn \in \Delta, \Delta_3. \overline{S} \subseteq P(pn) \wedge \overline{R} \subseteq P(pn)$.

By (T-SPORE) it follows from the previous three subgoals that $\Gamma; \Delta \vdash t' : T$.

□

2.5 Relation to spores in Scala

The soundness proof (see Section 2.4) of the formal type system guarantees several important properties for well-typed programs which closely correspond to the pragmatic model of spores in Scala:

1. Application of spores: for each property name pn , it is ensured that the dynamic types of all captured variables are contained in the type family pn maps to ($P(pn)$).
2. Dynamically, a spore only accesses its parameter and the variables in its header.
3. The properties computed for a composition of two spores is a safe approximation of the properties that are dynamically required.

2.6 Excluded types

This section shows how the formal model can be extended with excluded types as described above (see Section ??). Figure 6 shows the syntax extensions: first, spore terms and values are augmented with a sequence of excluded types; second, spore types and abstract spore types get another member $\mathbf{type} \mathcal{E} = \overline{T}$ specifying the excluded types.

$t ::= \dots$		terms
$\mathbf{spore} \{ \overline{x : T = t} ; \overline{T}; \overline{pn}; (x : T) \Rightarrow t \}$		spore
$v ::= \dots$		values
$\mathbf{spore} \{ \overline{x : T = v} ; \overline{T}; \overline{pn}; (x : T) \Rightarrow t \}$		spore value
$\mathcal{S} ::= T \Rightarrow T \{ \mathbf{type} \mathcal{C} = \overline{T} ; \mathbf{type} \mathcal{E} = \overline{T} ; \overline{pn} \}$		spore type
$T \Rightarrow T \{ \mathbf{type} \mathcal{C} ; \mathbf{type} \mathcal{E} = \overline{T} ; \overline{pn} \}$		abstract spore type

Figure 6: Core language syntax extensions

Figure 7 shows how the subtyping rules for spores have to be extended. Rule S-ESPORE requires that for each excluded type T' in the supertype, there must be an excluded type T in the subtype such that $T' <: T$. This means that by excluding type T , subtypes like T' are also prevented from being captured.

Figure 8 shows the extensions to the typing rules. Rule T-ESPORE additionally requires that none of the captured types \overline{S} is a subtype of one of the types contained in the excluded types \overline{U} . The excluded types are recorded in the type of the spore. Rule T-ECOMP computes a new set of excluded types \overline{V}

$$\begin{array}{c}
\text{S-ESPORE} \\
\frac{\overline{pn'} \subseteq \overline{pn} \quad T_2 <: T_1 \quad R_1 <: R_2 \quad M_1 = M_2 \vee M_2 = \mathbf{type} \mathcal{C} \quad \forall T' \in \overline{U'}. \exists T \in \overline{U}. T' <: T}{T_1 \Rightarrow R_1 \{ M_1 ; \mathbf{type} \mathcal{E} = \overline{U} ; \overline{pn} \} \\ <: T_2 \Rightarrow R_2 \{ M_2 ; \mathbf{type} \mathcal{E} = \overline{U'} ; \overline{pn'} \}} \\
\text{S-ESPOREFUN} \\
T_1 \Rightarrow R_1 \{ M ; E ; \overline{pn} \} <: T_1 \Rightarrow R_1
\end{array}$$

Figure 7: Subtyping extensions

based on both the excluded types and the captured types of t_1 and t_2 . Given that it is possible that one of the spores captures a type that is excluded in the other spore, the type of the result spore excludes only those types that are guaranteed not be captured.

$$\begin{array}{c}
\text{T-ESPORE} \\
\frac{\forall s_i \in \overline{s}. \Gamma; \Delta \vdash s_i : S_i \quad \overline{y : \overline{S}}, x : T_1; \Delta \vdash t_2 : T_2 \quad \forall pn \in \Delta, \Delta'. \overline{S} \subseteq P(pn) \quad \forall S_i \in \overline{S}. \forall U_j \in \overline{U}. \neg(S_i <: U_j)}{\Gamma; \Delta \vdash \mathbf{spore} \{ y : \overline{S} = s ; \overline{U}; \Delta'; (x : T_1) \Rightarrow t_2 \} : \\ T_1 \Rightarrow T_2 \{ \mathbf{type} \mathcal{C} = \overline{S} ; \mathbf{type} \mathcal{E} = \overline{U} ; \Delta, \Delta' \}} \\
\text{T-ECOMP} \\
\frac{\Gamma; \Delta \vdash t_1 : T_1 \Rightarrow T_2 \{ \mathbf{type} \mathcal{C} = \overline{S} ; \mathbf{type} \mathcal{E} = \overline{U} ; \Delta_1 \} \quad \Gamma; \Delta \vdash t_2 : U_1 \Rightarrow T_1 \{ \mathbf{type} \mathcal{C} = \overline{R} ; \mathbf{type} \mathcal{E} = \overline{U'} ; \Delta_2 \} \quad \Delta' = \{ pn \in \Delta_1 \cup \Delta_2 \mid \overline{S} \subseteq P(pn) \wedge \overline{R} \subseteq P(pn) \} \quad \overline{V} = (\overline{U} \setminus \overline{R}) \cup (\overline{U'} \setminus \overline{S})}{\Gamma; \Delta \vdash t_1 \mathbf{compose} t_2 : U_1 \Rightarrow T_2 \{ \mathbf{type} \mathcal{C} = \overline{S}, \overline{R} ; \mathbf{type} \mathcal{E} = \overline{V} ; \Delta' \}}
\end{array}$$

Figure 8: Typing extensions

Figure 9 shows the extensions to the operational semantics. Rule E-EAPPSPORE additionally requires that none of the captured types \overline{T} are contained in the excluded types \overline{U} . Rule E-ECOMP3 computes the set of excluded types of the result spore in the same way as in the corresponding type rule (T-ECOMP).

$$\begin{array}{c}
\text{E-EAPPSPORE} \\
\frac{\forall pn \in \overline{pn}. \overline{T} \subseteq P(pn) \quad \forall T_i \in \overline{T}. T_i \notin \overline{U}}{\text{spore } \{ \overline{x : T = v} ; \overline{U} ; \overline{pn} ; (x' : T) \Rightarrow t \} v' \rightarrow \quad \overline{[x \mapsto v][x' \mapsto v']}t} \\
\\
\text{E-ECOMP3} \\
\frac{\Delta = \{ p \mid p \in \overline{pn}, \overline{qn}. \overline{T} \subseteq P(p) \wedge \overline{S} \subseteq P(p) \} \quad \overline{V} = (\overline{U} \setminus \overline{S}) \cup (\overline{U}' \setminus \overline{T})}{\text{spore } \{ \overline{x : T = v} ; \overline{U} ; \overline{pn} ; (x' : T') \Rightarrow t \} \text{ compose} \\
\text{spore } \{ \overline{y : S = w} ; \overline{U}' ; \overline{qn} ; (y' : S') \Rightarrow t' \} \rightarrow \\
\text{spore } \{ \overline{x : T = v}, \overline{y : S = w} ; \overline{V} ; \Delta ; \\
(y' : S') \Rightarrow \text{let } z' = t' \text{ in } [x' \mapsto z']t \}}
\end{array}$$

Figure 9: Operational semantics extensions