A Type-and-Effect System for Object Initialization

FENGYUN LIU, EPFL, Switzerland

ONDŘEJ LHOTÁK, University of Waterloo, Canada

AGGELOS BIBOUDIS, EPFL, Switzerland

PAOLO G. GIARRUSSO, Delft University of Technology, Netherlands

MARTIN ODERSKY, EPFL, Switzerland

Every newly created object goes through several initialization states: starting from a state where all fields are uninitialized until all of them are assigned. Any operation on the object during its initialization process, which usually happens in the constructor via *this*, has to observe the initialization states of the object for correctness, i.e. only initialized fields may be used. Checking safe usage of *this* statically, without manual annotation of initialization states in the source code, is a challenge, due to aliasing and virtual method calls on *this*.

Mainstream languages either do not check initialization errors, such as Java, C++, Scala, or they defend against them by not supporting useful initialization patterns, such as Swift. In parallel, past research has shown that safe initialization can be achieved for varying degrees of expressiveness but by sacrificing syntactic simplicity.

We approach the problem by upholding *local reasoning about initialization* which avoids whole-program analysis, and we achieve *typestate polymorphism* via subtyping. On this basis, we put forward a novel type-and-effect system that can effectively ensure initialization safety while allowing flexible initialization patterns. We implement an initialization checker in the Scala 3 compiler and evaluate on several real-world projects.

CCS Concepts: • **Software and its engineering** → **Object oriented languages**; **Classes and objects**.

Additional Key Words and Phrases: Object initialization, Type-and-effect system

1 INTRODUCTION

Object-oriented programming is unsafe if objects cannot be initialized safely. The following code shows a simple initialization problem [1]:

```scala
class Hello {
  val message = "hello, " + name
  val name = "Alice"
}
println(new Hello().message))
```

---

[1]In the absence of special notes, the code examples are in Scala.

Authors' addresses: Fengyun Liu, EPFL, Switzerland; Ondřej Lhoták, University of Waterloo, Canada; Aggelos Biboudis, EPFL, Switzerland; Paolo G. Giarrusso, Delft University of Technology, Netherlands; Martin Odersky, EPFL, Switzerland.

The code above when run will print "hello, null" instead of "hello, Alice", as the field name is not initialized, thus holds the value null, when it is used in the second line.

The problem of safe initialization comes into existence since the introduction of object-oriented programming, and it is still a headache for programmers and language designers. Joe Duffy, in his popular blog post *on partially constructed objects* [Duffy 2010], wrote:

> Not only are partially-constructed objects a source of consternation for everyday programmers, they are also a challenge for language designers wanting to provide guarantees around invariants, immutability and concurrency-safety, and non-nullability.

1.1 Theoretical Challenges

Checking safe initialization of objects statically is becoming a challenge as the code in constructors is getting more complex. From past research [Fähndrich and Leino 2003; Fähndrich and Xia 2007; Gil and Shragai 2009; Qi and Myers 2009; Servetto et al. 2013; Summers and Müller 2011; Zibin et al. 2012], two initialization requirements are identified and commonly recognized.

Requirement 1: usage of "this" inside the constructor. The usage of already initialized fields in the constructor is safe and supported by almost all industrial languages. Based on an extensive study of over sixty thousand classes, Gil and Shragai [2009] report that over 8% of constructors include method calls on this. Method calls on this can be used to compute initial values for field initialization or serve as a private channel between the superclass and subclass.

Requirement 2: creation of cyclic data structures. Cyclic data structures are common in programming. For example, the following code shows the initialization of two mutually dependent objects:

```
1  class Parent { val child: Child = new Child(this) }
2  class Child(parent: Parent)
```

The objective is to allow cyclic data structures while preventing accidental premature usage of aliased objects. Accessing fields or calling methods on those aliased objects under initialization is an orthogonal concern, the importance of which is open to debate.

There are three theoretical challenges in addressing the requirements above.

Challenge 1: virtual method calls. While direct usage of already initialized fields via this is relatively easy to handle, indirect usage via virtual method calls poses a challenge. Such methods could be potentially overridden in a subclass, which makes it difficult to statically check whether it is safe to call such a method. This can be demonstrated by the following example:

```
1  abstract class AbstractFile {
2    def name: String
3    val extension: String = name.substring(4)
4  }
5  class RemoteFile(url:String) extends AbstractFile {
6    val localFile: String = url.hashCode // error
7    def name: String = localFile
8  }
```

According to the semantics of Scala (Java is the same), fields of a superclass are initialized before fields of a subclass, so initialization of the field extension proceeds before localFile. The field extension in the class AbstractFile is initialized by calling the abstract method name. The latter, implemented in the child class RemoteFile, accesses the uninitialized field localFile.

Challenge 2: aliasing. It is well-known that aliasing complicates program reasoning and it is challenging to develop practical type systems to support reasoning about aliasing [Clarke et al. 2013; Hogg et al. 1992]. It is also the case for safe initialization: if a field aliases an object under

initialization, we may not assume the field is fully initialized. This can be seen from the following example:

```
1  class Knot {
2    val self = this
3    val n: Int = self.n // error
4  }
```

In the code above, the field `self` is an alias of `this`, thus we may not use it as a fully initialized value. Aliasing may also happen indirectly through method calls, as the following code shows:

```
1  class Foo {
2    def f() = this
3    val n: Int = f().n // error
4  }
```

Challenge 3: typestate polymorphism. Every newly created object goes through several typestates [Strom and Yemini 1986]: starting from a state where all fields are uninitialized until all of them are assigned. If a method does not access any fields on *this*, then it should be able to be called on any typestate of *this*. For example, in the following class C, we should be able to call the method g regardless of the initialization state of `this`:

```
1  class C {
2    // ...
3    def g(): Int = 100
4  }
```

The challenge is how to support this feature succinctly without syntactic overhead.

1.2 Existing Work

1.2.1 Industrial Languages. Existing programming languages sit at two extremes. On one extreme, we find languages such as Java, C++, Scala, where programmers may use `this` as if it is fully initialized, devoid of any safety guarantee. On the other extreme, we find languages such as Swift, which ensures safe initialization, but is overly restrictive. The initialization of cyclic data structures is not supported, calling methods on `this` is forbidden, even the usage of already initialized fields is limited. For example, in the following Swift code, while the usage of x to initialize y is allowed, the usage of y to initialize f is illegal, which is a surprise:

```
1  class Position {
2      var x, y: Int
3      var f: () -> Int
4      init() {
5          x = 4
6          y = x * x                  // OK
7          f = { () -> Int in self.y } // error
8      }
9  }
```

1.2.2 Masked Types. Qi and Myers [2009] propose an expressive, flow-sensitive type-and-effect system [Lucassen and Gifford 1988] for safe initialization based on the concept of *masked types*.

A masked type $T \backslash f$ denotes objects of the type T, where the masked field f cannot be accessed. Each method has an effect signature of the form $\overline{M_1} \rightsquigarrow \overline{M_2}$, which means that the method can only

be called if `this` conforms to the masks $\overline{M_1}$, and the resulting masks for `this` after the call are $\overline{M_2}$. However, there are several obstacles to make the system practical.

First, the system incurs cognitive load and syntactic overhead. Many concepts are introduced in the system, such as *subclass masks*, *conditional masks*, *abstract masks*, each with non-trivial syntax. The paper mentions that inference can help remove the syntactic burden. However, it leaves open the formal development of such an inference system.

Second, the system, while expressive, is insufficient for simple and common use cases due to the missing support for *typestate polymorphism*. This can be seen from the following example, where we want the method g to be called for any initialization state of `this`:

```
1  class C { def g(): Int = 100 /* effect of g:  ∀M.M ⤳ M  */ }
```

As the method g can be called for `this` with any masks, we would like to give it the (imaginary) polymorphic effect signature $\forall M.M \rightsquigarrow M$, which is not supported. Even if an extension of the system supports polymorphic effect signatures, it will only incur more syntactic overhead.

1.2.3 The Freedom Model. Summers and Müller [2011] propose a light-weight, flow-insensitive type system for safe initialization, which we call *the freedom model*.

The freedom model classifies objects into two groups: *free*, that is under initialization, and *committed*, that is transitively initialized. Field accesses on free objects may get `null`, while committed objects can be used safely. To support typestate polymorphism, it introduces the typestate *unclassified*, which means either *free* or *committed*. In the system, typestate polymorphism becomes just *subtyping polymorphism*.

The freedom model supports the creation of cyclic data structures with light-weight syntax. However, the formal system does not address the usage of already initialized fields in the constructor. When an object is free, accessing its field will return a value of the type `unclassified C?`, which means the value could be `null`, free or committed. In the implementation, they introduce *committed-only fields* which can be assumed to be committed with the help of a dataflow analysis. However, the paper leaves open the formal treatment of the dataflow analysis. Our work will address the problem.

Moreover, the abstraction *free* is too coarse for some use cases. This is demonstrated by the following example:

```
1  class Parent {
2    var child = new Child(this)
3    var tag: Int = child.tag    // error in freedom model
4  }
5  class Child(parent: Parent @free) {
6    var tag: Int = 10
7  }
```

According to the freedom model, the expression `child` in line 3 will be typed as *free*, thus the type system cannot tell whether the field `child.tag` is initialized or not. But conceptually we know that all fields of `child` are initialized by the constructor of the class `Child`. In this work we propose a new abstraction to improve expressiveness in such cases.

1.3 Contributions

Our work makes contributions in four areas:

1. Better understanding of local reasoning about initialization. *Local reasoning about initialization* is a key requirement for simple and fast initialization systems. However, while prior work [Summers and Müller 2011] takes advantage of local reasoning about initialization to design

simple initialization systems, the concept of local reasoning about initialization is neither mentioned nor defined precisely. Identifying local reasoning about initialization as a concept with a better understanding enables it to be applied in the design of future initialization systems.

2. A more expressive type-based model. We propose a more expressive type-based model for initialization based on the abstractions *cold*, *warm* and *hot*. The introduction of the abstraction *warm* improves the expressiveness of the freedom model [Summers and Müller 2011], which classifies objects as either free (i.e. cold) or committed (i.e. hot).

3. A novel type-and-effect inference system. We propose a type-and-effect inference system for a practical fragment of the type-based model. Existing work usually depends on some unspecified inference or analysis to cut down syntactic overhead [Qi and Myers 2009; Summers and Müller 2011; Zibin et al. 2012]. We are the first to present a formal inference system on the problem of safe initialization. Meanwhile, to our knowledge, we are the first to demonstrate the technique of controlling aliasing in a type-and-effect system.

4. Implementation in Scala 3. We implement an initialization system in the Scala 3 compiler and evaluate it on several real-world projects. The system is capable of handling complex language features, such as inner classes, traits and functions.

2 LOCAL REASONING ABOUT INITIALIZATION

An important insight in the work of Summers and Müller [2011] is that *if a constructor is called with only transitively initialized arguments, the resulting object is transitively initialized*. We give this insight a name, *local reasoning about initialization*; it enables reasoning about initialization without the global analysis of a program, which is the key for simple and fast initialization systems. The insight can be generalized to the following:

> In a transitively initialized environment, the result value of an expression must be transitively initialized.

But how can we justify the insight? While a justification can be found in the soundness proof of the freedom model, it is obscured in a monolithic proof structure (see Lemma 1 of Summers and Müller [2011]). We provide a modular understanding of local reasoning about initialization by identifying three semantic properties, which we call *weak monotonicity*, *stackability* and *scopability*. Identifying local reasoning about initialization as a concept with a better understanding enables it to be applied in the design of future initialization systems. The properties can be explained roughly as follows:

- weak monotonicity: initialized fields continue to be initialized.
- stackability: all fields of a class should be initialized at the end of the class constructor.
- scopability: objects under initialization can only be accessed via static scoping.

To study the properties more formally, we first introduce a small language.

2.1 A Small Language

Our language resembles a subset of Scala having only top-level classes, mutable fields and methods.

$$
\begin{array}{rcll}
\mathcal{P} \in \text{Program} & ::= & (\overline{C}, D) \\
C \in \text{Class} & ::= & class\ C(\hat{f}{:}T)\ \{\ \overline{\mathcal{F}}\ \overline{\mathcal{M}}\ \} \\
\mathcal{F} \in \text{Field} & ::= & var\ f{:}T = e \\
e \in \text{Exp} & ::= & x\ \mid\ this\ \mid\ e.f\ \mid\ e.m(\overline{e})\ \mid\ new\ C(\overline{e})\ \mid\ e.f = e; e \\
\mathcal{M} \in \text{Method} & ::= & def\ m(\overline{x{:}T}) : T = e \\
S, T, U \in \text{Type} & ::= & C
\end{array}
$$

A program $\mathcal{P}$ is composed of a list of class definitions and an entry class. The entry class must have the form $class\ D\ \{\ def\ main()\ :\ T\ =\ e\ \}$. The program runs by executing e.

A class definition contains class parameters ($\hat{f}$:T), field definitions ($var\ f$:$T\ =\ e$) and method definitions ($def\ m(\overline{x{:}T})\ :\ T\ =\ e$). Class parameters are also fields of the class. All class fields are mutable. As a convention, we use f to range over all fields, and $\hat{f}$ to only range over class parameters.

An expression (e) can be a variable (x), a self reference ($this$), a field access ($e.f$), a method call ($e.m(\overline{e})$), a class instantiation ($new\ D(\overline{e})$), or a block expression ($e.f = e; e$). The block expression is used to avoid introducing the syntactic category of statements in the presence of assignments, which simplifies the presentation and meta-theory.

A method definition is standard. The body of a method is an expression, which could be a block expression to express a sequence of computations.

The following constructs are used in defining the semantics:

$$
\begin{aligned}
\Xi \in \text{ClassTable} &= \text{ClassName} \rightharpoonup \text{Class} \\
\sigma \in \text{Store} &= \text{Loc} \rightharpoonup \text{Obj} \\
\rho \in \text{Env} &= \text{Variable} \rightharpoonup \text{Value} \\
o \in \text{Obj} &= \text{ClassName} \times (\text{FieldName} \rightharpoonup \text{Value}) \\
l, \psi \in \text{Value} &= \text{Loc}
\end{aligned}
$$

We use ψ to denote the value of $this$, σ to denote the heap, and ρ to denote the local variable environment of the current stack frame.

The big-step semantics is expressed in the form $[\![e]\!]\,(\sigma, \rho, \psi) = (l, \sigma')$, which means that given the heap σ, environment ρ and value ψ for $this$, the expression e evaluates to the value l with the updated heap σ'. The semantics is standard, thus we omit detailed explanation and refer the reader to the technical report [Liu et al. 2020]. The only note is that non-initialized fields are represented by missing keys in the object, instead of a $null$ value. Newly initialized objects have no fields, and new fields are gradually inserted during initialization until all fields defined by the class have been assigned.

Note that this language does not enjoy initialization safety, and it is the task of later sections to make it safe. However, the language enjoys local reasoning about initialization.

2.2 Definitions

Definition 2.1 (reachability). We write $\sigma \vDash l \rightsquigarrow l'$ to say that an object l' is reachable from l in the heap σ. Reachability is formally defined according to the following rules:

$$
\frac{l \in dom(\sigma)}{\sigma \vDash l \rightsquigarrow l} \qquad \frac{\sigma \vDash l_0 \rightsquigarrow l_1 \quad (\_, \omega) = \sigma(l_1) \quad \exists f.\ \omega(f) = l_2 \quad l_2 \in dom(\sigma)}{\sigma \vDash l_0 \rightsquigarrow l_2}
$$

Definition 2.2 (reachability for set of locations).

$$
\begin{aligned}
\sigma \vDash L \rightsquigarrow l &\triangleq \exists l' \in L.\sigma \vDash l' \rightsquigarrow l \\
\sigma \vDash l \rightsquigarrow L &\triangleq \exists l' \in L.\sigma \vDash l \rightsquigarrow l'
\end{aligned}
$$

Definition 2.3 (cold). An object is cold if it exists in the heap, formally

$$
\sigma \vDash l : cold \quad \triangleq \quad l \in dom(\sigma)
$$

Definition 2.4 (warm). An object is warm if all its fields are assigned, formally

$$
\sigma \vDash l : warm \quad \triangleq \quad \exists (C, \omega) = \sigma(l) \bigwedge fields(C) \subseteq dom(\omega)
$$

Definition 2.5 (hot). An object is hot if all reachable objects are warm, formally

$$\sigma \vDash l : hot \quad \triangleq \quad l \in dom(\sigma) \bigwedge \forall l'.\sigma \vDash l \rightsquigarrow l' \implies \sigma \vDash l' : warm$$

From the definitions, it is easy to see that *hot* implies *warm* and *warm* implies *cold*.

2.3 Weak Monotonicity

The idea of *monotonicity* dates back to *heap monotonic typestates* by Fähndrich and Leino [2003]. There are, however, at least three different concepts of monotonicity.

Weak monotonicity means that initialized fields continue to be initialized. More formally, we may prove the following theorem:

THEOREM 2.6 (WEAK MONOTONICITY).

$$[\![e]\!] (\sigma, \rho, \psi) = (l, \sigma') \implies \sigma \preceq \sigma'$$

In the above, the predicate *weak monotonicity* ($\sigma \preceq \sigma'$) is defined below:

Definition 2.7 (Weak Monotonicity).

$$\sigma \preceq \sigma' \quad \triangleq \quad \forall l \in dom(\sigma). \, (C, \omega) = \sigma(l) \quad \implies \quad (C, \omega') = \sigma'(l) \bigwedge dom(\omega) \subseteq dom(\omega')$$

While weak monotonicity is sufficient to justify local reasoning about initialization, stronger monotonicity is required for initialization safety. For example, the freedom model [Summers and Müller 2011] enforces *strong monotonicity*:

$$\sigma \preceq \sigma' \quad \triangleq \quad \forall l \in dom(\sigma). \, \sigma \vDash l : \mu \implies \sigma' \vDash l : \mu$$

In the above, we abuse the notation by using μ to denote either *cold*, *warm* or *hot*. Strong monotonicity additionally ensures that hot objects continue to be hot. Therefore, it is always safe to use hot objects freely. However, to enforce safer usage of already initialized fields of non-hot objects, we need an even stronger concept, *perfect monotonicity*:

$$\sigma \preceq \sigma' \quad \triangleq \quad \forall l \in dom(\sigma). \quad (C, \omega) = \sigma(l) \implies$$
$$(C, \omega') = \sigma'(l) \quad \bigwedge \quad \forall f \in dom(\omega).\sigma \vDash \omega(f) : \mu \implies \sigma' \vDash \omega'(f) : \mu$$

In the above, we abuse the notation by writing directly $\omega'(f)$ to require that $dom(\omega) \subseteq dom(\omega')$. Perfect monotonicity in addition ensures that initialization states of object fields are monotone. It would be problematic if a field is initially assigned a hot value and later reassigned to a non-hot value.

2.4 Stackability

Conceptually, stackability ensures that all newly created objects during the evaluation of an expression e are *warm*, i.e. all fields of the objects are assigned. Formally, the insight can be proved as a theorem:

THEOREM 2.8 (STACKABILITY).

$$[\![e]\!] (\sigma, \rho, \psi) = (l, \sigma') \implies \sigma \ll \sigma'$$

The predicate $\sigma \ll \sigma'$ is defined below; it says that for any object in the heap σ', either the object is *warm*, or the object pre-exists in the heap σ.

Definition 2.9 (Stacking).

$$\sigma \ll \sigma' \quad \triangleq \quad \forall l \in dom(\sigma').\sigma' \vDash l : warm \bigvee l \in dom(\sigma)$$

Fig. 1. Each block represents the initialization duration of an object, i.e., from the creation of the object to the point where all fields are assigned.

Definite assignment [Gosling et al. 2015] can be used to enforce stackability in programming languages. Java, however, enforces definite assignment only for final fields.

If we push an object onto a stack when it comes into existence, and remove it from the stack when all its fields are assigned, we will find that the object to be removed is always at the top of the stack. This is illustrated in Figure 1.

2.5 Scopability

Scopability says that the access to uninitialized objects should be controlled by static scoping. Intuitively, it means that a method may only access pre-existing uninitialized objects through its environment, i.e. method parameters and `this`.

Objects under initialization are dangerous when used without care, therefore the access to them should be controlled. Scopability imposes discipline on accessing uninitialized objects. If we regard uninitialized objects as capabilities, then scopability restricts that there should be no side channels for accessing those capabilities. All accesses have to go through the explicit channel, i.e. method parameters and `this`. In contrast, global variables or control-flow effects such as algebraic effects may serve as side channels for teleporting values under initialization. To maintain local reasoning about initialization, an initialization system needs to make sure that only initialized values may travel by side channels.

More formally, we can prove the following theorem:

THEOREM 2.10 (SCOPABILITY).

$$\llbracket e \rrbracket \, (\sigma, \rho, \psi) = (l, \sigma') \Longrightarrow (\sigma, codom(\rho) \cup \{\, \psi \,\}) \prec (\sigma', \{\, l \,\}) \tag{1}$$

In the above, the predicate $(\sigma, L) \prec (\sigma', L')$ is defined below:

Definition 2.11 (Scoping). A set of addresses $L' \subseteq dom(\sigma')$ is *scoped* by a set of addresses $L \subseteq dom(\sigma)$, written $(\sigma, L) \prec (\sigma', L')$, is defined as follows

$$(\sigma, L) \prec (\sigma', L') \quad \triangleq \quad \forall l \in dom(\sigma). \quad \sigma' \vDash L' \rightsquigarrow l \Longrightarrow \sigma \vDash L \rightsquigarrow l$$

The theorem means that if e evaluates to l, then every location l' reachable from l in the new heap is either fresh, in that it did not exist in the old heap, or it *was* reachable from $codom(\rho) \cup \psi$ in the *old heap*.

Note that in the definition of *scoping*, we use $\sigma \vDash L \rightsquigarrow l$ instead of $\sigma' \vDash L \rightsquigarrow l$. This is because in a language with mutation, l may no longer be reachable from L in σ' due to reassignment. This can be seen in Figure 2.

Fig. 2. Each circle represents an object and numbers are locations. An arrow means that an object holds a reference to another object. The thick circles and links on the right heap are new objects and links created during the execution. Due to scopability, we have $(\sigma, \{2\}) \lessdot (\sigma', \{8\})$. This means that if the result object 8 reaches any object which pre-exists in the heap σ, then the object must be reachable from object 2 in the heap σ. The object 7 which is reachable from the object 2 in the heap σ, is no longer reachable from object 2 in the heap σ' due to the removal of the link from object 1 to object 7.

The property of scopability holds intuitively, but its proof is not obvious at all. The subtlety is in proving the case $e_1.m(e_2)$. Suppose we have $[\![e_1]\!](\sigma_1, \rho, \psi) = (l_1, \sigma_2)$ and $[\![e_2]\!](\sigma_2, \rho, \psi) = (l_2, \sigma_3)$. By the induction hypothesis, we have $(\sigma_1, codom(\rho) \cup \{\psi\}) \lessdot (\sigma_2, l_1)$ and $(\sigma_2, codom(\rho) \cup \{\psi\}) \lessdot (\sigma_3, l_2)$. However, we do not know that $(\sigma_1, codom(\rho) \cup \{\psi\}) \lessdot (\sigma_3, l_1)$. We need some invariant saying that scoping relations are preserved. That invariant has to be carefully defined, as not all scoping relations are preserved due to reassignment. We refer the reader to the technical report for more detailed discussions [Liu et al. 2020].

2.6 Local Reasoning about Initialization

With weak monotonicity, stackability and scopability, we may prove the theorem of local reasoning about initialization, which says that only hot objects can be produced from an expression in a hot environment.

LEMMA 2.12 (LOCAL REASONING). *If* $(\sigma, L) \lessdot (\sigma', L')$, $\sigma \lll \sigma'$, $\sigma \preceq \sigma'$ *and* $\sigma \vDash L : hot$, *then we have* $\sigma' \vDash L' : hot$.

PROOF. Let's consider any object l that is reachable from L', i.e. $\sigma' \vDash L' \rightsquigarrow l$. Depending on whether $l \in dom(\sigma)$, there are two cases.

- **Case** $l \notin dom(\sigma)$.
  Using the fact that $\sigma \lll \sigma'$, we know $\sigma' \vDash l : warm$.

- **Case** $l \in dom(\sigma)$.
  Using the fact that $(\sigma, L) \lessdot (\sigma', L')$, we have $\sigma \vDash L \rightsquigarrow l$. From the premise $\sigma \vDash L : hot$, we have $\sigma \vDash l : warm$. From $\sigma \preceq \sigma'$, we have $\sigma' \vDash l : warm$.

In both cases, we have $\sigma' \vDash l : warm$. Then by definition, we have $\sigma' \vDash L' : hot$. □

THEOREM 2.13 (LOCAL REASONING). *If* $[\![e]\!](\sigma, \rho, \psi) = (l, \sigma')$ *and* $\sigma \vDash \{\psi\} \cup codom(\rho) : hot$, *then we have* $\sigma' \vDash l : hot$.

PROOF. Immediate from Lemma 2.12, whose preconditions are satisfied by Theorem 2.10, Theorem 2.6 and Theorem 2.8. □

This theorem echoes the insight in the freedom model [Summers and Müller 2011]: if a constructor is called with all arguments committed, then the constructed object is also committed.

Table 1. The three abstractions of initialization states.

cold	A cold object *may* have uninitialized fields.
warm	A warm object has all its fields initialized.
hot	A hot object has all its fields initialized and only reaches hot objects.

3 THE BASIC MODEL

In this section, we take advantage of local reasoning about initialization to develop a type system that ensures initialization safety of objects.

3.1 Types

From the last section, we see that there are three natural abstractions of initialization states, as summarized in Table 1. If we posit the abstractions *cold*, *warm* and *hot* as types, we arrive at a type system for safe initialization of objects, which we call *the basic model*. Types in the language have the form C^μ:

$$
\begin{aligned}
\Omega &::= \quad \{\, f_1, f_2, \dots \,\} \\
\mu &::= \quad cold \mid warm \mid hot \mid \Omega \\
T &::= \quad C^\mu
\end{aligned}
$$

The type C^Ω is introduced to support the usage of already initialized fields — Ω denotes the set of initialized fields. The type is well-formed if Ω contains only fields of the class C. In languages that are equipped with an annotation system, such as Java, the type C^μ can be written using annotations (e.g. C @*warm* and C @*cold*), while a type without annotation can be assumed to be *hot*. Types like C^Ω are mainly used internally in the type system, thus there is no need to write them explicitly.

A type C^{μ_1} is a subtype of another type C^{μ_2}, written $C^{\mu_1} <: C^{\mu_2}$, if $\mu_1 \sqsubseteq \mu_2$. The lattice for modes μ is defined below:

$$
hot \sqsubseteq \mu \qquad\qquad warm \sqsubseteq \Omega \qquad\qquad \Omega_1 \cup \Omega_2 \sqsubseteq \Omega_1 \qquad\qquad \mu \sqsubseteq cold
$$

The modes *hot* and *cold* are respectively bottom and top of the lattice, and Ω is in the middle.

Methods are now annotated with modes, i.e., in @μ *def* $m(\overline{x{:}T}) : T = e$, the mode μ means *this* has the type C^μ inside the method m of the class C. We will propose an inference system to avoid the annotations in Sections 4 and 5. The semantics of the language remain the same as the language introduced in section 2.

3.2 Type System

We present expression typing and definition typing in Figure 3 and Figure 4. In an expression typing judgment $\Gamma; T \vdash e : U$, T is the type for *this*. Note that for simplicity of presentation, the class table Ξ is omitted in expression typing judgments.

Both the rules T-New and T-Invoke take advantage of *local reasoning about initialization*. The rule T-SelHot capitalizes on the fact that a hot object may only reach hot objects. The rules T-SelWarm and T-SelObj enforce that field selection takes the declared type of the field.

The rule T-Block demands that we only reassign *hot* values to fields; that is how we enforce *perfect monotonicity* in the system. It also restricts that only initialized fields may be mutated. The motivation for this restriction is to reject programs like class C { a = 10; var a = 5 }.

When type checking a program $(\overline{C}, D)$, the rule T-Prog ensures that every class is well-typed, and the entry class D has the expected form. In checking a class, the rule T-Class first checks the field definitions, assuming the type C^{Ω_i} for *this*, where Ω_i contains already initialized fields. It also

Expression Typing

$$\boxed{\Gamma; T \vdash e : T}$$

$$\frac{\Gamma; T \vdash e : T_1 \qquad T_1 <: T_2}{\Gamma; T \vdash e : T_2} \tag{T-Sub}$$

$$\frac{x : U \in \Gamma}{\Gamma; T \vdash x : U} \tag{T-Var}$$

$$\Gamma; T \vdash this : T \tag{T-This}$$

$$\frac{\Gamma; T \vdash e : D^{hot} \qquad C^\mu = fieldType(D, f)}{\Gamma; T \vdash e.f : C^{hot}} \tag{T-SelHot}$$

$$\frac{\Gamma; T \vdash e : D^{warm} \qquad U = fieldType(D, f)}{\Gamma; T \vdash e.f : U} \tag{T-SelWarm}$$

$$\frac{\Gamma; T \vdash e : D^\Omega \qquad f \in \Omega \qquad U = fieldType(D, f)}{\Gamma; T \vdash e.f : U} \tag{T-SelObj}$$

$$\frac{\overline{T_i} = constrType(C) \qquad \Gamma; T \vdash e_i : C_i^{\mu_i} \qquad C_i^{\mu_i} <: T_i \qquad \mu = (\sqcup \mu_i) \sqcap warm}{\Gamma; T \vdash new\ C(\overline{e}) : C^\mu} \tag{T-New}$$

$$\frac{\begin{array}{c}\Gamma; T \vdash e : C^{\mu_0} \qquad (\mu_m, \overline{T_i}, D^{\mu_r}) = methodType(C, m) \\ \mu_0 \sqsubseteq \mu_m \qquad \Gamma; T \vdash e_i : D_i^{\mu_i} \qquad D_i^{\mu_i} <: T_i \qquad \mu = (\sqcup \mu_i = hot)?hot : \mu_r\end{array}}{\Gamma; T \vdash e.m(\overline{e}) : D^\mu} \tag{T-Invoke}$$

$$\frac{\Gamma; T \vdash e_1.f : C^\mu \qquad \Gamma; T \vdash e_2 : C^{hot} \qquad \Gamma; T \vdash e : T_1}{\Gamma; T \vdash e_1.f = e_2; e : T_1} \tag{T-Block}$$

Fig. 3. Expression typing of the basic model

ensures that each method is well-typed. In type checking a field definition *var f:T = e*, the rule
T-Field ensures that the expression *e* can be typed as *T* in an empty environment. In checking a
method, the rule T-Method checks that the method body *e* conforms to the method return type *S*,
assuming *this* to take the mode of the method.

The soundness theorem says that a well-typed program does not get stuck at runtime.

Theorem 3.1 (Soundness). *If* $\vdash \mathcal{P}$, *then* $\forall k.\ [\![\mathcal{P}]\!](k) \neq Error$

The meta-theory takes the approach of step-indexed definitional interpreters [Amin and Rompf
2017]. For a step-indexed interpreter, there are three possible outcomes: (1) time out; (2) error; (3) a
resulting value and an updated heap. Initialization safety is implied by soundness, as initialization
errors will cause the program to fail at runtime. We refer the reader to the technical report for more
details about the meta-theory [Liu et al. 2020].

Program Typing $\boxed{\vdash \mathcal{P}}$

$$\frac{\Xi = \overline{C \to C} \qquad \Xi(D) = class\ D\ \{\ @hot\ def\ main() : T = e\ \} \qquad \overline{\Xi \vdash C}}{\vdash (\overline{C}, D)}$$ (T-Prog)

Class Typing $\boxed{\Xi \vdash C}$

$$\frac{\Omega_0 = \overline{\hat{f}} \qquad \overline{\Xi; C^{\Omega_i} \vdash \mathcal{F}_i \quad \Omega_{i+1} = \Omega_i \cup \{\ f_i\ \}} \qquad \overline{\Xi; C \vdash \mathcal{M}}}{\Xi \vdash class\ C(\overline{\hat{f}:T})\ \{\ \overline{\mathcal{F}}\ \overline{\mathcal{M}}\ \}}$$ (T-Class)

Field Typing $\boxed{\Xi; C^{\Omega} \vdash \mathcal{F}}$

$$\frac{\emptyset; C^{\Omega} \vdash e : T}{\Xi; C^{\Omega} \vdash var\ f : T = e}$$ (T-Field)

Method Typing $\boxed{\Xi; C \vdash \mathcal{M}}$

$$\frac{\overline{x:T}; C^{\mu} \vdash e : S}{\Xi; C \vdash @\mu\ def\ m(\overline{x:T}) : S = e}$$ (T-Method)

Fig. 4. Definition typing of the basic model

3.3 Typestate Polymorphism

A key design decision of the type system is to embrace *flow-insensitivity*. This follows an insight from Summers and Müller [2011] that we may achieve *typestate polymorphism* via subtyping in a flow-insensitive system.

Otherwise, if the system were flow-sensitive, we would have to track the change of typestates of this inside a method. Suppose we track the changes of a method m with $C^{\mu_1} \to C^{\mu_2}$, which means that the method m requires this to conform to C^{μ_1} before the call, and this takes the typestate C^{μ_2} after the call, similar to what is done by Qi and Myers [2009]. This creates a difficulty for methods that can be called for any typestates of this, as the following example shows:

```
1  class C {
2    // ...
3    def g(): Int = 100 //  ∀μ.Cᵘ → Cᵘ
4  }
```

In the code above, the method g can be called for any typestate of this. Representing the fact in the system would require *parametric polymorphism*, which complicates the solution. In fact, the system proposed by Qi and Myers [2009] does not support typestate polymorphism and thus invalidates such simple use cases.

3.4 The Discipline of Authority

An important discipline that our type system follows implicitly is what we call *the discipline of authority*. The discipline says that we may only consider the initialization state of an object to be advanced at field initialization points in the class constructor, but not at arbitrary assignments to fields.

The issue can be illustrated with the following example:

```
1  class C(a: A @cold) {
2    var x: A @cold = { foo(); a }
3    def foo() = { this.x = new A; this.a = new A }
4  }
```

The @cold annotation on the field x says that it may be initialized with a cold object. However, this annotation does not remove the transitive nature of a hot reference: given a reference r of type C, although r.x is *allowed* to hold a cold object, the reference r can be considered hot *only when* r.x actually contains a hot object.

In this example, the method foo() assigns a new (hot) object to fields x and a, so one might think that this should be considered hot (i.e., fully initialized) after a call to foo(). The initializer of x calls foo() and then returns the cold value a as the initial value for x. Although this.x holds a hot object immediately after the call to foo() in the initializer, this is no longer the case after the field initialization of x. The discipline of authority prevents us from considering this to be hot after the call to foo(). Without it, this example would violate monotonicity, in that the state of this would transition from hot to warm after the assignment.

Our type system by design guarantees that the this object still takes the typestate $C^{\{a\}}$ after the call to foo() (indicating that field a is initialized but x is not), as it only advances the initialization state of an object at field initialization points, never at field assignment points. This can be seen in the typing rule T-BLOCK: suppose e_1 is this, after the assignment we do not change the typestate of this in checking e (it remains to be T).

While the discipline of authority is followed implicitly by our type system, it arises as an explicit proof obligation in the meta-theory, and we do verify it for our type system. The meta-theory is based on *store typing* [Pierce 2002, Chapter 13], which is an abstraction of the concrete heap. We use Σ to range over store typings, which are maps from locations to types, i.e. $Loc \rightharpoonup Type$. In the meta-theory, we need to define the following predicate:

$$\Sigma \triangleright \Sigma' \quad \triangleq \quad \forall l \in dom(\Sigma).\Sigma(l) = C^\Omega \implies \Sigma'(l) = C^\Omega$$

In the above, Σ and Σ' refer to the store typings before and after evaluating an expression when the predicate is used. In the proof, we need to show that in evaluating an expression, if the object at location l is considered to have the initialization state C^Ω before the evaluation of an expression, it must be considered to still have the initialization state C^Ω after the evaluation of the expression.

Note that the definition of the predicate $\Sigma \triangleright \Sigma'$ only talks about types of the form C^Ω. The store typing never contains types like C^{cold}; a value takes such a type by subtyping. For the type C^{warm}, the only possible next monotone state is C^{hot}, so it is impossible for monotonicity to fail. For the type C^{hot}, monotonicity guarantees that the type stays the same.

The discipline of authority prevents us from advancing the initialization states of existing objects during the evaluation of an expression. It leaves only the possibility to advance the initialization state of objects at field initialization points in the constructor. At the end of the class body when all fields are initialized, we promote the type of the fresh object to be warm. Its promotion to hot may be delayed until a group of cyclic objects becomes hot together, which is called a *commitment point* by Summers and Müller [2011].

We believe the discipline of authority is already necessary for a system that enforces strong monotonicity, such as the freedom model [Summers and Müller 2011], but it has not been made explicit in previous work.

4 TYPE-AND-EFFECT INFERENCE, INFORMALLY

The type system proposed in the last section depends on verbose annotations, which are an obstacle for its adoption in practice. In this section, we propose a type-and-effect inference system [Lucassen and Gifford 1988; Nielson et al. 1999] to significantly cut down the syntactic overhead.

We first discuss the design of the type-and-effect inference system informally with examples. The formal presentation of the inference system comes in Section 5.

4.1 Potentials and Effects

The general idea is to track field accesses in the program with *effects*, and then check that only initialized fields are accessed. Effects over-approximate the uninitialized fields that an expression could possibly access. If the effect of an expression is empty, it means that at runtime, it is impossible for the expression to access an uninitialized field.

To deal with aliasing of potentially uninitialized objects, we introduce the concept of *potentials*. Potentials over-approximate the set of uninitialized objects that an expression appearing in the program could evaluate to. A potential encodes an uninitialized object in the form of a path, such as $C.this$, $C.this.f$ or $C.this.m$. If an expression appearing in the program can evaluate to an uninitialized object, then the potential of the expression must include a path to that object. If an expression in the program has an empty set of potentials, then, at runtime, the expression can only evaluate to an object that is hot.

In this section, we motivate effects and potentials informally using examples. The full formal syntax of effects and potentials is given in Section 5.2.

Consider the following erroneous program, which accesses the field y before it is initialized:

```
1  class C {
2    var x: Int = this.y      // C.this.y!
3    var y: Int = 10
4  }
```

A natural idea to ensure safe initialization is to analyze the fields that are accessed at each step of initialization, and check that only initialized fields are accessed. This leads to the fundamental effect in initialization: **field access effect**, e.g. $C.this.f!$.

Fields may also be accessed *indirectly* through method calls, as the following code shows:

```
1  class C {
2    var x: Int = this.m() // C.this.m<>
3    var y: Int = 10
4    def m(): Int = this.y // C.this.y!
5  }
```

For this case, we may introduce method calls as effects, which act as placeholders for the actual effects that happen in the method: **method call effects**, e.g. $C.this.m\Diamond$.

If we first analyze effects of the method m and map the effect $C.this.m\Diamond$ to the set of effects $\{C.this.y!\}$, then we may effectively check the initialization error in the code above.

One subtlety is how to handle aliasing. We illustrate with the following example:

```
1  class C {
2    var self  = this       // potentials of "self": { C.this }
3    var x: Int = self.x  // effects of "self.x": { C.this.self.x!, C.self! }
4  }
```

In the code above, the field x is used via the alias self before it is initialized. To check such errors, we need a way to represent and track the aliasing information in the system. That leads us to the

concept of **potentials** introduced in the beginning of this section. In the code example above, we can represent the aliasing information of the field self as the potential $C.this$. Now an initialization checker may take advantage of the aliasing information and report an error for the code self.x. The system knows that self.x is essentially this.x, but the field x is not yet initialized, the code is thus rejected.

To enforce that we may only assign *hot* values to a field in field assignment (not field initialization), we introduce **promotion effects** that promote potentials to be hot, e.g., $C.this\uparrow$. The checking system will check that only hot objects are promoted. The following example illustrates the usage of the effect:

```
1  class C(ctx: Context} {
2    val buffer: Buffer = { this.m(); new Buffer }  // C.this.m<>
3    def m(): Unit      = { ctx.reporter = this }   // C.this↑
4  }
```

In the code above, the method call effect $C.this.m\Diamond$ incurs the promotion effect $C.this\uparrow$. That is, the method this.m() may only be called when this is hot, because it assigns this to a field. The system finds that at the point of the call this.m(), the value of this is not hot, so such promotion is illegal.

Semantically, potentials keep track of objects possibly under initialization in order to maintain a *directed segregation* of initialized objects and objects under initialization: objects under initialization may point to initialized objects, but not vice versa. A promotion effect means that the object pointed to by the potential ascends to the initialized world, and the system gives up on tracking it. The system will have to ensure that by the time this happens, the object is hot.

Note that field access $C.this.a!$ and field promotion $C.this.a\uparrow$ are different effects, because a field access does not necessarily promote the field, as demonstrated by the following example:

```
1  class C {
2    var a = this
3    var b = this.a // C.this.a! , but no promotion
4  }
```

In the code above, the field selection this.a produces the field access effect $C.this.a!$, but not the promotion effect $C.this.a\uparrow$, because the potential $C.this.a$ does not leak from the uninitialized world.

Aliasing and promotion may also happen through methods, as the following example shows:

```
1  class Reporter(ctx: Context) {
2    ctx.reporter = this.m()      //  Reporter.this.m↑
3    def m() = this               //  potentials of m: { Reporter.this }
4  }
```

The type-and-effect system knows that the return value of the method m aliases this, thus the promotion of this.m() at line 2 indirectly promotes this to hot.

A similar distinction is drawn on methods: (1) the method invocation effect $C.this.m\Diamond$ means that the method m is called with the receiver this; (2) the method promotion effect $C.this.m\uparrow$ means that the return value of the call this.m is promoted to hot.

4.2 Two-Phase Checking

A common issue in program analysis is how to deal with recursive methods. We tackle the problem with *two phase checking*.

Table 2. Effect summaries for the methods h and g.

method	effects	potentials
h	{ $Foo.this.g\Diamond$ }	{ $Foo.this.g$ }
g	{ $Foo.this.h\Diamond$ }	{ $Foo.this.h$ }

In the first phase, the system computes effect summaries for methods and fields. For example, given the following program:

```
1  class Foo {
2    var a: Int = h()
3    def h(): Int = g()
4    def g(): Int = h()
5  }
```

Effect summaries for the methods h and g are computed as shown in Table 2.

In the second phase, the system checks the class constructor to ensure that only initialized fields are accessed. While checking the method call h(), the analysis propagates the effects associated with the method h until it reaches the fixed point { $Foo.this.g\Diamond$, $Foo.this.h\Diamond$ }. As the set does not contain accesses to any uninitialized fields of this nor invalid promotions, the program passes the check. Note that the domain of effects has to be finite for the existence of the fixed point.

4.3 Full-Construction Analysis

Another common issue in analysis is how to handle virtual method calls. The approach we take is *full-construction analysis*: we treat the constructors of concrete classes as entry points, and check all super constructors as if they were inlined. The analysis spans the full duration of object construction. This way, all virtual method calls on this can be resolved statically. From our experience, full-construction analysis greatly improves user experience, as no annotations are required for the interaction between subclasses and superclasses.

The following problem also motivates us to check the full construction duration of an object, which is also known as the *fragile base class problem*:

```
1  class Base { def g(): String = "hello" }
2  class Foo extends Base { val a = this.g() }
3  class Bar extends Base {
4    val b: String = "b"
5    override def g(): String = this.b
6  }
```

This program is correct. However, if we follow a type-based approach like the freedom model [Summers and Müller 2011], in order to call g() in the class Foo, the method Base.g has to be annotated @free, so that it may not access any fields on this. For soundness, the overriding method Bar.g has to be annotated @free too: but now it may not access the field this.b in the body of the method Bar.g. This unnecessarily restricts expressiveness of the system.

Moreover, we believe it is the only practical way to handle complex language features such as properties and traits. In languages such as Scala and Kotlin, fields are actually properties, accesses of public field are dynamic method calls, as the following code shows:

```
1  class A { val a = "Bonjour"; val b: Int = a.size }
2  class B extends A { override val a = "Hi" }
3  new B
```

In the code above, when the constructor of class B calls the constructor of class A, the expression a.size will dynamically dispatch to read the field a declared in class B, not the field a declared in class A. This results in a null-pointer exception at runtime because at the time, the field a in class B is not yet initialized. Without full-construction analysis, it is difficult to make the analysis sound for the code above.

Closed World Assumption. Full-construction analysis does not assume a *closed world* in the sense that it does not depend on the program entry as the analysis entry point. In contrast, it takes constructors of concrete classes as analysis entry points. The analysis does require the code of constructors of superclasses to be available.

Modularity. While full-construction analysis is capable of handling language features like traits and properties, it pays the price of modularity in the sense that if a superclass is changed, the subclasses have to be recompiled. We believe this is a worthy price to pay. First, the coupling between a superclass and its subclasses is well-known in object-oriented programming. For example, if a superclass adds a new method, then all its subclasses have to be recompiled to check proper overriding. Second, the ideal granularity for modular checking is not classes, but projects. From our experience with real-world projects, most subtle initializations happen within the same project. Third, the type system presented in Section 3 can serve as a coarse-grained type specification at project boundaries.

4.4 Cyclic Data Structures

Cyclic data structures are supported with an annotation @*cold* on class parameters, as the following example demonstrates:

```
1  class Parent { val child: Child = new Child(this) }
2  class Child(parent: Parent @cold) {
3    val friend: Friend = new Friend(this.parent)
4  }
5  class Friend(parent: Parent @cold) { val tag = 10 }
```

The annotation @cold indicates that the actual argument to parent during object construction might not be initialized. The type-and-effect system will ensure that the field parent is not used directly or indirectly when instantiating Child. However, aliasing the field to another cold class parameter is fine, thus the code new Friend(this.parent) at line 3 is accepted by the system. This allows programmers to create complex aliasing structures during initialization.

Our system tracks the return value of new Child(this) as the set of potentials { *warm*[*Child*] }. All fields of a warm value are assigned, but they may hold values that are not fully initialized. The inference system also takes advantage of local reasoning about initialization (Section 2): the whole cyclic data structure becomes hot at the same time when the first object in the group, i.e. the instance of Parent, becomes warm. This is called *commitment point* in the work of Summers and Müller [2011].

4.5 Relationship with the Type System

The type-and-effect system is intended to serve as an inference system for the type system in Section 3. Although simpler, the type system there requires annotations and thus forms an obstacle for adoption in practice. Meanwhile, the type-and-effect system scales better to complex language features like properties, inner classes and functions, and integrates better with compilers as no changes to the type system of the compiler are needed.

That said, the type-and-effect system is based on the type system in Section 3, and can be regarded as an inference system for a fragment of the type system there. For example, consider the following code:

```
1  class C {
2    val d: D = new D(this)
3    def foo = this.n
4    foo
5    val n = 10
6  }
7  class D(c: C @cold) {
8    val tag = 10
9  }
```

The field d is associated with the potentials { $warm[D]$ }, so it may take the type D^{warm}. The method foo is associated with the effects { $this.n!$ }, which suggests that this should conform to the type $C^{\{\,n\,\}}$ when the method foo is called.

In practice, the type-and-effect system does not bother to compute the exact type annotations nor elaborate the program with such type annotations, because the type elaboration is not useful in later compiler phases. Instead, it only checks that all the effects are safe in the constructor.

The fragment of the type system that we identify demands that (1) *method arguments must be hot*, and (2) *non-hot class parameters must be annotated*. Note that the receiver of a method call, e.g. this, is not considered as an argument of the call. The fragment supports calling methods on *this* in the constructor, as well as creation of cyclic data structures. There are several considerations for the restrictions.

First, from practical experience, there is little need to use non-hot values as method arguments. Meanwhile, virtual method calls on this are allowed, which covers most use cases in practice [Gil and Shragai 2009].

Second, it agrees with good programming practices that values under initialization should not escape [Bloch 2008]. Therefore, when there is the need to pass non-hot arguments to a constructor, it is a good practice to mark them explicitly.

Third, demanding method arguments to be hot saves us from changing the core type system of a language to check safe overriding of virtual methods.

5 FORMALIZING TYPE-AND-EFFECT INFERENCE

In this section, we formalize the type-and-effect system presented informally in the last section. The full soundness proof of the system is presented in the technical report [Liu et al. 2020].

5.1 Syntax and Semantics

Our language is almost the same as the language introduced in section 2, except for the definition of class parameters. In a class definition like $class\ C(\overline{\hat{f}:T})\ \{\ \overline{\mathcal{F}}\ \overline{\mathcal{M}}\ \}$, we introduce *cold class parameters*, which have the syntax $\tilde{f}$. Cold class parameters may take a value that is not transitively initialized. A class parameter $\hat{f}$ is also a field of its defining class. By default, we use f to range over all fields, $\hat{f}$ over class parameters, and $\tilde{f}$ over cold class parameters.

The tilde annotation $\tilde{f}$ is only used in the type-and-effect system; it does not have runtime semantics. That is the only annotation that is required in the source code.

The semantics is the same as the language in section 2, we thus omit the details.

5.2 Effects and Potentials

As seen from Figure 5, the definition of potentials (π) and effects (ϕ) depends on *roots* (β). Roots are the shortest path that represents an alias of a value that may not be transitively initialized. There are three roots in the system:

- $C.this$ represents an alias of *this* inside class C.
- $warm[C]$ represents an alias of a value of class C, all fields of which are assigned, but which may not be transitively initialized.
- *cold* represents a value whose initialization status is unknown. It is used to represent the potentials of cold class parameters. Field accesses or method calls on such an object are forbidden.

Potentials (π) represent aliasing information. They extend roots with field aliasing $\beta.f$ and method aliasing $\beta.m$. Field aliasing $\beta.f$ represents aliasing of the field f of β, while method aliasing $\beta.m$ represents aliasing of the return value of the method m with the receiver β.

Effects (ϕ) include field accesses, method calls, and promotions of possibly uninitialized values. A promotion effect is represented with $\pi\uparrow$, which enforces that the potential π is transitively initialized. The field access effect $\beta.f!$ means that the field f is accessed on β. The method call effect $\beta.m\Diamond$ means that the method m is called on β.

There are three helpers for the creation of potentials and effects:

- Field selection: $select(\Pi, f)$
- Method call: $call(\Pi, m)$
- Class instantiation: $init(C, \overline{\hat{f}_i = \Pi_i})$

They are used in expression typing to summarize the potentials and effects of expressions. A key to understanding the definitions is that the promotion effect $\pi\uparrow$ is the same as saying that π should be hot, and the result of an expression that has the empty set of potentials is hot.

Bounded Length. To make sure that the domain of effects and potentials is finite, the current system restricts the maximum length of potentials and effects to be two. In the implementation (Section 6), the maximum length of effects is 3. The bound is chosen in order to support calling methods on inner class instances, which is relatively common in Scala.

If the length of potentials exceeds the limit, the system checks that the potential is *hot* by producing a promotion effect. This can be seen from the last line of the definitions of the helper methods *select* and *call*.

Limiting the length will lead to incompleteness relative to the type system presented in Section 3. This does not pose a problem in practice (Section 7), due to the fact that fields usually hold hot values and methods return hot values. On the other hand, if it becomes an issue, the user may write explicit type annotations and the inference system can be extended to take advantage of the explicit type annotations.

5.3 Expression Typing

Expression typing (Figure 6) has the form $\Gamma; C \vdash e : D \,!\, (\Phi, \Pi)$, which means that the expression e in class C under the environment Γ, can be typed as D, and it produces effects Φ and has the potentials Π. Generally, when typing an expression, the effects of sub-expressions will **accumulate**, while potentials may be **refined** (via selection), or **promoted** (used as arguments to methods or assigned to a field).

The definitions assume several helper methods, such as $fieldType(C, f)$, $methodType(C, m)$ and $constrType(C)$, to look up in class table Ξ the type, respectively, of field $C.f$, of method $C.m$ and of the constructor of C.

Potentials and Effects

$$
\begin{array}{rclr}
T & ::= & C \mid D \mid E \mid \cdots & \text{type} \\
\beta & ::= & C.this \mid warm[C] \mid cold & \text{root} \\
\pi & ::= & \beta \mid \beta.f \mid \beta.m & \text{potential} \\
\Pi & ::= & \{ \pi_1, \pi_2, \cdots \} & \text{potentials} \\
\phi & ::= & \pi\uparrow \mid \beta.f! \mid \beta.m\Diamond & \text{effect} \\
\Phi & ::= & \{ \phi_1, \phi_2, \cdots \} & \text{effects} \\
\Omega & ::= & \{ f_1, f_2, \cdots \} & \text{fields} \\
\Delta & ::= & \overline{f_i \mapsto (\Phi_i, \Pi_i)} & \text{field summary} \\
\mathcal{S} & ::= & \overline{m_i \mapsto (\Phi_i, \Pi_i)} & \text{method summary} \\
\mathcal{E} & ::= & \overline{C \mapsto (\Delta, \mathcal{S})} & \text{effect table}
\end{array}
$$

Select

$$
\begin{array}{rcl}
select(\Pi, f) & = & \Pi.map(\pi \Rightarrow select(\pi, f)).reduce(\oplus) \\
select(\beta, \tilde{f}) & = & (\emptyset, \{cold\}) \\
select(\beta, \hat{f}) & = & (\emptyset, \emptyset) \\
select(\beta, f) & = & (\{\beta.f!\}, \{\beta.f\}) \text{ where } \beta \neq cold \\
select(cold, f) & = & (\{cold\uparrow\}, \emptyset) \\
select(\pi, f) & = & (\{\pi\uparrow\}, \emptyset) \text{ where } \pi = \beta.f \text{ or } \pi = \beta.m
\end{array}
$$

Call

$$
\begin{array}{rcl}
call(\Pi, m) & = & \Pi.map(\pi \Rightarrow call(m, \pi)).reduce(\oplus) \\
call(\beta, m) & = & (\{\beta.m\Diamond\}, \{\beta.m\}) \text{ where } \beta \neq cold \\
call(cold, m) & = & (\{cold\uparrow\}, \emptyset) \\
call(\pi, m) & = & (\{\pi\uparrow\}, \emptyset) \text{ where } \pi = \beta.f \text{ or } \pi = \beta.m
\end{array}
$$

Init

$$
\begin{array}{rcl}
init(C, \overline{\hat{f_i} = \Pi_i}) & = & (\cup\overline{\Pi_{k \neq j}\uparrow}, \{warm[C]\}) \text{ if } \exists \tilde{f_j}, \Pi_j \neq \emptyset \\
init(C, \overline{\hat{f_i} = \Pi_i}) & = & (\cup\overline{\Pi_i\uparrow}, \emptyset)
\end{array}
$$

Helpers

$$
\begin{array}{rcl}
\Pi\uparrow & = & \{ \pi\uparrow \mid \pi \in \Pi \} \\
(A_1, A_2) \oplus (B_1, B_2) & = & (A_1 \cup B_1, A_2 \cup B_2)
\end{array}
$$

Fig. 5. Potentials and Effects

5.4 Definition Typing

Definition typing (Figure 7) defines how programs, classes, fields and methods are checked. The checking happens in two phases:

(1) *first phase*: conventional type checking is performed and effect summaries are computed;
(2) *second phase*: effect checking is performed to ensure initialization safety.

The two-phase checking is reflected in the typing rule T-Prog. To type check a program $(\overline{C}, D)$, first each class is type checked separately for well-typing and the effect summary for fields Δ_c and methods $\mathcal{S}_c$ is computed using class typing $\Xi \vdash C \,!\, (\Delta, \mathcal{S})$. The result of class typing is stored in the effect table $\mathcal{E}$, which is then used for modular effect checking of each class. Effect checking is

Expression Typing

$$\boxed{\Gamma; C \vdash e : D \ ! \ (\Phi, \Pi)}$$

$$\frac{x : D \in \Gamma}{\Gamma; C \vdash x : D \ ! \ (\emptyset, \emptyset)} \quad \text{(T-Var)}$$

$$\Gamma; C \vdash this : C \ ! \ (\emptyset, \{C.this\}) \quad \text{(T-This)}$$

$$\frac{\Gamma; C \vdash e : D \ ! \ (\Phi, \Pi) \qquad (\Phi', \Pi') = select(\Pi, f) \qquad E = fieldType(D, f)}{\Gamma; C \vdash e.f : E \ ! \ (\Phi \cup \Phi', \Pi')} \quad \text{(T-Sel)}$$

$$\frac{\begin{array}{c} \Gamma; C \vdash e_0 : E_0 \ ! \ (\Phi, \Pi) \qquad \overline{\Gamma; C \vdash e_i : E_i \ ! \ (\Phi_i, \Pi_i)} \\ (\overline{x_i{:}E_i}, D) = methodType(E_0, m) \qquad (\Phi', \Pi') = call(\Pi, m) \end{array}}{\Gamma; C \vdash e_0.m(\overline{e}) : D \ ! \ (\Phi \cup \overline{\Phi_i} \cup \overline{\Pi_i\uparrow} \cup \Phi', \Pi')} \quad \text{(T-Call)}$$

$$\frac{\overline{\hat{f_i}{:}E_i} = constrType(C) \qquad \overline{\Gamma; C \vdash e_i : E_i \ ! \ (\Phi_i, \Pi_i)} \qquad (\Phi', \Pi') = init(C, \overline{\hat{f_i} = \Pi_i})}{\Gamma; C \vdash new \ C(\overline{e}) : C \ ! \ (\cup \overline{\Phi_i} \cup \Phi', \Pi')} \quad \text{(T-New)}$$

$$\frac{\begin{array}{c} \Gamma; C \vdash e_0 : E_0 \ ! \ (\Phi_0, \Pi_0) \qquad E_1 = fieldType(E_0, f) \\ \Gamma; C \vdash e_1 : E_1 \ ! \ (\Phi_1, \Pi_1) \qquad \Gamma; C \vdash e_2 : E_2 \ ! \ (\Phi_2, \Pi_2) \end{array}}{\Gamma; C \vdash e_0.f = e_1; e_2 : E_2 \ ! \ (\Phi_0 \cup \Phi_1 \cup \Pi_1\uparrow \cup \Phi_2, \Pi_2)} \quad \text{(T-Block)}$$

Fig. 6. Expression Typing

performed modularly on each class with the help of the effect table $\mathcal{E}$. The typing rule T-Prog also checks that the entry class D is well-typed.

When type checking a class, the rule T-Class checks that the body fields and methods are well-typed and computes the associated effects and potentials. The effects and potentials associated with a field are the effects and potentials of its initializer (the right-hand-side expression). The effects and potentials associated with a method are the effects and potentials of the body expression of the method. The effect summaries are used during the second phase in T-Check, which checks that given the already initialized fields, the effects on the right-hand-side of each field are allowed.

The typing rule T-Field checks the right-hand-side expression e in an empty typing environment, as there are no variables in a class body (class parameters are fields of their defining class). In the typing rule T-Method, the method parameters $\overline{x : D}$ are used as the typing environment to check the method body.

5.5 Effect Checking

The effect checking judgment $\mathcal{E}; C^{\Omega} \vdash \Phi$ (Figure 8) means that the effects Φ are permitted inside class C when the fields in Ω are initialized. It first checks that there is no promotion of *this* in the closure of the effects, as the underlying object is not transitively initialized, so promotion is illegal. Then it checks that each accessed field is in the set Ω, i.e., only initialized fields are used.

The closure of effects is presented in a declarative style for clarity, but it has a straight-forward algorithmic interpretation: it just propagates the effects recursively until a fixed-point is reached. The fixed-point always exists as the domain of effects and potentials is finite for any given program.

Program Typing $\boxed{\vdash \mathcal{P}}$

$$\frac{\Xi = \overline{C \mapsto C} \qquad \Xi(D) = class\ D\ \{\ def\ main()\ :\ T = e\ \}}{\overline{\Xi \vdash C\ !\ (\Delta_c, \mathcal{S}_c)} \qquad \mathcal{E} = \overline{C \mapsto (\Delta_c, \mathcal{S}_c)} \qquad \overline{\Xi; \mathcal{E} \vdash C}}{\vdash (\overline{C}, D)} \qquad \text{(T-Prog)}$$

Effect Checking $\boxed{\Xi; \mathcal{E} \vdash C}$

$$\frac{(\Delta, \_) = \mathcal{E}(C) \qquad \overline{(\Phi, \_) = \Delta(f_i) \qquad \mathcal{E}; C^{\{f_1, \cdots, f_{i-1}\}} \vdash \Phi}}{\Xi; \mathcal{E} \vdash class\ C(\widehat{f}:T)\ \{\ \overline{\mathcal{F}}\ \overline{\mathcal{M}}\ \}} \qquad \text{(T-Check)}$$

Class Typing $\boxed{\Xi \vdash C\ !\ (\Delta, \mathcal{S})}$

$$\frac{\overline{\Xi; C \vdash \mathcal{F}_i\ !\ (\Phi_i, \Pi_i)} \qquad \Delta = \overline{f_i \mapsto (\Phi_i, \Pi_i)} \qquad \overline{\Xi; C \vdash \mathcal{M}_i\ !\ (\Phi_i, \Pi_i)} \qquad \mathcal{S} = \overline{m_i \mapsto (\Phi_i, \Pi_i)}}{\Xi \vdash class\ C(\widehat{f}:T)\ \{\ \overline{\mathcal{F}}\ \overline{\mathcal{M}}\ \}\ !\ (\Delta, \mathcal{S})}$$

$$\text{(T-Class)}$$

Field Typing $\boxed{\Xi; C \vdash \mathcal{F}\ !\ (\Phi, \Pi)}$

$$\frac{\emptyset; C \vdash e : D\ !\ (\Phi, \Pi)}{\Xi; C \vdash var\ f : D = e\ !\ (\Phi, \Pi)} \qquad \text{(T-Field)}$$

Method Typing $\boxed{\Xi; C \vdash \mathcal{M}\ !\ (\Phi, \Pi)}$

$$\frac{\overline{x:D}; C \vdash e : E\ !\ (\Phi, \Pi)}{\Xi; C \vdash\ def\ m(\overline{x:D}) : E = e\ !\ (\Phi, \Pi)} \qquad \text{(T-Method)}$$

Fig. 7. Definition Typing

The main step in fixed-point computation is the propagation of effects and potentials. In effect propagation $\mathcal{E} \vdash \phi \rightsquigarrow \Phi$, field access $\beta.f!$ is an atomic effect, thus it propagates to the empty set. For a promotion effect $\pi\uparrow$, we first propagate the potential π to a set of potentials Π, and then promote each potential in Π. For a method call effect $C.this.m\Diamond$, we look up the effects associated with the method from the effect table.

In potential propagation $\mathcal{E} \vdash \pi \rightsquigarrow \Pi$, root potentials like $C.this$ propagate to the empty set, as they do not contain *proxy* aliasing information in the effect table. For a field potential like $C.this.f$, propagation just looks up the potentials associated with the field f from the effect table. For a method potential $C.this.m$, propagation looks up the potentials associated with the method m from the effect table.

The soundness theorem says that a well-typed program does not get stuck at runtime.

THEOREM 5.1 (SOUNDNESS). *If* $\vdash \mathcal{P}$, *then* $\forall k.\ [\![\mathcal{P}]\!](k) \neq Error$

The meta-theory takes the approach of step-indexed definitional interpreters [Amin and Rompf 2017]. Initialization safety is implied by soundness, as initialization errors will cause the program

Propagate Potentials $\boxed{\mathcal{E} \vdash \pi \rightsquigarrow \Pi}$

$$\mathcal{E} \vdash \beta \rightsquigarrow \emptyset$$

$$\frac{(\Delta, \_) = \mathcal{E}(C) \qquad (\_, \Pi) = \Delta(f)}{\mathcal{E} \vdash C.this.f \rightsquigarrow \Pi} \qquad \frac{\mathcal{E} \vdash C.this.f \rightsquigarrow \Pi \qquad \Pi' = [C.this \mapsto warm[C]]\Pi}{\mathcal{E} \vdash warm[C].f \rightsquigarrow \Pi'}$$

$$\frac{(\_, \mathcal{S}) = \mathcal{E}(C) \qquad (\_, \Pi) = \mathcal{S}(m)}{\mathcal{E} \vdash C.this.m \rightsquigarrow \Pi} \qquad \frac{\mathcal{E} \vdash C.this.m \rightsquigarrow \Pi \qquad \Pi' = [C.this \mapsto warm[C]]\Pi}{\mathcal{E} \vdash warm[C].m \rightsquigarrow \Pi'}$$

Propagate Effects $\boxed{\mathcal{E} \vdash \phi \rightsquigarrow \Phi}$

$$\mathcal{E} \vdash \beta.f! \rightsquigarrow \emptyset \qquad\qquad \frac{(\_, \mathcal{S}) = \mathcal{E}(C) \qquad (\Phi, \_) = \mathcal{S}(m)}{\mathcal{E} \vdash C.this.m\Diamond \rightsquigarrow \Phi}$$

$$\frac{\mathcal{E} \vdash \pi \rightsquigarrow \Pi}{\mathcal{E} \vdash \pi\uparrow \rightsquigarrow \Pi\uparrow} \qquad \frac{\mathcal{E} \vdash C.this.m\Diamond \rightsquigarrow \Phi \qquad \Phi' = [C.this \mapsto warm[C]]\Phi}{\mathcal{E} \vdash warm[C].m\Diamond \rightsquigarrow \Phi'}$$

Closure

$$\frac{\Phi \subseteq \Phi' \qquad \forall \phi \in \Phi'.\mathcal{E} \vdash \phi \rightsquigarrow \Phi'' \implies \Phi'' \subseteq \Phi'}{\Phi^c = \Phi'}$$

Check $\boxed{\mathcal{E}; \Omega; C \vdash \Phi}$

$$\frac{\beta\uparrow \notin \Phi^c \qquad \forall C.this.f! \in \Phi^c.\ f \in \Omega}{\mathcal{E}; C^\Omega \vdash \Phi}$$

Fig. 8. Effect Checking

to fail at runtime. We refer the reader to the technical report for more details about the meta-theory [Liu et al. 2020].

6 IMPLEMENTATION

Based on the type-and-effect inference system, we have implemented an initialization system for Scala. The implementation is already integrated in the Scala 3 compiler [Odersky et al. 2013] and available to Scala programmers via the compiler option `-Ycheck-init`.

The implementation supports inner classes, first-class functions, traits and properties. Instantiation of inner classes is supported without any annotations, as the following example shows:

```
1    class Trees {
2      private var counter = 0
3      class ValDef { counter += 1 }    // ok, counter is initialized
4      class EmptyValDef extends ValDef
5      val theEmptyValDef = new EmptyValDef
6    }
```

To make the example above work, a warm potential in the system takes the form $warm(C, \pi)$, where C is the concrete class of the object and π is the potential for the immediate outer of C. The current version of the system only allows creating cyclic data structures via inner classes; passing *this* as an argument to new-expressions is disallowed. To support the usage would require adding an annotation to the language, which involves the language improvement process, which we want to avoid in the initial version. We plan to support this in the next version following the solution outlined in the theory (Section 5).

To support first-class functions, we introduce the potential $Fun(\Phi, \Pi)$, where Φ is the set of effects to be triggered when the function is called, while Π is the set of potentials for the result of the function call. For example, this enables the following code, which is rejected in Swift:

```
1  class Rec {
2    val even = (n: Int) => n == 0 || odd(n - 1)
3    val odd = (n: Int) => n == 1 || even(n - 1)
4    val flag: Boolean  = odd(6)
5  }
```

In functional programming, the recursive binding construct *letrec* may introduce similar initialization patterns as the code above. With the latest checker [Reynaud et al. 2018], OCaml still does not support the code below in the construct *let rec*:

```
1  let rec even n = if n = 0 then true else odd (x - 1)
2      and odd n = if n = 0 then false else even (x - 1)
3      and flag = odd 3
```

Naive extension of the type-and-effect system can easily lead to non-termination of effect checking in practice. This can be demonstrated by the following example:

```
1  class B {
2    class C extends B
3    val c: C = new C
4  }
```

The code above involves an infinite sequence of constructor call effects of the form $\pi_i.init(C)$, where $\pi_0 = warm(C, this)$ and $\pi_i = warm(C, \pi_{i-1})$. To prevent infinite regression, we bound the depth of nested potentials. When the bound is exceeded, we over-approximate the nested potentials that exceed the bound with the potential *cold*.

One advantage of the type-and-effect system is that it integrates well with the compiler without changing the core type system. In contrast, integrating a type-based system in the compiler poses an engineering challenge, as the following example demonstrates:

```
1  class Knot {
2    val self: Knot @cold = this
3  }
```

In the code above, the type of the field self depends on *when* we ask for its type. If it is queried during the initialization of the object, then it has the type Knot @cold. Otherwise, it has the type Knot. We do not see a principled way to implement the type-based solution in the Scala 3 compiler.

7 EVALUATION

We evaluate the implementation on a significant number of real-world projects, with zero changes to the source code. The results of the experiment are shown in Figure 3. The first three columns show the size of the projects and warnings reported for each project:

- **KLOC** - the number of lines of code (KLOC) in the project checked by the system

Table 3. Experiment result. The column W/K is the number of warnings per KLOC, and the column W is the number of warnings issued for the corresponding project. Other columns are explained in the text.

Project	KLOC	W/K	W	X1	X2	X3	X4	A	B	C	D	E	F	G	H
DOTTY	106.0	0.73	77	742	447	146	350	7	16	2	32	0	3	4	13
INTENT	1.8	39.53	71	10	290	0	1	0	0	0	71	0	0	0	0
ALGEBRA	1.3	4.70	6	1	6	0	0	0	0	0	0	0	0	6	0
STDLIB213	43.6	0.62	27	231	104	8	99	14	0	4	2	0	1	6	0
SCALACHECK	5.5	1.08	6	39	70	6	83	0	0	0	6	0	0	0	0
SCALATEST	378.9	0.39	149	1037	718	18	664	0	0	8	114	0	8	19	0
SCALAXML	6.8	0.15	1	36	13	0	0	0	0	0	0	0	0	1	0
SCOPT	0.3	0.00	0	6	4	0	0	0	0	0	0	0	0	0	0
SCALAP	2.2	5.43	12	62	57	2	108	0	0	0	7	5	0	0	0
SQUANTS	14.1	0.00	0	9	0	0	0	0	0	0	0	0	0	0	0
BETTERFILES	2.8	0.00	0	17	1	0	0	0	0	0	0	0	0	0	0
SCALAPB	16.2	0.31	5	28	10	0	6	4	0	0	1	0	0	0	0
SHAPELESS	2.5	0.79	2	5	0	0	0	0	0	0	0	2	0	0	0
EFFPI	5.7	0.53	3	15	5	0	12	0	0	0	3	0	0	0	0
SCONFIG	21.8	0.60	13	70	43	0	8	13	2	2	0	0	1	6	2
MUNIT	2.7	1.13	3	32	73	1	13	0	0	0	2	0	0	0	1
SUM	612.1	0.61	375	2340	1841	181	1344	38	18	16	238	7	13	42	16

- **W/K** - the number of warnings issued by the system per KLOC
- **W** - the number of warnings issued by the system

We can see that for over 0.6 million lines of code, the system reports 375 warnings in total and the average is 0.61 warnings per KLOC. We can better interpret the data in conjunction with the following columns:

- **X1** - the number of field accesses on `this` during initialization
- **X2** - the number of method calls on `this` during initialization
- **X3** - the number of field accesses on *warm* objects during initialization
- **X4** - the number of method calls on *warm* objects during initialization

The data for the columns above are censused by the initialization checker, one per source location. Without type-and-effect inference, the system would have to issue one warning for each method call on `this` and warm objects [2], i.e., the counts in columns X2 and X4 would all become warnings. This would contribute more than 3000 warnings, an 8-fold increase in the number of warnings.

We manually analyzed all the warnings and classified them into 8 categories:

- **A** - Use `this` as constructor arguments, e.g. `new C(this)`
- **B** - Use `this` as method arguments, e.g. `call(this)`
- **C** - Use inner class instance as constructor arguments, e.g. `new C(innerObj)`
- **D** - Use inner class instance as method arguments, e.g. `call(innerObj)`
- **E** - Use uninitialized fields as by-name arguments
- **F** - Access non-initialized fields
- **G** - Call external Java or Scala 2 methods
- **H** - others

The warnings in category **A** and **C** are related to the creation of cyclic data structures. From Section 5, we know that such code patterns can be supported by declaring a class parameter to be

---

[2]If we ignore the fact that non-private field accesses are also method calls in Scala.

cold. The current implementation does not support any annotations yet, but we plan to introduce explicit annotations in the next version of the system.

Most of the warnings lie in the category **D**, which refers to cases like the following:

```scala
1   object Foo {
2     case class Student(name: String, age: Int)
3     call(Student("Jack", 30)              // should be OK, currently a warning
4   }
```

For the code above, our system currently issues a warning, as it only knows that the object created by `Student("Jack", 20)` is warm, while method arguments are required to be hot. Checking whether an inner class instance may be safely promoted to hot or not can be expensive if the inner class contains many fields and methods. However, these cases suggest that the system could be improved for common use cases that only involve small classes, such as the example above.

The category **E** refers to cases like the following, which is not supported currently:

```scala
1   def foo(x: => Int) = new A(x)
2   class A(init: => Int)
3   class Foo {
4     val a: A = foo(b) // category E
5     val b: Int = 100
6   }
```

As an over-approximation, we expect the warnings in category **F** are all false positives. However, to our delight, the system actually finds 8 true positives in ScalaTest, and one true positive in the Scala standard library. It also discovers two bugs in the Scala 3 compiler. We reported the bugs, and the bugs in the Scala 3 compiler have been fixed. The technical report contains more details about the bugs [Liu et al. 2020].

The category **G** involves method calls on `this` in the constructor, but the target method is compiled by Java or the Scala 2 compiler. The category **H** involves code that performs pattern matching on `this` or calls methods on `cold` values.

8 RELATED WORK

Our work takes inspiration from several milestone papers on the problem of initialization.

Fähndrich and Leino [2003] introduce raw types like $T^{\mathrm{raw}(S)}$ — a value of such a type is possibly under initialization, and all fields up to the superclass S are initialized. Class fields may not hold raw values, thus it does not support creating cyclic data structures. To overcome the limitation, they introduce *delayed types* [Fähndrich and Xia 2007]. The system ensures that the initialization of objects forms stacked time regions.

Qi and Myers [2009] introduce a flow-sensitive type-and-effect system for initialization based on masked types. The system is expressive, however, it leaves open the problem of typestate polymorphism and type-and-effect inference. Our work can be seen as an attempt to address the problems.

Summers and Müller [2011] show that initialization of cyclic data structures can be supported in a light-weight, flow-insensitive type system. The system cleverly uses subtyping to achieve typestate polymorphism. However, it leaves open the design of a dataflow analysis that enables the usage of already initialized fields. Our work effectively addresses the problem.

There is another main difference: our system favors *perfect monotonicity*, while the freedom model favors *strong monotonicity*. There are design trade-offs in both approaches. In our case, perfect monotonicity enables us to remove the abstraction *unclassified* and it is easy to safely use already initialized fields in the constructor. In contrast, the freedom model enables assigning a free

object to the field of another free object anywhere, while in our system it is only possible in the constructor at initialization points. More concretely, the following example is supported by the freedom model, but not by our system:

```
1  class A {
2    m(this)
3    var b: b = new B(this)
4    def m(a: A @free): Unit = { a.b = new B(a) } // !!
5  }
6  class B(a: A @free)
```

The assignment `a.b = new B(a)` in the method `m` will be rejected by our system, as `new B(a)` is a value under initialization (it holds a reference to a free value a). In our system, it is only possible to assign hot values to fields of cold objects, while in the freedom model it is possible to assign non-committed values to fields of non-committed values. Our design is based on our experience with Scala projects, where an object rarely escapes from its constructor and has its fields initialized elsewhere. Summers and Müller [2011] have similar observations (Section 8.1).

The Checker Framework enables many useful checkers for various properties of Java programs [Ernst and Ali 2010]. In particular, it implements and extends the freedom model. One major extension is the introduction of the annotation `UnknownInitialization`, which is in the same spirit as `warm`. A difference is that `warm` in our type-based model enjoys transitivity — a warm object may in turn contain warm fields. The initialization model in Checker Framework does not enjoy this kind of transitivity enabled by `warm`, despite the introduction of 4 annotations: `Initialized`, `UnderInitialization`, `UnknownInitialization` and `NotOnlyInitialized`.

The initialization in X10 [Zibin et al. 2012] employs an inter-procedural analysis to ensure safe initialization, which removes the annotation burden required when calling final or private methods on *this*. However, the analysis algorithm is not presented in the paper. To call virtual methods on *this*, annotations are required on method definitions.

The Billion-Dollar Fix [Servetto et al. 2013] introduces a new linguistic construct *placeholders* and *placeholder types* to support initialization of circular data structures. The work is orthogonal to the current work, in that we are constrained from introducing new language constructs and semantics.

ACKNOWLEDGMENTS

REFERENCES

Nada Amin and Tiark Rompf. 2017. Type soundness proofs with definitional interpreters. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, Giuseppe Castagna and Andrew D. Gordon (Eds.). ACM, 666–679. http://dl.acm.org/citation.cfm?id=3009866

Joshua Bloch. 2008. *Effective Java (2nd Edition) (The Java Series)* (2 ed.). Prentice Hall PTR, Upper Saddle River, NJ, USA.

Dave Clarke, James Noble, and Tobias Wrigstad (Eds.). 2013. *Aliasing in Object-Oriented Programming. Types, Analysis and Verification*. Lecture Notes in Computer Science, Vol. 7850. Springer. https://doi.org/10.1007/978-3-642-36946-9

Joe Duffy. 2010. On partially-constructed objects. http://joeduffyblog.com/2010/06/27/on-partiallyconstructed-objects/.

Michael D. Ernst and Mahmood Ali. 2010. Building and using pluggable type systems. In *Proceedings of the 18th ACM SIGSOFT International Symposium on Foundations of Software Engineering, 2010, Santa Fe, NM, USA, November 7-11, 2010*, Gruia-Catalin Roman and André van der Hoek (Eds.). ACM, 375–376. https://doi.org/10.1145/1882291.1882356

Manuel Fähndrich and K. Rustan M. Leino. 2003. Declaring and checking non-null types in an object-oriented language. In *Proceedings of the 2003 ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications, OOPSLA 2003, October 26-30, 2003, Anaheim, CA, USA*, Ron Crocker and Guy L. Steele Jr. (Eds.). ACM, 302–312. https://doi.org/10.1145/949305.949332

Manuel Fähndrich and K Rustan M Leino. 2003. Heap monotonic typestates. In *International Workshop on Aliasing, Confinement and Ownership in object-oriented programming (IWACO)*.

Manuel Fähndrich and Songtao Xia. 2007. Establishing object invariants with delayed types. In *Proceedings of the 22nd Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2007, October 21-25, 2007, Montreal, Quebec, Canada*, Richard P. Gabriel, David F. Bacon, Cristina Videira Lopes, and Guy L. Steele Jr. (Eds.). ACM, 337–350. https://doi.org/10.1145/1297027.1297052

Joseph Gil and Tali Shragai. 2009. Are We Ready for a Safer Construction Environment?. In *ECOOP 2009 - Object-Oriented Programming, 23rd European Conference, Genoa, Italy, July 6-10, 2009. Proceedings (Lecture Notes in Computer Science, Vol. 5653)*, Sophia Drossopoulou (Ed.). Springer, 495–519. https://doi.org/10.1007/978-3-642-03013-0_23

James Gosling, Bill Joy, Guy Steele, Gilad Bracha, and Alex Buckley. 2015. The Java Language Specification, Java SE 8 Edition.

John Hogg, Doug Lea, Alan Cameron Wills, Dennis de Champeaux, and Richard C. Holt. 1992. The Geneva convention on the treatment of object aliasing. *OOPS Messenger* 3, 2 (1992), 11–16. https://doi.org/10.1145/130943.130947

Fengyun Liu, Ondrej Lhoták, Aggelos Biboudis, Paolo G. Giarrusso, and Martin Odersky. 2020. Safe Initialization of Objects. (2020), 141. http://infoscience.epfl.ch/record/279970

John M. Lucassen and David K. Gifford. 1988. Polymorphic Effect Systems. In *Conference Record of the Fifteenth Annual ACM Symposium on Principles of Programming Languages, San Diego, California, USA, January 10-13, 1988*, Jeanne Ferrante and P. Mager (Eds.). ACM Press, 47–57. https://doi.org/10.1145/73560.73564

Flemming Nielson, Hanne Riis Nielson, and Chris Hankin. 1999. *Principles of program analysis*. Springer. https://doi.org/10.1007/978-3-662-03811-6

Martin Odersky et al. 2013. Dotty Compiler: A Next Generation Compiler for Scala. https://dotty.epfl.ch/.

Benjamin C. Pierce. 2002. *Types and programming languages*. MIT Press.

Xin Qi and Andrew C. Myers. 2009. Masked types for sound object initialization. In *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, Savannah, GA, USA, January 21-23, 2009*, Zhong Shao and Benjamin C. Pierce (Eds.). ACM, 53–65. https://doi.org/10.1145/1480881.1480890

Alban Reynaud, Gabriel Scherer, and Jeremy Yallop. 2018. A right-to-left type system for mutually-recursive value definitions. *CoRR* abs/1811.08134 (2018). arXiv:1811.08134 http://arxiv.org/abs/1811.08134

Marco Servetto, Julian Mackay, Alex Potanin, and James Noble. 2013. The Billion-Dollar Fix - Safe Modular Circular Initialisation with Placeholders and Placeholder Types. In *ECOOP 2013 - Object-Oriented Programming - 27th European Conference, Montpellier, France, July 1-5, 2013. Proceedings (Lecture Notes in Computer Science, Vol. 7920)*, Giuseppe Castagna (Ed.). Springer, 205–229. https://doi.org/10.1007/978-3-642-39038-8_9

Robert E. Strom and Shaula Yemini. 1986. Typestate: A Programming Language Concept for Enhancing Software Reliability. *IEEE Trans. Software Eng.* 12, 1 (1986), 157–171. https://doi.org/10.1109/TSE.1986.6312929

Alexander J. Summers and Peter Müller. 2011. Freedom before commitment: a lightweight type system for object initialisation. In *Proceedings of the 26th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2011, part of SPLASH 2011, Portland, OR, USA, October 22 - 27, 2011*, Cristina Videira Lopes and Kathleen Fisher (Eds.). ACM, 1013–1032. https://doi.org/10.1145/2048066.2048142

Yoav Zibin, David Cunningham, Igor Peshansky, and Vijay A. Saraswat. 2012. Object Initialization in X10. In *ECOOP 2012 - Object-Oriented Programming - 26th European Conference, Beijing, China, June 11-16, 2012. Proceedings (Lecture Notes in Computer Science, Vol. 7313)*, James Noble (Ed.). Springer, 207–231. https://doi.org/10.1007/978-3-642-31057-7_10