

Actor Capabilities for Message Ordering

Colin S. Gordon¹[0000–0002–9012–4490]

Drexel University, Philadelphia PA 19104, USA csgordon@drexel.edu

Abstract. Actor systems are a flexible model of concurrent and distributed programming, which are efficiently implementable, and avoid many classic concurrency bugs by construction. However they must still deal with the challenge of messages arriving in unexpected orders.

We describe an approach to restricting the orders in which actors send messages to each other, by equipping actor references — the handle used to address another actor — with a protocol restricting which message types can be sent to another actor and in which order using that particular actor reference. This endows the actor references with the properties of (flow-sensitive) static capabilities, which we call actor capabilities.

By sending other actors only restricted actor references, actors may control which messages are sent in which orders by other actors. Rules for duplicating (splitting) actor references ensure that these restrictions apply even in the presence of delegation. The capabilities themselves restrict message send ordering, which may form the foundation for stronger forms of reasoning. We demonstrate this by layering an effect system over the base type system, where the relationships enforced between the actor capabilities and the effects of an actor’s behaviour ensure that an actor’s behaviour is always prepared to handle any message that may arrive.

Keywords: Actor systems · Capability systems · Type-and-effect systems

1 Introduction

Actor systems have long been of interest since Hewitt’s original conception of the idea [40], both for anthropomorphic, intuitive appeal similar to that which boosted early interest in object-oriented programming (leading to work on active objects [58,5]) — and because the programming model enjoys many natural benefits that make it particularly well-suited to structuring concurrent and distributed programs. Compared to popular shared-memory concurrency models like POSIX threads, the actor model inverts control so that all actions in the system occur in response to asynchronous messages. This avoids classic issues with data races¹ and deadlocks. Instead actor systems suffer from difficulties ensuring that messages are sent or received in expected orders [4,38,67], including

¹ Though actor implementations embedded in imperative programming languages may need additional techniques [35,14,51,66] to share data that was once or may again be mutable.

the possibility that a message arrives before or after an actor is using a behaviour ready to process that particular message.

Static checking of message ordering in message passing systems is a classic problem, studied in many settings. Nielson, Nielson, and Amtoft [2,57] examined this problem for Concurrent ML (CML) in the 1990s, proposing what we now recognize as sequential [65] (a.k.a., flow-sensitive) type-and-effect systems for ensuring sends and receives in that synchronous message passing environment were in agreement between different threads. Work on session types began around the same time [44,43], initially for pairs of processes, and later for arbitrary numbers of processes. (And as it turns out, those can be formulated as effect systems as well [59].) Each of these has a language for describing *all* communication behaviour of a process, yielding possibly-complex specifications. (We discuss this further in Section 5.)

Effect systems are closely related to static *capability systems*, which typically trade off some expressivity and precision for dramatic increases in brevity of specifications [30]. Rather than allowing a program to do whatever it pleases and attempting to analyze that after the fact with effects, these systems require a specific capability for the program to perform any operation of interest. Dynamic capability systems do this via the dynamic semantics of the language: the operation literally cannot execute without the availability of a capability value [49,54]. Actor systems are already a form of dynamic capability system — an actor X can only send messages to another actor Y if it possesses the actor reference for Y, and provides it to the send operation. Static capability systems instead use static checking — usually a type(-and-effect) system — to restrict what a region of code can do by restricting which (static) capabilities are accessible to that code. In this way, the set of input capabilities can give an upper bound on the effect of a program fragment [15,30].

A less widely-acknowledged aspect of static capability systems is that by restricting the creation or duplication of capabilities, it becomes possible to enforce *global* invariants with *local* checks [30], which has been employed throughout decades of work on what are now collectively known as *reference capabilities*, where the capability restrictions are attached to each reference to an object. This encompasses many ownership (and universe) type systems [7,6,22], read-only reference systems [68,70,71,35,28], and rely-guarantee reasoning applied between references [33,53,52,34], to name just a few variants.

In this paper we explore this idea adapted to *actor references* rather than heap references, in a version of these ideas we call *actor capabilities*. Each actor reference carries a restriction, given as a formal language over message types, restricting the order in which messages may be sent using that particular reference; different references to the same actor may carry different capabilities/restrictions, allowing highly *localized* checks for message ordering. We also make an unusual choice compared to most capability systems: rather than assuming each actor is created with a specific fixed language, we allow each actor to dynamically create *new* capabilities for itself — tracked with an effect sys-

tem. The result is a type system which combines flow-sensitive capabilities with flow-sensitive (sequential [65]) effects in a highly-integrated way.

2 An Intuitive Explanation of Actor Capabilities

Current actor system implementations come in two varieties: those which allow actors to send any message type to any actor at any time, and those which attach a type parameter to actor references (handles used to address another actor) which restrict the types of messages that can be sent. The former category includes early versions of Akka and modern versions of Akka.NET, as well as all dynamically-typed actor languages and libraries (Erlang, Elixir, etc.). The latter category includes modern Akka, which we focus on because it is closest to what we propose. *All* actor systems in either category effectively handle messages by case analysis of the message’s dynamic type. Akka (now) uses the Scala type `ActorRef[T]` to represent a handle to send a particular actor messages of type `T`. This is both a dynamic capability and a static capability:

- It is impossible to send an actor a message without obtaining an `ActorRef[T]` referring to it. These actor references cannot be forged by a program. Thus an actor reference is effectively a form of dynamic object capability [54].²
- To send an actor a message of type `U`, code must use an `ActorRef[T]` where `U <: T` (`U` is a subtype of `T`). To send some other message of type `B`, the code would still need an `ActorRef[B]` — even if a developer knew (correctly!) that the recipient would *always* accept a `B`, the corresponding actor reference would be required by the type system. Multiple actor references to the same actor with different (incomparable!) message type bounds can exist. Thus this already has much in common with reference capability systems.

Akka also ensures that if an `ActorRef[T]` exists for a specific actor, then that actor will *always* accept messages of type `T`. With the capability properties above, this invariant ensures that all messages that can be sent will be handled. The code of Figure 1 offers a useful starting point to dig into this a bit more, to call out the specific aspects of this situation we must generalize to deal with the set of valid messages changing over time.

The `ExampleActor` is an object which when invoked (i.e., by `ExampleActor()` implicitly calling the `apply` method) creates a new behaviour, defined in the presence of an Akka-selected `ActorContext` for the newly-created actor; for our purposes, this is effectively a container for the new actor’s self reference `context.self` of type `ActorRef[MessageBase]` (though of course Akka uses it for much more). The new actor is created by `Behaviors.setup`. `beh` is a recursively defined `Behavior[MessageBase]`, indicating that this is a behavior which accepts messages of type `MessageBase` (not shown). Akka’s functional API (the one used in this example³) requires each behaviour to return a new behavior

² We are intentionally side-stepping discussions of capability-safety and ambient authority raised by the ability to look up an actor reference by the actor’s name.

³ Akka also has an object-oriented API, and both are heavily used, but the functional API is quite close to both classic actor calculi [1] and our formal development.

```

import akka.actor.typed.*
import akka.actor.typed.scaladsl.*
object ExampleActor {
  def apply(): Behavior[MessageBase] =
    Behaviors.setup((context: ActorContext[MessageBase]) => {
      val beh : Behavior[MessageBase] = // Recursively-defined behaviour
      Behaviors.receiveMessage((msg: MessageBase) => {
        msg match { // Inspect the message type
          case Heartbeat(sender) => // if it is a heartbeat...
            val returnAddress : ActorRef[Heartbeat] = context.self
            sender ! HeartbeatResponse(returnAddress) // reply
            beh // Return next behaviour from this case
          ... // other message cases omitted
        } })
      beh // Make beh the initial behaviour
    })
}

```

Fig. 1. A simple Scala program using the Typed Akka API

to handle the next message, which this code does in each branch of the `match` statement, returning itself (so after handling one message, the behavior explicitly signals that it will be used to handle the next message as well).

The one shown case of the behavior responds to the `Heartbeat` subclass of `MessageBase` by first extracting a new reference for the sender to directly message this actor in the future (the message may have been routed indirectly through other actors). The type of the “return address” reference is specialized to its intended use: `ActorRef[Heartbeat]`. `ActorRef[MessageBase]` is a subtype of `ActorRef[Heartbeat]` — `ActorRef` is contravariant in its type parameter. So via subtyping the `context.self` reference, which can be used to send any subtype of `MessageBase`, can be coerced to a more restrictive type, which can only be used to send further `Heartbeat` messages. This more restrictive reference is then sent back to the sender.

2.1 A Capability View of Safety Guarantees

Notice that we derive important safety guarantees from the type system: the original requestor can never use the `ActorRef` sent back with the response to send the `ExampleActor` a message type it does not handle, and the actor is statically known to handle all messages it receives.⁴ We obtain this from the (sometimes implicit) three-way agreement across:

- The `ActorRef[MessageBase]` (from `context.self`) specifies the expected type of messages sent to the actor. In particular, all other actor references to this actor are derived from `context.self` or the spawn operation (which yields an `ActorRef` of the same type), and by contravariance specify only *subtypes* of `ActorRef[MessageBase]` — for any `ActorRef[T]` referring to this actor, `T <: MessageBase`.

⁴ We are using handle in a loose sense here; the behaviour might simply ignore some subtypes of `MessageBase` (intentionally or due to a developer’s mistake).

- The `Behavior[MessageBase]` specifies the expected message (super)type for incoming messages, and `context.self`'s type parameter to `ActorRef` will always match the behaviour's type parameter.
- The send operation (`!` operator) at the recipient of the response (`sender`) only allows messages of type `Heartbeat` to be sent via an `ActorRef[Heartbeat]`

`ActorRef[T]` can be seen as a family of static capabilities. Owning one for a particular `T` *only* permits sending values of (subtypes of) type `T`, even if the destination actor can accept additional message types. The fact that each new capability to a certain actor is created with *at most* the permissions of the capability it was copied from means that creation of new capabilities / actor references collectively preserve a global invariant that no part of the program possesses a capability to send anything but a `MessageBase` in this example. This property of creating new `ActorRefs` from old is similar to a cross-cutting theme in much of the static capability literature, though only some work makes this fully explicit [10,33,34,53,52] (as we will need to do later). As all actor references are either returned from the actor creation (not shown) or derived from the actor's self-reference (`context.self`), both of which have an `ActorRef` parameter matching the `Behavior` parameter, this collectively maintains the global property [30] that no actor reference can be used to send a message that will not be handled by the installed behaviour. While each behaviour returns a possibly-different behaviour for handling future incoming messages, the type system ensures the new behaviour's message type matches the current behaviour, so previously-shared references to the actor never become invalid (i.e., never permit sending messages unexpected by some future behaviour of the same target actor).

This is a valuable property, and the reason that newer versions of Akka switched *to* this typed API, rather than using a general `ActorRef` type without a type parameter (called the “classic” API), which could be sent messages of any type. However, it does come at a cost. It is common for actors to move through multiple lifecycle phases, where some messages are (temporarily or permanently) irrelevant, yet a correct program must still explicitly include code to handle those unexpected messages. The older untyped Akka API allowed code to simply switch to a new behaviour which might not handle *any* of the message types it previously received! This was convenient for code that was correct, but nothing statically prevented bugs where one actor would send another a message of a type it either expected to never see again or to not see yet. Technically, the same can still happen with the typed API if `MessageBase` is actually `Object`, it is simply less pronounced because developers are encouraged to use smaller type hierarchies. It would be desirable to preserve the kinds of safety guarantees targeted by this implicitly capability-based API, while extending them to *fully* resolve the question of certain message types arriving at unexpected times.

2.2 Extending the Capability View for Message Ordering

This capability-based view of how Akka (and similar actor systems) ensure all messages are handled can be extended to deal with the case where actors' set of accepted messages change arbitrarily over time.

The key idea we will employ is to replace the single type parameter of Akka’s `ActorRef` with a formal language over types — thus specifying a set of allowable *sequences* of messages that may be sent to that actor. This allows the specification of actor references that permit sending just a single message, a fixed number of messages, sending a certain message type just once and others an unbounded number of times, alternating message types in a particular way, or other interesting restrictions — the notion of a formal language is quite general.

On top of this, we must preserve and extend the three-way agreement described above. When an actor reference is duplicated, we cannot (in general) allow the reference to be duplicated at the same type. Otherwise an `ActorRef` that permitted sending exactly one message of type `T` could first be copied, then used to send a `T`, then the copy could be used to send a `T`, circumventing the intended restriction of the original capability. So we must adapt machinery for carefully tracking splitting of capabilities [33,34,53,52,9,10] to ensure the number and relative ordering of message sends are preserved when actor references are duplicated — including handling the subtleties from the new copies being usable in any (possibly-interleaved) order. Fortunately there are existing tools in formal language theory we can use to guide the design here.

The other pieces of preserving and extending this three-way correspondence are the send operation, and agreement between actor references and behaviour typing. The send operation clearly must be typed in a flow-sensitive manner, to track when a send “uses up” one allowed message type. The more complex piece is maintaining agreement between these consumable actor reference capabilities and a more flexible behaviour typing that allows actors to start and stop accepting messages of various types. The challenge is to ensure that previously-shared capabilities continue to remain valid when an actor changes behaviours.

If we treat Akka’s existing `ActorRef[T]` as a capability to send a sequence of messages of types in $\{\tau \mid \tau <: T\}^*$, then ensuring every future handler for the actor must handle the same types is sufficient to ensure old capabilities remain valid, because the message types are fixed permanently (invariant), so ensuring that the initial actor references and initial behaviour agree on `T` at actor creation time is a sufficient check. But if both the permissions of remaining capabilities, and the future behaviours’ domains change over time, we will need additional machinery to synchronize those changes. We will use an effect system for this, where duplicating/splitting a capability for the current actor will create obligations for future message handling tracked in effects.

3 An Actor Capability Language

Figure 2 gives the syntax, and dynamic semantics of a simple actor language (abbreviated to omit entirely standard aspects). As in the most classic versions of semantics for the actor model [1], an actor is a reactive program, executing only in response to an incoming message. The actor’s message handling is defined by a *behaviour*, which we model as a partial function from types to handling expressions for each accepted incoming message type. In response to a message,

Expressions $e ::= n \mid b \mid \langle e, e \rangle \mid e \oplus e \mid \neg e \mid \text{if } e \text{ e } e \mid (\lambda_f x. e) \mid \rho \mid \text{beh}_L(\bar{\tau} \rightsquigarrow e) \mid \text{spawn}(e)$
 Operators $\oplus \in + \mid - \mid \times \mid \div \mid \vee \mid \wedge$
 Actor Stores $\sigma \in \text{Location} \rightarrow \text{beh}_L(\bar{\tau} \rightsquigarrow e)$
 Message Queues $\mu \in \text{list}(\text{Value} \times \text{Type} \times \text{Location})$
 Global Message Queues $\gamma \in \text{Location} \times \text{Location} \rightarrow \text{list}(\text{Value} \times \text{Type})$

$\boxed{\mu; \sigma; e \xrightarrow{L}_\ell \mu'; \sigma'; e'}$ Booleans, natural numbers, pairs, and (recursive) functions omitted

$$\begin{array}{c}
 \text{E-SELF} \frac{}{\mu; \sigma; \text{self}_L \xrightarrow{\{L\}}_\ell \mu; \sigma; \ell} \quad \text{E-SPAWN} \frac{\text{IsBehaviour}(v) \quad \text{fresh } \ell'}{\mu; \sigma; \text{spawn}(v) \xrightarrow{\{e\}}_\ell \mu; \sigma \uplus \ell \mapsto v; \ell'} \\
 \\
 \text{E-SEND} \frac{l \Downarrow_\ell \ell_{\text{dst}}; \mu; \sigma \quad e \Downarrow_\ell v; \mu'; \sigma'}{\mu; \sigma; \text{send}_\tau(\ell_{\text{dst}}, v) \xrightarrow{\{e\}}_\ell \mu + (v, \tau, \ell_{\text{dst}}); \sigma; ()} \\
 \\
 \text{E-MSGHANDLED} \frac{\sigma; \gamma \Rightarrow \sigma'; \gamma' \quad \tau \in \bar{\tau} \quad \emptyset; \emptyset; e_\tau[\text{msg} \mapsto v] \xrightarrow{L}_\ell^* \mu'; \sigma'; v'}{\sigma[\ell \mapsto \text{beh}(\bar{\tau} \rightsquigarrow e)]; \gamma[\ell, \text{sender} \mapsto (v, \tau) :: \bar{v}s] \Rightarrow \sigma[\ell \mapsto v'] \uplus \sigma'; \text{enqueue}(\ell, \gamma, \mu)}
 \end{array}$$

where $\text{enqueue}(\ell, \gamma, \mu)$ extends each $\gamma(\ell_{\text{destination}}, \ell)$ with $\mu(\ell)$ by appending, thus preserving message ordering between endpoints, without committing to delivery order across actors.

Fig. 2. An actor language, partial operational semantics

the message handler may create other actors and/or send messages to existing (or newly-allocated) actors. $\mu; \sigma; e \xrightarrow{L}_\ell \mu'; \sigma'; e'$ says that with initial outgoing message queue μ and locally-spawned actors σ , expression e running as actor ℓ steps to e' , resulting in message queue μ' (which always extends μ) and locally-spawned actors σ' (which always extends σ). The L indicates the cumulative self-capabilities generated by the reduction, which is the empty string except when reducing an annotated `self` expression. Iterated single stepping shuffles these formal languages together — an operation discussed in detail with the type system. The language L is purely an instrumentation and does not influence execution. Local actor reduction is deterministic in this calculus.

The full program steps by reducing an individual actor's reaction to any message ready for delivery: $\sigma; \gamma \Rightarrow \sigma'; \gamma'$ has one rule, delivering an available message to an actor *which is able to handle it* in the sense of running a matching handler to completion (i.e., termination). The handler body is expected to reduce to another behaviour, which may be the same behaviour or a new behaviour. (Note that this is also how many modern actor frameworks, such as Akka, work as well.) Message delivery is tracked per send-receive pair and is in-order between any pair of actors in a single direction, but incoming messages from different actors may be interleaved arbitrarily. Configurations where all actors either diverge⁵ for any of the next incoming messages or do not handle them are stuck. Configurations where at least one actor may continue executing are *not* stuck.

⁵ Non-terminating message handlers are always a bug in actor programs, as they prevent the actor from handling any messages in the future. Another static analysis could be used to ensure termination.

Failures In typical actor implementations, if an actor does not support a particular message type, either the message is dropped (the runtime simply doesn't deliver the message at all), explicitly ignored / passes through a generic handler which matches on message type, or an exception is delivered to the receiving actor. Regardless of typical handling in implementations, sending a message to an actor that does not expect that message type is typically an error we would like to statically prevent. We abstract the three possibilities above as a stuck transition in our formalization.

Actor State and Recursive Behaviours In our calculus, there is no mutable state, so changes to the state of an actor between messages must be encoded by returning a new behaviour which somehow captures the updated state. Because our functions are always recursively defined, we can for example encode a count of messages received as

$$(\lambda_f \text{state}. \text{beh}(\mathbf{0} \rightsquigarrow f(\text{state} + 1))) \mathbf{0} \quad (1)$$

That is, *state* binds the current count, and in response to receiving a unit message, the new behaviour is defined by recursively calling the function with the count incremented by one. The issues involved in direct mutable state are discussed in Section 5.

Note that 1 above also highlights the construction of a recursive behaviour *indirectly* through the use of a recursive function. Real actor framework implementations frequently include direct self-references in their behaviours (e.g., if a behaviour is modeled as an object, the standard `this` refers to the behaviour) or have ways to implicitly repeat a behaviour if a change in behaviour is not explicitly requested. Both of these can be encoded in our calculus.

Actor Self-Reference and Recursive Behaviours Actors may retrieve their own address using the `self` operator. Note that in our calculus, `self` is an *operator*, not a variable bound to the current actor's reference. While `self` could be modeled as a variable, our type system treats `self` somewhat differently from variables, for reasons explained in Section 4.2.

4 Actor Capabilities For Guaranteed Message Handling

The language in the previous section permits actors to send messages to other actors that they are not prepared to handle — messages of types that may not be dealt with by the target actor's behaviour at the time of message delivery. This section develops a type-and-effect system to ensure such errors are detected statically by the type checker. Conceptually, there are two aspects. The first aspect, which gives rise to the name of our technique, is to equip actor references with a capability structure that restricts how they are used. Section 4.1 explains that structure and how the use is enforced. Section 4.2 then explains how the system ensures consistency between an actor capability sent to another actor and the subsequent behaviours of a given actor.

Capabilities and effects are both formal languages over the set of types.

$$\begin{array}{l}
\text{Types } \tau ::= \mathbb{B} \mid \mathbb{N} \mid \mathbf{1} \mid \tau \times \tau \mid \tau \xrightarrow{\chi} \tau \mid \text{ActorRef}(L) \mid \text{Beh}(L) \\
\text{Effects } \chi \in \mathcal{L}(\tau) \\
\text{(Formal) Languages } L \in \mathcal{L}(\tau) \\
\text{Paths } p ::= x \mid p.1 \mid p.2
\end{array}$$

$$\boxed{\Gamma \vdash e : \tau \dashv \Gamma' \mid \chi} \quad \text{Rules for booleans, conditionals, numbers, and pairs omitted}$$

$$\begin{array}{c}
\text{T-APP} \quad \frac{\Gamma \vdash e_1 : \tau \xrightarrow{\chi} \tau' \dashv \Gamma' \mid \chi_1 \quad \Gamma' \vdash e_2 : \tau \dashv \Gamma'' \mid \chi_2}{\Gamma \vdash e_1 e_2 : \tau' \dashv \Gamma'' \mid \chi_1 \sqcup \chi_2 \sqcup \chi} \quad \text{T-LAM} \quad \frac{\Gamma \vdash \Gamma \prec \Gamma * \Gamma \quad \Gamma, f : \tau \xrightarrow{\chi} \tau', x : \tau \vdash e : \tau' \dashv \Gamma \mid \chi}{\Gamma \vdash (\lambda_f x. e) : \tau \xrightarrow{\chi} \tau' \dashv \Gamma \mid \{\epsilon\}} \\
\text{T-CONTRACT} \quad \frac{\Gamma \vdash \tau \prec \tau' * \tau'' \quad \Gamma, x : \tau', x : \tau'', \Gamma' \vdash e : \tau \dashv \Gamma' \mid \chi}{\Gamma, x : \tau, \Gamma' \vdash e : \tau \dashv \Gamma' \mid \chi} \quad \text{T-WEAKEN} \quad \frac{\Gamma, \Gamma' \vdash e : \tau \dashv \Gamma'' \mid \chi}{\Gamma, x : \tau, \Gamma' \vdash e : \tau \dashv \Gamma'' \mid \chi} \\
\text{T-VAR} \quad \frac{}{\Gamma, x : \tau \vdash x : \tau \dashv \Gamma \mid \{\epsilon\}} \quad \text{T-SELF} \quad \frac{}{\Gamma \vdash \text{self}_L : \text{ActorRef}(L) \dashv \Gamma \mid L} \\
\text{T-SEND} \quad \frac{\Gamma \vdash e : \tau \dashv \Gamma' \mid \chi \quad \Gamma' \vdash_{\text{path}} p : \text{ActorRef}(L) \rightsquigarrow \text{ActorRef}(\tau^{-1}L) \vdash \Gamma'' \quad \tau^{-1}L \neq \emptyset}{\Gamma \vdash \text{send}_\tau(p, e) : \mathbf{0} \dashv \Gamma'' \mid \chi} \\
\text{T-SPAWN} \quad \frac{\Gamma \vdash e : \text{Beh}(L) \dashv \Gamma' \mid \chi \quad L' \subseteq L}{\Gamma \vdash \text{spawn}(e) : \text{ActorRef}(L') \dashv \Gamma' \mid \chi} \quad \text{T-IF} \quad \frac{\Gamma \vdash e_c : \mathbf{B} \dashv \Gamma' \mid L_c \quad \forall b \in \{t, f\}. \Gamma' \vdash e_b : \tau \dashv \Gamma_b \mid L_b}{\Gamma \vdash \text{if } e_c \text{ } e_t \text{ } e_f : \tau \dashv \Gamma_t \cap \Gamma_f \mid L_c \sqcup (L_t \cup L_f)} \\
\text{T-BEH} \quad \frac{\forall i \in 1..n. \Gamma, \text{msg} : \tau_i \vdash e_i : \text{Beh}((\tau_i^{-1}L) \sqcup L'_i) \dashv \Gamma'_i \mid L'_i \quad \text{disjoint}(\bar{\tau})}{\Gamma \vdash \text{beh}_L(\bar{\tau} \rightsquigarrow \bar{e}^n) : \text{Beh}(L) \dashv \epsilon \mid \{\epsilon\}} \\
\boxed{\Gamma \vdash \tau \prec \tau * \tau} \quad \text{S-ACTORREF} \quad \frac{L' \sqcup L'' \subseteq L}{\Gamma \vdash \text{ActorRef}(L) \prec \text{ActorRef}(L') * \text{ActorRef}(L'')} \\
\text{S-SELF SPLIT} \quad \frac{\tau \in \{\mathbb{B}, \mathbb{N}, \mathbf{1}\} \vee \tau = \tau' \xrightarrow{\chi} \tau''}{\Gamma \vdash \tau \prec \tau * \tau} \quad \text{S-PAIR} \quad \frac{\Gamma \vdash \tau \prec \tau * \tau \quad \Gamma \vdash \sigma \prec \sigma * \sigma}{\Gamma \vdash \tau \times \sigma \prec \tau \times \sigma * \tau \times \sigma}
\end{array}$$

Fig. 3. A capability type-and-effect system for our actor language. A standard exchange rule is omitted. Behaviour types are *not* splittable.

4.1 Substructural Actor Capabilities and Splitting

Actor capabilities in our type system are equipped with a *formal language* L drawn from Type^* , and thus describe the sequence of message types that may be sent via that actor reference.⁶ Because the language describes both a restriction on the use of the value and a guarantee to the possessor that uses within that restriction should be valid, we call these actor references capabilities. An invariant we maintain is that the combined use of any and all actor capabilities (references) to some actor, in any interleaved order, should be allowed and handled by the

⁶ Technically the set of all types forms an infinite alphabet (one may have, at a minimum, n -ary tuples of unit for each $n \in \mathbb{N}$), but in the absence of polymorphic recursion, any given program only mentions a finite number of types (those used in behaviour case labels), and we assume Type denotes that set, avoiding subtleties of infinite alphabets.

actor they refer to — after all, if each capability is held by a different actor, they could potentially be used in any order! Such interleaving can be represented by the shuffle (\sqcup) operator. The shuffle of two words $x = x_1 \dots x_k$ and $y = y_1 \dots y_k$ (with x_i and y_i being possibly-empty subwords, not necessarily characters) is defined as

$$x_1 \dots x_k \sqcup y_1 \dots y_k = x_1 y_1 \dots x_k y_k$$

which is lifted to languages as

$$L_1 \sqcup L_2 = \{x \sqcup y \mid x \in L_1, y \in L_2\}$$

That is, the shuffle of two languages consists of all possible interleavings of words from each language.

This gives us a way to control duplication of capabilities, to share actor references without increasing the permissions available to some code. Consider the case of an actor reference with type $\text{ActorRef}(\text{nop}^* \cdot \text{act} \cdot \text{nop}^*)$ (momentarily assuming our languages are described by regular expressions, which is not a requirement of the formal development). This reference can be used to send as many nop messages as desired or necessary, but may only be used to send an act message once — exactly once. Simply duplicating that reference at the same type results in *two* actor references which could *each* be used to send an act message, and repeated duplication would then permit sending unlimited such messages. Instead, we borrow an idea from the reference capability literature, of *splitting* types with such restrictions [33,34,53,52,9,10]. We borrow Gordon’s notation [33,34] in Figure 3: $\Gamma \vdash \tau \prec \tau' * \tau''$ justifies splitting a variable of type τ into two copies, one of type τ' and the other of type τ'' . This is used by the contraction rule (T-CONTRACT). We call primitive types and closures *reflexively splittable* because they may be split without change (truly copied). Pairs split element-wise (S-PAIR). Actor references split using language shuffle (S-ACTORREF). Thus we could split our example actor reference into one of type $\text{ActorRef}(\text{act})$ and one of type $\text{ActorRef}(\text{nop}^*)$, because the shuffle of those two protocol languages is contained in the original:

$$\frac{\text{act} \sqcup \text{nop}^* \subseteq \text{nop}^* \cdot \text{act} \cdot \text{nop}^*}{\Gamma \vdash \text{ActorRef}(\text{nop}^* \cdot \text{act} \cdot \text{nop}^*) \prec \text{ActorRef}(\text{act}) * \text{ActorRef}(\text{nop}^*)} \text{S-ACTORREF}$$

Note that behaviours are *not* splittable — they are linear, and the absence of a rule for splitting behaviours is by design. This allows a behaviour to capture a non-reflexively-splittable variable and use it. *Recursive* behaviours are encoded via recursive lambda expressions which return behaviours, and lambdas may only capture reflexively-splittable values — so any substructural values used in such behaviours are freshly created each time the recursive function is invoked.

Consumption The point of use for an actor reference is via sending a message (T-SEND). This rule types the message to send, and uses an *updating* typing judgment on paths to both retrieve the type of the path prior to use (a variable, or “field chasing” from a variable bound to a tuple), and update that type (either

updating the type of the variable directly, or updating the type of appropriate — possibly nested — tuple components) to reflect sending a message of type τ . The judgment returns the updated type environment. This structure is chosen because our type system uses contraction deliberately, and may have multiple bindings for the same variable; separately typing the path and updating the output type environment would introduce difficulties ensuring the binding used for the pure typing was the same variable updated after the send, and it is simpler to do both at once.

$\tau^{-1}L$ is the derivative operator on formal languages, which includes all possible *completions* of words in L that start with the chosen element of the alphabet:

$$\alpha^{-1}L = \{w \mid \alpha \cdot w \in L\}$$

Of course the purpose of our actor reference type is two-fold: to enforce a usage discipline on code that holds an actor reference (addressed above), and to guarantee that uses according to that discipline are valid. The second purpose is partly addressed by splitting (which ensures an actor cannot increase the set of message sequences it can send to another actor) and partly addressed by effects (which allow relating the capabilities that exist for an actor to its intended behaviours).

Program Typing If actors only run in response to messages, how does the system begin? A program is given by the behaviour e for an initial actor, which must accept (at least) the unit type as a message. The initial configuration for a system has actor pool $\{\ell_0 \mapsto e\}$ and global message queue $\{\ell_0, \ell_0 \mapsto ((), \mathbf{0}) :: \mathbf{nil}\}$, from which point the only available reduction (which *is* available) is to deliver the unit message to the initial actor. So top-level programs are typed by

$$\text{T-PROGRAM} \frac{\emptyset \vdash e : \text{Beh}(\bar{\tau}) \quad \mathbf{0} \in \bar{\tau}}{\vdash e \text{ prog}}$$

4.2 Consistency of Actor Capabilities, Behaviours and Effects

To discuss consistency between behaviours and capabilities, we must discuss the typing of behaviours. The type $\text{Beh}(L)$ characterizes a behaviour which will accept message sequences in L , though this does not fully characterize the future behaviour of the actor using this behaviour. L is intuitively an over-approximation of how all *existing* capabilities to send to that behaviour could collectively be used — all capabilities that existed *prior* to that behaviour handling a message. Note that this does *not* require knowing the exact history of the program: it is limited to the actor using the behaviour, and refers only to the *unused* capabilities, not past usage of capabilities. This information is necessary for typing behaviours because they represent *unresolved promises* for future message acceptance. At the same time, an actor with such a behaviour installed may end up accepting more than L in the future, because when the behaviour handles a message, it may create additional obligations for its actor to handle in the future, which we track using a static effect system.

The Effect of Capability Creation The starting point for all effects is T-SELF, which captures the creation of a new self-reference capability that is usable according to some particular language. Other rules simply accumulate an over-approximation of created capabilities. Subexpressions that run together in order (such as a function, argument, and latent function effect in an application expression per T-APP) have their effects shuffled together with \sqcup , as collectively all created capabilities could be used in an interleaved fashion. Expressions that may run alternatively are combined with union — the rule for conditionals combines the conditional effect L_c with the true- and false-branch effects L_t and L_f as $L_c \sqcup (L_t \cup L_f)$, capturing that the condition always runs before either branch, but it is statically unknown which branch will execute.

Note that \sqcup distributes over \cup on both sides, so with $\{\epsilon\}$ as the unit for \sqcup , these effects form a total commutative effect quantale [31]. It follows both from that framework’s soundness results (and from our later proofs) that the effect of an expression over-approximates the shuffle-interleaving of all self-capabilities created during a local expression’s reduction.

Note that the system tracks the *creation* of new capabilities for the current actor, *not* which capabilities are *actually sent* to other actors! This is both simpler than distinguishing between sending a self-reference or another actor’s reference, and also reflects the fact that an actor can send itself messages, in which case those would need to be handled by later behaviours.⁷

Typing Behaviours T-BEH in Figure 3 is subtle, but captures the relationship between past promises, locally-created capabilities, and future message acceptance. Unsurprisingly, the rule requires typing each branch of the behaviour’s body, with the expectation that the result is a new behaviour (which will be used to handle future messages). Those messages the actor handles in the future may come from one of two sources, each reflected in an argument to the shuffle (\sqcup) operator. The left argument $\tau_i^{-1}L$ models the fact that each branch of the behaviour handles a different incoming message type (which must have been justified by a previously-existing capability accounted for by L), but any remaining capabilities that existed before might still exist and be usable to send further messages. So the left argument captures that the obligation to accept an initial τ has been discharged. The right argument L_i corresponds to new capabilities to the current actor which this message-handler may have created, which are then *additional* obligations for the future which did not previously exist, so are not reflected in the original L . L_i is guaranteed to include all new obligations because it is also the *effect* of the message handler. So requiring the behaviour returned after handling τ_i to have type $\text{Beh}((\tau_i^{-1}L) \sqcup L'_i)$ ensures that both previously-existing obligations and newly-created obligations will be supported by the next handler.

⁷ It is of course also conservative: an actor may create a self-reference that it then discards, but the discarded capability will still be captured in the effect and become a new (un-dischargeable) obligation for the actor.

For example, in the **nop/act** example of the previous section, if the actor has only given out that one capability discussed previously, L would be $\text{nop}^* \cdot \text{act} \cdot \text{nop}^*$. In typing the handler for **act**, the derivative of that history assumption by **act** would be nop^* , reflecting that the one promised handling of **act** was settled, and if that was the only such capability distributed, then no future behaviour would necessarily be forced by that old capability (which was now partly consumed) to handle additional **act** messages. The derivative handles this. However, if in response to the **act** message the handler created a new capability for the actor which granted permission to send another **act** message, then the future behaviour would need to reflect that, hence the requirement that the next behaviour also handle the languages of all newly-created capabilities — incorporated via \sqcup .

Spawning New Actors, and Initial Capabilities The one remaining matter to attend to is creating an actor with a specific behaviour — the rules discussed so far ensure that *if the history assumptions are accurate they remain accurate*, and T-SPAWN establishes the initial accuracy. T-SPAWN requires that the assumed protocol L' for the initial external capability is a subset of the language L that the spawned behaviour assumes exists prior to its execution.

4.3 Soundness

Due to space constraints we omit details of the type soundness proof, and give only a brief overview here; see the associated technical report [32] for more detail. We use the syntactic approach to type safety. Soundness of effect tracking follows typical form for syntactic proofs of sequential (flow-sensitive) effect systems [31], but in this case integrated with tracking of a flow-sensitive runtime actor capability map mirroring Γ 's usage, to track each actor's unused capabilities to send messages to other actors. There is both a local (per-actor) and global version of each of preservation and progress.

The essence of ensuring that outstanding capabilities match the expectations of their respective target actors is that locally, the shuffle of all capabilities an actor possess to another actor ℓ before handling a message ($\Sigma(\ell)$) is related to the remaining capabilities after handling it ($\Sigma'(\ell)$) by a derivative of messages sent. If the handler sent messages \bar{v} of types $\bar{\tau}$ to ℓ , then $\Sigma'(\ell) \sqcup \text{caps}(\ell, (\bar{v}, \bar{\tau})) = \bar{\tau}^{-1} \Sigma(\ell)$, where $\text{caps}(\ell, (\bar{v}, \bar{\tau}))$ is the shuffle of all capabilities as passed through outgoing messages (though the formal proof deals with them as an additional level of existentials, not as something computed from the values and message types). The right side of that equality reflects what is left of the original capabilities after using some to send values, and the left side reflects what is remaining to type the reduced expression, plus the capabilities that were transferred out via messages. Thus if a capability with protocol $\tau_1 \cdot \tau_2 \cdot \tau_3$ were used once to send a τ_1 and the remainder (now a capability for $\tau_2 \cdot \tau_3$) were then sent in an outgoing message, the send of τ_1 would be reflected in the derivative on the right, and the message-bound remainder would be accounted for in the **caps** (if the leftover capability were retained rather than transferred, it would instead be accounted for in $\Sigma'(\ell)$).

5 Related Work

Our system draws on many areas of prior work, including substructural type systems and reference capabilities, but those were discussed earlier. We focus this section not on an exhaustive explanation of the many sources of inspiration for this work (largely covered by related work discussions of the previously-cited papers), but instead on comparing our work against other approaches to ensuring that actors (or active objects) and processes handle specific message sequences and send compatible message sequences to others.

A classic point of reference for this is Nierstrasz’s work [58] on *regular types*, assigning finite-state-based specifications of object behaviour, as the sequence of messages accepted and the resultant responses. His exposition explores the structure of these specifications and informally how they might apply to objects, but provides no application of these types to a concrete programming language, and focuses on composing applications by way of one object using another as its sole client. This focus on the behaviour of the active object itself over time is also nearly dual to ours: regular types for active objects describe the process behaviour summatively, while in our (source-level) type system, no summative language-oriented description of an actor’s accepted messages exists! Instead static knowledge of an actor’s accepted language can only be assembled from various capabilities. Active objects, as nicely captured in de Koster et al.’s taxonomy of actor language variants [19], do not permit changing the set of accepted messages over time. Support for such changes is an explicit goal of this work.

Our calculus is pure, but many popular actor frameworks are in languages with mutable state, and permit sharing references across actors on the same machine [11,36]. The introduction of direct mutable heaps raises classic issues of balancing safety and efficiency of sending messages. Deep copying of state is always sound, but highly inefficient, especially when the data being sent is either immutable or the sender discards its own access to the mutable data after sending. This has motivated a great deal of work on various flavors of uniqueness and mutation control, including a great deal of work applying the reference capability notions discussed in the introduction to track forms of unique ownership transfer or object immutability [35,34,9,10,14,60,28]. Some of this work [35,14] has specifically been applied to actor systems.⁸ However, incorporating two orthogonal flavors of capabilities — the actor capabilities proposed here, along with the reference capabilities from existing work — would distract from the main technical development of capabilities for controlling message ordering in actor systems.

An alternative approach to managing state sharing and exchange between actors is the introduction of specially-shared capsules. De Koster et al. [20,18] introduce *domains*, which are shared-memory constructs which actors asynchronously request access to (akin to asynchronous mutex acquisition). Similarly, Cheeseman et al.’s *couns* [12] (named for *concurrent ownership*) are shared resources

⁸ Now that the relevant NDAs are long-expired, we can finally say in writing that the context for [35] was the Midori [23] derivative of Singularity [24,45], where processes were essentially actors.

accessed by asynchronously-spawned computations that run only when all required cowns are available at once. Some of the most closely related work to ours also falls into this group: Caldwell et al. [8] describe a specialized data exchange environment called a *dataspace*, as well as an effect system describing behaviour over time for actors and data — effectively dual to our approach of restricting actor behaviours with capabilities. Similar to Skalka et al. [64,63], their effects give abstract summaries of actor behaviour, which can then be used to model-check properties of the program — the effect types for both Skalka et al.’s work and Caldwell et al.’s work give a model that is verified to soundly abstract the program code. Establishing a formal relationship between our capabilities and behavioural effects like Caldwell et al.’s, akin to Craig et al.’s formal relationship between flow-insensitive capabilities and flow-insensitive effects [15] would be intriguing future work.

Session types [44,48], choreographic programming [62,16,27], and multitier programming [69] solve the issue of agreement between the messages accepted by a given process and message sent to it by others in the inverse way of what we describe here. Instead of analyzing correctness of processes in isolation, these approaches start with a single program describing all global behaviour, and a compiler generates code (or in the case of session types, synchronization skeletons) for individual processes via a projection process. These approaches are generally applied to *synchronous* message passing, though Harvey et al. [37] and Neykova et al. [56] have adapted session types to actors (though still using a global description of the full orchestration which is then projected for each of a fixed set of roles). By contrast, our lack of a global type is a trade-off. On one hand, there is no complete closed description of behaviour separate from the code itself, which can be useful for understanding the system as a whole. On the other hand, actors in our approach can be specified and checked in isolation (e.g., as library components), without needing to know a priori how they fit into a larger system, because any messages accepted by the actor beyond those described in the initial capability returned from `spawn` are chosen by the actor itself. Like Caldwell et al.’s work [8], the type of a process in these approaches describes *all* of the process’s future behaviour, rather than only the promises it has made to its environment about its behaviour (as in our work).

Originating around the same time as classic binary session types [43], Nielson and Nielson [57] described one of the earliest flow-sensitive effect systems, for ensuring Concurrent ML (CML) processes sent and received *synchronous* messages in an expected order. This early similarity turned out to be quite relevant, as session types can be recast as an effect system (and vice versa) [59]. Thus the idea of applying effect systems — or an equivalent formalism — to ensure correct message request ordering is itself far from new. And the idea of using capabilities to restrict effects is also not new, going back past recent formal connections [15] to many earlier expositions of effects [39,50] in an implicit form.

Two aspects of our approach distinguish it from this established body of work. First, this is to the best of our knowledge the first use of flow sensitive *capabilities* (a form of coeffect [13]) to achieve safe ordering. While the relation-

ship to a putative effect system achieving the same ends has not been formally pursued, and this likely carries some of the downside of capabilities (vs. effects) outlined in our prior work [30] (which applied only to flow-insensitive effect systems), it also carries a unique advantage in system design not discussed in that paper. While effect systems may obtain greater precision in reasoning about side effects [30], they do so at the cost of needing to “reconstruct” behaviour and check compatibility after the fact. In contrast, our capability formulation of the problem inverts this checking, and places each actor (and its original allocation site) in firm control of how other actors interact with it — the protocols of the self-capabilities an actor sends to others directly constrain those others’ interactions a priori, localizing checks to a far greater extent than in CML or session types — no direct global accounting of behaviour is required to exist.

Second — and a complementary factor in not requiring a global behavioural description to exist — is our use of an effect system to track capability *creation*. This differs from most static capability work, which typically assumes all possible capabilities for something exist upon its creation, and must be passed through the program. Instead we treat actors’ retrieval of a new self-reference as a *dynamic* capability creation — via effects. While it has long been recognized that both coeffects (e.g., capabilities) and effects are required to reason about various combinations of independently-interesting static analyses and language semantics [26,17], ours is one of very few examples of mixing coeffects and effects *as part of a single static analysis*.

Another body of work with a more nuanced relationship to the work reported here is the concept of *mailbox types*, as originally proposed by de’Liguoro and Padovani [21] and later implemented for an actor language with explicit mailboxes by Fowler et al. [25]. The core idea there is to drastically restrict aliasing of mailbox references, and use the near-linearity to reason about the contents of the *current actor’s* mailbox, using commutative regular expressions. Beyond their implementation of their system, and experience applying it to some benchmark programs, there are three key differences between this idea and actor capabilities. First is that the behaviour tracking is almost inverted: rather than restricting senders to only send what the recipient expects, the recipient is instead restricted to only *request* things from the mailbox which are known to have been sent to it. Mailbox types are predicated on the *selective receive* operation present in some actor frameworks, which retrieves a message of specified type. Thus mailbox types prevent an actor from issuing a selective receive which would fail. Second is that only message multiplicity is tracked, because of the commutative regular expressions; the use of selective receive makes this the appropriate modeling choice, though most actor frameworks include both normal and selective receive, making integration of our ideas and theirs interesting future work. Third, they use a combination of second-class values and a mild relaxation of linearity to deal with aliasing of mailboxes and strong updates to mailbox types. By contrast, we permit many aliases to each actor whose usage restrictions we track, by ensuring that each alias is independently usable (via

\sqcup); again, their choice is a reasonable one for their setting, as each mailbox is expected to be consumed by only one actor.

Program logics have also been used for reasoning about correctness of actor messages. Gordon [29] proposed modal assertions of the form $@_l(P)$ inspired by hybrid logic [3], indicating that P was true of the state of the actor at l . Thus assertions true at the sending actor could be packaged in this way, and communicated as part of message invariants; a rely-guarantee-style relation constraining state evolution worked together with a restriction to only send assertions stable with respect to that relation ensured the assertions about other actors were never invalidated, but limited what could be proved to assertions over monotonically-evolving portions of state (though CRDTs have shown this can be made more flexible than it initially appears [61]). More recently some program logics (instantiations of IRIS [46]) have tied together ideas from session types and verification for simple actor languages [41,42]. This work handles protocols similar to binary (for Actris) or multiparty (for Multris) session types, but each message has not only a value but a set of assertions that the recipient may assume upon receipt of the corresponding value. This is much more powerful and flexible than Gordon’s approach, at the cost of complexity. Gordon’s system was prototyped as a library for the automated verification tool Dafny, while the modern approaches rely on IRIS embedded in an interactive proof assistant. The key point of comparison to the present work is that Actris and Multris do ensure the absence of unhandled messages, but the specifications for doing so are essentially traditional session types. Compatibility when delegating use of a channel to another process is checked via an ad hoc simulation search procedure akin to Militão et al.’s [52]. In contrast, our use of shuffle uncovers a clean algebraic structure underlying common uses of sharing and resharing the addresses of other processes, though it is less general than Actris and Multris, where the correct continuation of a protocol may depend on the specific value sent.

Acknowledgments. Thanks to Gul Agha for his influential work on the actor model, which forms the basis for much of my research career (via the work I am most known for [35] and some of the work I am most fond of [29]); as well as some of what I most enjoy teaching (using actors when teaching distributed computing). Agha’s work [47,1] has been the source for a significant amount of my own learning to think about concurrent programs and their implementation, and by chance he has had a direct influence on individual people who have influenced me during time in industry and academia [55]. Thanks also to Ivan Beschastnikh, for asking me in 2012 how I could apply reference capabilities to distributed systems. It took 13 years, but I think I’ve finally articulated a good answer, even if handling faults remains future work. Audience members at IWACO 2024 provided useful feedback on a more abstract version of the ideas presented here. Thanks to Jonathan Schuster for bringing mailbox types to my attention, and to Lindsey Kuper for bringing my preprint to Jonathan’s attention. This work was supported by NSF Award #CCF-2007582.

References

1. Agha, G.: *Actors: a model of concurrent computation in distributed systems*. MIT press (1986)
2. Amtoft, T., Nielson, F., Nielson, H.R.: *Type and Effect Systems: Behaviours for Concurrency*. Imperial College Press, London, UK (1999)
3. Areces, C., ten Cate, B.: Hybrid logics. In: *Studies in Logic and Practical Reasoning*, vol. 3, pp. 821–868. Elsevier (2007)
4. Bagherzadeh, M., Fireman, N., Shawesh, A., Khatchadourian, R.: Actor concurrency bugs: a comprehensive study on symptoms, root causes, api usages, and differences. *Proceedings of the ACM on Programming Languages* **4**(OOPSLA), 1–32 (2020)
5. Boer, F.D., Serbanescu, V., Hähnle, R., Henrio, L., Rochas, J., Din, C.C., Johnsen, E.B., Sirjani, M., Khamespanah, E., Fernandez-Reyes, K., et al.: A survey of active object languages. *ACM Computing Surveys (CSUR)* **50**(5), 1–39 (2017)
6. Boyapati, C., Lee, R., Rinard, M.: Ownership Types for Safe Programming: Preventing Data Races and Deadlocks. In: *OOPSLA* (2002)
7. Boyapati, C., Rinard, M.: A Parameterized Type System for Race-Free Java Programs. In: *OOPSLA* (2001)
8. Caldwell, S., Garnock-Jones, T., Felleisen, M.: Programming and reasoning about actors that share state. *Journal of Functional Programming* **34** (2024). <https://doi.org/10.1017/S0956796824000091>
9. Castegren, E., Wrigstad, T.: Reference capabilities for concurrency control. In: *30th European Conference on Object-Oriented Programming, ECOOP 2016* (2016)
10. Castegren, E., Wrigstad, T.: Relaxed linear references for lock-free programming. In: *31st European Conference on Object-Oriented Programming, ECOOP 2017* (2017)
11. Charousset, D., Hiesgen, R., Schmidt, T.C.: Caf-the c++ actor framework for scalable and resource-efficient applications. In: *Proceedings of the 4th International Workshop on Programming based on Actors Agents & Decentralized Control*. pp. 15–28 (2014)
12. Cheeseman, L., Parkinson, M.J., Clebsch, S., Kogias, M., Drossopoulou, S., Chisnall, D., Wrigstad, T., Liétar, P.: When concurrency matters: Behaviour-oriented concurrency. *Proceedings of the ACM on Programming Languages* **7**(OOPSLA2), 1531–1560 (2023)
13. Choudhury, V., Krishnaswami, N.: Recovering purity with comonads and capabilities. *Proceedings of the ACM on Programming Languages* **4**(ICFP), 1–28 (2020)
14. Clebsch, S., Drossopoulou, S., Blessing, S., McNeil, A.: Deny capabilities for safe, fast actors. In: *Proceedings of the 5th International Workshop on Programming Based on Actors, Agents, and Decentralized Control*. pp. 1–12. ACM (2015)
15. Craig, A., Potanin, A., Groves, L., Aldrich, J.: Capabilities: Effects for Free. In: *International Conference on Formal Engineering Methods (ICFEM)* (2018)
16. Cruz-Filipe, L., Montesi, F., Peressotti, M.: A formal theory of choreographic programming. *Journal of Automated Reasoning* **67**(2), 21 (2023)
17. Dal Lago, U., Gavazzo, F.: A relational theory of effects and coeffects. *Proceedings of the ACM on Programming Languages* **6**(POPL), 1–28 (2022)
18. De Koster, J., Marr, S., D’Hondt, T., Van Cutsem, T.: Domains: Safe sharing among actors. *Science of Computer Programming* **98**, 140–158 (2015)
19. De Koster, J., Van Cutsem, T., De Meuter, W.: 43 years of actors: a taxonomy of actor models and their key properties. In: *Proceedings of the 6th International*

- Workshop on Programming Based on Actors, Agents, and Decentralized Control. pp. 31–40 (2016)
20. De Koster, J., Van Cutsem, T., D’Hondt, T.: Domains: Safe sharing among actors. In: Proceedings of the 2nd edition on Programming systems, languages and applications based on actors, agents, and decentralized control abstractions. pp. 11–22 (2012)
 21. de’Liguoro, U., Padovani, L.: Mailbox types for unordered interactions. In: 32nd European Conference on Object-Oriented Programming (ECOOP 2018). pp. 15–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2018)
 22. Dietl, W., Drossopoulou, S., Müller, P.: Generic Universe Types. In: ECOOP (2007)
 23. Duffy, J.: (2015), <https://joeduffyblog.com/2015/11/03/blogging-about-midori/>
 24. Fähndrich, M., Aiken, M., Hawblitzel, C., Hodson, O., Hunt, G., Larus, J.R., Levi, S.: Language support for fast and reliable message-based communication in Singularity os. In: Proceedings of the 1st ACM SIGOPS/EuroSys European Conference on Computer Systems 2006. pp. 177–190. EuroSys ’06, ACM (2006). <https://doi.org/10.1145/1217935.1217953>
 25. Fowler, S., Attard, D.P., Sowul, F., Gay, S.J., Trinder, P.: Special delivery: programming with mailbox types. Proceedings of the ACM on Programming Languages **7**(ICFP), 78–107 (2023)
 26. Gaboardi, M., Katsumata, S.y., Orchard, D., Breuvar, F., Uustalu, T.: Combining effects and coeffects via grading. In: Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming. pp. 476–489 (2016)
 27. Giallorenzo, S., Montesi, F., Peressotti, M.: Choral: Object-oriented choreographic programming. ACM Transactions on Programming Languages and Systems **46**(1), 1–59 (2024)
 28. Giannini, P., Servetto, M., Zucca, E., Cone, J.: Flexible recovery of uniqueness and immutability. Theoretical Computer Science **764**, 145–172 (2019)
 29. Gordon, C.S.: Modal assertions for actor correctness. In: Proceedings of the 9th ACM SIGPLAN International Workshop on Programming Based on Actors, Agents, and Decentralized Control. pp. 11–20 (2019)
 30. Gordon, C.S.: Designing with static capabilities and effects: Use, mention, and invariants (pearl). In: 34th European Conference on Object-Oriented Programming (ECOOP 2020). Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2020)
 31. Gordon, C.S.: Polymorphic iterable sequential effect systems. ACM Transactions on Programming Languages and Systems (TOPLAS) **43**(1), 1–79 (2021)
 32. Gordon, C.S.: Actor capabilities for message ordering (extended version). Tech. rep. (2025), <https://arxiv.org/abs/2502.07958>
 33. Gordon, C.S., Ernst, M.D., Grossman, D.: Rely-guarantee references for refinement types over aliased mutable data. In: Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation. pp. 73–84 (2013)
 34. Gordon, C.S., Ernst, M.D., Grossman, D., Parkinson, M.J.: Verifying invariants of lock-free data structures with rely-guarantee and refinement types. ACM Transactions on Programming Languages and Systems (TOPLAS) **39**(3), 1–54 (2017)
 35. Gordon, C.S., Parkinson, M.J., Parsons, J., Bromfield, A., Duffy, J.: Uniqueness and reference immutability for safe parallelism. In: Proceedings of the ACM international conference on Object oriented programming systems languages and applications. pp. 21–40 (2012)
 36. Haller, P.: On the integration of the actor model in mainstream technologies: the Scala perspective. In: Proceedings of the 2nd edition on Programming systems, languages and applications based on actors, agents, and decentralized control abstractions, pp. 1–6 (2012)

37. Harvey, P., Fowler, S., Dardha, O., Gay, S.J.: Multiparty session types for safe run-time adaptation in an actor language. In: 35th European Conference on Object-Oriented Programming (ECOOP 2021). Schloss Dagstuhl-Leibniz-Zentrum für Informatik (2021)
38. Hedden, B., Zhao, X.: A comprehensive study on bugs in actor systems. In: Proceedings of the 47th International Conference on Parallel Processing. pp. 1–9 (2018)
39. Henglein, F., Makhholm, H., Niss, H.: Effect types and region-based memory management. In: Pierce, B.C. (ed.) *Advanced Topics in Types and Programming Languages*, chap. 3, pp. 87–136. MIT Press (2005)
40. Hewitt, C., Bishop, P., Steiger, R.: A universal modular actor formalism for artificial intelligence. In: *Proceedings of the 3rd international joint conference on Artificial intelligence*. pp. 235–245 (1973)
41. Hinrichsen, J.K., Bengtson, J., Krebbers, R.: Actris: Session-type based reasoning in separation logic. *Proceedings of the ACM on Programming Languages* **4**(POPL), 1–30 (2019)
42. Hinrichsen, J.K., Jacobs, J., Krebbers, R.: Multiris: Functional verification of multiparty message passing in separation logic. *Proceedings of the ACM on Programming Languages* **8**(OOPSLA2), 1446–1474 (2024)
43. Honda, K., Vasconcelos, V.T., Kubo, M.: Language primitives and type discipline for structured communication-based programming. In: *European Symposium on Programming*. pp. 122–138. Springer (1998)
44. Honda, K., Yoshida, N., Carbone, M.: Multiparty Asynchronous Session Types. In: *Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. POPL '08 (2008). <https://doi.org/10.1145/1328438.1328472>
45. Hunt, G.C., Larus, J.R.: Singularity: Rethinking the software stack. *SIGOPS Oper. Syst. Rev.* **41**(2), 37–49 (Apr 2007). <https://doi.org/10.1145/1243418.1243424>
46. Jung, R., Krebbers, R., Jourdan, J.H., Bizjak, A., Birkedal, L., Dreyer, D.: Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *Journal of Functional Programming* **28**, e20 (2018)
47. Karmani, R.K., Shali, A., Agha, G.: Actor frameworks for the jvm platform: a comparative analysis. In: *Proceedings of the 7th International Conference on Principles and Practice of Programming in Java*. pp. 11–20 (2009)
48. Lagaillardie, N., Neykova, R., Yoshida, N.: Implementing multiparty session types in rust. In: *International Conference on Coordination Languages and Models*. pp. 127–136. Springer (2020)
49. Levy, H.M.: *Capability-based computer systems*. Digital Press (1984)
50. Marino, D., Millstein, T.: A generic type-and-effect system. In: *TLDI* (2009)
51. Milano, M., Turcotti, J., Myers, A.C.: A flexible type system for fearless concurrency. In: *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation*. pp. 458–473 (2022)
52. Militão, F., Aldrich, J., Caires, L.: Composing interfering abstract protocols. In: 30th European Conference on Object-Oriented Programming, ECOOP 2016 (2016). <https://doi.org/10.4230/LIPIcs.ECOOP.2016.16>
53. Militao, F., Aldrich, J., Caires, L.: Rely-guarantee protocols. In: *ECOOP 2014–Object-Oriented Programming: 28th European Conference, Uppsala, Sweden, July 28–August 1, 2014*. Proceedings 28. pp. 334–359. Springer (2014)
54. Miller, M.: *Robust composition: Towards a unified approach to access control and concurrency control*. Johns Hopkins University (2006)

55. Miller, M.S., Tribble, E.D., Shapiro, J.: Concurrency among strangers: Programming in e as plan coordination. In: Trustworthy Global Computing: International Symposium, TGC 2005, Edinburgh, UK, April 7-9, 2005. Revised Selected Papers. pp. 195–229. Springer (2005)
56. Neykova, R., Yoshida, N.: Multiparty session actors. *Logical Methods in Computer Science* **13** (2017)
57. Nielson, F., Nielson, H.R.: From CML to process algebras. In: CONCUR (1993)
58. Nierstrasz, O.: Regular types for active objects. In: Proceedings of the Eighth Annual Conference on Object-Oriented Programming Systems, Languages, and Applications. p. 1–15. OOPSLA '93, Association for Computing Machinery, New York, NY, USA (1993). <https://doi.org/10.1145/165854.167976>, <https://doi.org/10.1145/165854.167976>
59. Orchard, D., Yoshida, N.: Effects as sessions, sessions as effects. In: Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. pp. 568–581 (2016)
60. Servetto, M., Pearce, D.J., Groves, L., Potanin, A.: Balloon types for safe parallelisation over arbitrary object graphs. In: Workshop on Determinism and Correctness in Parallel Programming (WoDet) (2013)
61. Shapiro, M., Preguiça, N., Baquero, C., Zawirski, M.: Conflict-free replicated data types. In: Stabilization, Safety, and Security of Distributed Systems: 13th International Symposium, SSS 2011, Grenoble, France, October 10-12, 2011. Proceedings 13. pp. 386–400. Springer (2011)
62. Shen, G., Kashiwa, S., Kuper, L.: Haschor: Functional choreographic programming for all (functional pearl). *Proceedings of the ACM on Programming Languages* **7**(ICFP), 541–565 (2023)
63. Skalka, C.: Types and trace effects for object orientation. *Higher-Order and Symbolic Computation* **21**(3) (2008)
64. Skalka, C., Smith, S., Van Horn, D.: Types and trace effects of higher order programs. *Journal of Functional Programming* **18**(2) (2008)
65. Tate, R.: The sequential semantics of producer effect systems. In: POPL (2013)
66. Team, S.L.: Data Race Safety, <https://www.swift.org/migration/documentation/swift-6-concurrency-migration-guide/dataracesafety/>
67. Torres Lopez, C., Marr, S., Gonzalez Boix, E., Mössenböck, H.: A study of concurrency bugs and advanced development support for actor-based programs, pp. 155–185. Springer (2018)
68. Tschantz, M.S., Ernst, M.D.: Javari: Adding Reference Immutability to Java. In: OOPSLA (2005). <https://doi.org/10.1145/1094811.1094828>
69. Weisenburger, P., Wirth, J., Salvaneschi, G.: A survey of multitier programming. *ACM Computing Surveys (CSUR)* **53**(4), 1–35 (2020)
70. Zibin, Y., Potanin, A., Ali, M., Artzi, S., Kiezun, A., Ernst, M.D.: Object and Reference Immutability Using Java Generics. In: ESEC-FSE (2007). <https://doi.org/10.1145/1287624.1287637>
71. Zibin, Y., Potanin, A., Li, P., Ali, M., Ernst, M.D.: Ownership and Immutability in Generic Java. In: OOPSLA (2010). <https://doi.org/10.1145/1869459.1869509>