

Typed Dataspace Actors

SAM CALDWELL and TONY GARNOCK-JONES and MATTHIAS FELLEISEN
Northeastern University, Boston, Mass.

(*e-mail*: samc@ccs.neu.edu)

Abstract

Actors collaborate via message exchanges to reach a common goal. Experience has shown, however, that pure message-based communication is limiting and forces developers to use design patterns.

The recently introduced dataspace actor model borrows ideas from the tuple space realm. It offers a tightly controlled, shared storage facility for groups of actors. In this model, actors assert facts that they wish to share and interests in such assertions. The dataspace notifies interested parties of changes to the set of assertions that they are interested in.

Although it is straightforward to add the dataspace model to untyped languages, adding a typed interface is both necessary and challenging. Without restrictions on exchanged data, a faulty actor may propagate erroneous data through a malformed assertion, causing an otherwise well-behaved actor to crash—violating the key principle of failure-isolation. A properly designed type system can prevent this scenario and rule out other kinds of uncooperative actors. This paper presents the first structural type system for the dataspace model of actors; it does not address the question of behavioral types for assertion-oriented protocols.

1 Coordinating Conversing Actors

Languages in the tuple space family (Gelernter, 1985; Carriero *et al.*, 1994; Murphy *et al.*, 2006; Mostinckx *et al.*, 2007) address the need for manageable sharing between concurrent components (Tasharofi *et al.*, 2013).¹ Generally speaking, these languages connect concurrent processes to a storage medium, through which they communicate indirectly. Rather than sending a message along a channel or to a mailbox, a process deposits it in the central repository, from which another process may later retrieve it. This indirection provides decoupling from identity (e.g. address or channel) and across time, allowing a process to receive a message “sent” before its own creation.

Over the past few years, Garnock-Jones *et al.* (2014; 2016; 2017) have developed the dataspace model, a promising addition to the actor family based on tuple spaces. Roughly speaking, their dataspace actors utilize a storage medium for sharing knowledge among a group of participants (see section 2 for a recap). A dataspace actor shares facts by asserting them in this common space. These assertions remain associated with the actor. A dataspace actor interested in facts states its interest in assertions about those facts. It is the dataspace’s responsibility to inform an actor of the appearance and disappearance of facts in which it

¹ The Blackboard idea (Newell & Simon, 1972; Englemore & Morgan, 1988) is a similar approach to the problem in an object-oriented setting.

has stated interest. More precisely, the dataspace sends actors notifications of changes to the state of facts, who in response may add or withdraw assertions or spawn other actors. Finally, the dataspace withdraws an actor's assertions by force when the latter raises an exception.

We have implemented the dataspace model as extensions of ECMAScript (ECMA, 2015) and Racket (Flatt & PLT, 2010), and we have created several applications in these extended languages. This experience shows that dataspace are particularly well-suited for situations where actors participate in several conversations at once and where many actors participate in the same conversation:

- Assertions lend themselves towards group communication. New participants can seamlessly join an ongoing conversation. Dataspace routing keeps each actor up-to-date with a minimum of boilerplate.
- Dataspace decouple actors from each other, allowing multicast communication without the burden of maintaining a handle (address or channel) for each peer.
- The link between actor and assertion provides a uniform and powerful mechanism for resource management in the presence of partial failure, a notoriously difficult aspect of concurrent programming.

These features eliminate many design patterns necessitated by conventional coordination mechanisms, say, in Erlang/OTP (Armstrong, 2003). The need for design patterns in traditional models becomes especially prevalent when code must deal with potentially faulty actors. See section 2.2 for more details.

Unfortunately, implementations of the tuple space model, including dataspace, have largely been untyped or untyped—embedded in a typed host language in principle, but circumventing the type system by assigning the same type to all messages and requiring casts. This situation deprives programmers of the potential for type-based error prevention, documentation, optimization, design, and so on. The decoupling of components and general open-endedness of communication poses a challenge for type systems, which must predict the shape of incoming data. Our experience has also drawn our attention to fault isolation and scalability, two essential principles of actor programming, that suffer from the absence of a type discipline:

- Localizing faults to the originating component is an important tool for developing applications that detect and recover from exceptional situations. Typically, however, only some errors receive this treatment, leaving others open to transmission between actors via exchanged messages. Sadly, this can be as basic as a type mismatch in a message, such as a faulty actor that sends a two-element tuple where the recipient expects a two-element list. Only when the receiver tries to use that tuple as a list is the mistake detected, but now the exception occurs inside an actor that works as intended.
- Actor programming calls for concurrency—spawning an individual actor for each independent task. This philosophy implies certain architectural decisions for language implementations. In particular, the potential for a great number of simultaneously live actors means that they must share the underlying resources for computation, such as CPUs and threads. And where there is sharing, there is potential for abuse. Actors that take an overly-long time to respond to messages can deprive peers of

their own share of time. In the worst case, the event handler of a running actor never terminates, permanently hampering the system and potentially rendering it unresponsive. Languages with type-based termination guarantees are a first step to face this challenge.

However, providing a typed interface to dataspace proves a challenge. Type systems, by rule, control the exchange of data between components. By necessity, type system design is intimately connected to the underlying semantics of data exchange, such as function call or method dispatch. When a new communication medium is introduced, such as a dataspace, it must be reconciled with the type system. Tuple space-style communication is particularly challenging, because the decoupling of components and indirect communications make the flow of data less apparent.

In response, we develop a structural type system that accounts for communication based on dataspace (see section 3). Our approach should be generalizable to other languages in the tuple space family. With the type system in hand, programmers may enjoy the benefits of both dataspace and a type discipline. The type system for the base actor language satisfies both a type soundness property and a termination guarantee (see section 4). Using these theorems, we can finally verify the (type) soundness of a complete dataspace model (see section 5). The properties of the model have a number of practical implications for programming, as well as limitations (see section 6). A prototype implementation of our type system confirms its basic usefulness, both in terms of convenience of use and of programming support (see section 7).

2 Programming in the Dataspace Model

The dataspace model of actors is parameterized over the base language, that is, the language for programming each individual actor. For an introduction using illustrative examples it is still best to use some concrete syntax. Here we use Racket syntax, falling back on one of the two untyped prototype implementations. This section provides an overview of the basic concepts (section 2.1), illustrates it with examples (sections 2.2 and 2.3), and makes the case for introducing a structural type system (section 2.4).

2.1 Racket Dataspaces

In the dataspace model, an individual actor combines a private state with a behavior function. The private state is any piece of data from the underlying language. In a chat server program, for example, a room-list may have the shape

```
(list "FBI" "CIA")
```

recording the name of each active chat room. Intuitively, an actor behaves as if it were a function that maps events and the actor's current state value to a new state value and some instructions to carry out on its behalf:

$$\text{Event} \times \text{State} \rightarrow \text{State} \times \text{Instructions}$$

Instructions are interpreted by the surrounding dataspace, which connects actors with the goal of enabling conversation and managing their state.

A dataspace equips a group of actors with a means of aggregating and sharing items of information, called *assertions*. In our chat server dataspace, we create one actor for each connection to an end user. This *user-agent* actor is responsible for relaying chat messages and carrying out commands from the remote user. For example, a connected user under the alias "Mo" may issue a request to join the room named "FBI". In response, the user's agent actor states the following assertion:

```
(in-room "Mo" "FBI")
```

In the syntax of our Racket prototype, this assertion is stated as a Racket `struct` form. Racket structures are fixed-length tuples declared by the programmer, where the name (`in-room`) serves as an identifying tag. When the user leaves the room, the user-agent withdraws this assertion. Assertions range over basic values—numbers, strings, and so on—as well as first-order, immutable data structures, including lists and program-specific structures. Assertions have the same status in dataspace programs as facts in a Prolog database (Clocksin & Mellish, 1981).

The dataspace model links an actor with its assertions in a particular way. An assertion is *read-only*, and only the actor making an assertion may remove it. While multiple actors may make identical assertions, dataspace hide such redundancy by providing a set view of an underlying bag of assertions. Additionally, the lifetime of an actor bounds the lifetime of its assertions. When an actor terminates, the dataspace removes its current assertions.

In order to observe the appearance and disappearance of assertions, an actor makes an *assertion of interest*. An assertion expressing interest in the above chat-room assertion uses the built-in structure `observe` (figure 2 summarizes the extensions to base Racket):

```
(observe (in-room "Mo" "FBI"))
```

This particular assertion of interest could originate from an actor managing the room list. The dataspace notifies this actor—and any other actor expressing interest—when the user-agent asserts `(in-room "Mo" "FBI")`, and again when it withdraws the assertion.

Actors often assert interest in entire families of related assertions. For example, the room-list actor tracks the presence of every user in the "FBI" room, not just "Mo". To express interest in *all* such values, the actor uses the wildcard `*` to assert

```
(observe (in-room * "FBI"))
```

In response, the actor receives a notification each time a matching `in-room` assertion appears or disappears. Semantically, a wildcard creates infinite sets of assertions; routing of notifications can thus be understood in terms of set intersection.

Figure 1 provides a visual intuition for the chat server dataspace described thus far. The "Mo" user has recently joined the "FBI" room. Accordingly, the user-agent actor places an `in-room` assertion in the dataspace. Because the room-list actor's assertions include an expressed interest in `in-room` assertions, the dataspace routing mechanism duly recognizes the overlap and sends an event.

An expression of interest is itself an assertion, observable by other actors. A common use of this recursion concerns *demand matching*. An actor capable of matching demand for `c` asserts `(observe (observe c))`, while actors interested in `c` assert `(observe c)`. The dataspace routing mechanism duly informs the first actor of the matching interest. In

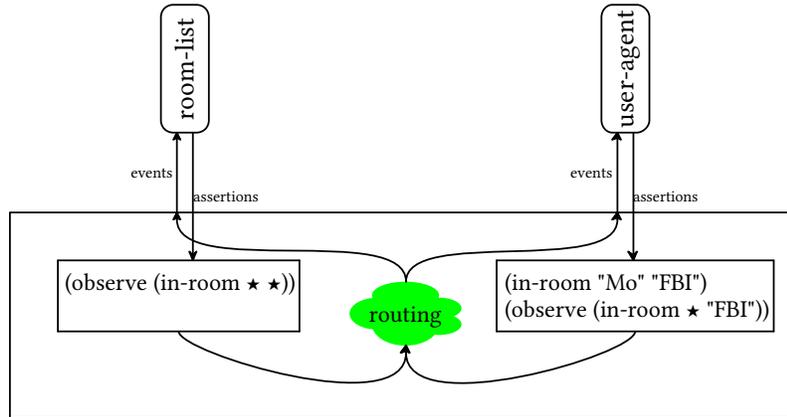


Fig. 1. Dataspace With Actors Engaged In a Conversation

response, it performs the computation to produce the assertion c . Finally, dataspace routing again kicks in to notify the interested actor in the appearance of c .

Actors receive notification *Events* regarding changes to their dataspace's assertions in the form of a *patch*, which contains two sets. The first is a set of assertions newly *added* to the dataspace matching the actor's expressed assertions of interest. The second is a set of relevant assertions that have been *removed* from the dataspace. Hence, the room-list actor receives the notification

$$(\text{patch } \{(\text{in-room } \text{"Mo"} \text{"FBI"})\} \emptyset)$$

when user "Mo" joins. When the user leaves the room the actor receives the corresponding notification $(\text{patch } \emptyset \{(\text{in-room } \text{"Mo"} \text{"FBI"})\})$.²

Expressions between curly braces $\{e \dots\}$ create sets of assertions while actors analyze incoming sets using *project*, which resembles set comprehension. Projection takes the form $(\text{project } [\text{pattern set}] \text{ body})$, as in

$$(\text{project } [(\text{in-room } \$\text{name } \text{"FBI"}) \{(\text{in-room } \text{"Mo"} \text{"FBI"})\}] \text{ name})$$

The pattern serves to filter and destructure assertions in *set*. Evaluation instantiates the body of the *project* form once for each assertion in the set matching the given pattern, yielding a list. Here, it yields $(\text{list } \text{"Mo"})$. In addition to binding variables, prefixed with $\$$, patterns may include the discard symbol $_$ which matches any value. Since assertion sets may be conceptually infinite, some *project* expressions may yield infinite lists, as in $(\text{project } [\$x \{*\}] x)$. The implementation signals an error when this happens.³ The discard pattern $_$ is especially useful for analyzing such sets; the expression $(\text{project } [_ \{*\}] \text{"match"})$ yields the value $(\text{list } \text{"match"})$.

² Dataspaces additionally support *message broadcast*, though we focus here on communication through assertions. Messages may be thought of as a special case of assertion-based communication: sending a message atomically combines the assertion and immediate retraction of its value, i.e., $(\text{message } c)$ behaves roughly as $(\text{patch } c \emptyset)$ atomically followed by $(\text{patch } \emptyset c)$.

³ An alternative interpretation is that of a diverging computation, instantiating x an infinite number of times. Since these infinite sets are recognizable, we opt for signaling an error.

6 *S. Caldwell and T. Garnock-Jones and M. Felleisen*

<code>(dataspace <i>expr</i>)</code>	Launch a dataspace program with some initial actors
<code>(actor <i>expr expr expr</i>)</code>	An actor with a behavior plus initial state and assertions
<code>{ <i>expr ...</i> }</code>	Create a set of assertions
<code>(observe <i>expr</i>)</code>	An assertion of interest
<code>*</code>	Describes all possible assertions
<code>(patch <i>expr expr</i>)</code>	Create a patch of two assertion sets
<code>(transition <i>expr expr</i>)</code>	Update private state plus some instructions to perform
<code>(quit <i>expr</i>)</code>	Termination with some final instructions
<code>(project [<i>expr pattern</i>] <i>expr</i>)</code>	Iterate over the matching assertions in a set
<code>\$<i>name</i></code>	Pattern that matches anything and binds it to <i>name</i>
<code>-</code>	Pattern that matches anything

Fig. 2. Extensions to Racket for Dataspace Programming

Actor behavior functions actually return values belonging to a disjoint sum, as opposed to the simplistic product described above. The sum allows an actor to shutdown in an orderly, non-exceptional manner. When an actor has fulfilled its purpose, it submits the first form of response, a quit record, to the dataspace:

```
(quit instructions)
```

In response, the dataspace removes the actor, withdrawing each of its assertions, and carries out some final *instructions* on its behalf. In the course of processing a notification, an actor may raise an exception. Uncaught exceptions are translated to `(quit empty)` by the dataspace sub-language.

The second form of return value makes or withdraws assertions via `patch-es` or augments the program with an entire new actor. In either case, the actor submits

```
(transition state instructions)
```

Such a `transition` record provides a new, updated private state for the actor and a list of *instructions* for the dataspace to carry out, in order. An instruction might be a simple patch. For example, the user represented by the "Mo" actor may issue a request to switch to the "CIA" chat room. To this end, the "Mo" actor would update both its `in-room` assertion, thus notifying other actors of the change, and its assertion of interest, allowing it to monitor other users present in the new room. It accomplishes this goal with this transition:

```
(transition
  "CIA"
  (list (patch {(in-room "Mo" "CIA") (observe (in-room * "CIA"))}
              {(in-room "Mo" "FBI") (observe (in-room * "FBI"))})))
```

The first item updates the actor's private state value to "CIA", recording the user's current room. The transition provides a single instruction, a patch manipulating the actor's assertions. The first set in the patch places the assertion `(in-room "Mo" "CIA")` in the dataspace as well as an assertion of interest in other users in the "CIA" room. The second set in the patch withdraws the corresponding assertion of presence and interest for the previous "FBI" room.

An actor might also request the creation of a new actor. A corresponding instruction is an actor specification, which takes the form

```
(actor behavior state assertions)
```

providing a description of the actor's behavior function as well as its initial private state and assertions.

Finally, the form

(dataspace *actors*)

launches a dataspace program with the given initial actors. Figure 2 summarizes the extensions to Racket for dataspace programming.

2.2 Putting the Pieces Together

Figure 3 shows a sketch of the user-agent actor implementation for a chat room, plus some of the context. The `main` function launches a dataspace with the described user-agent, room-list, and any other potential actors. The actor's initial state establishes the presence of user "Mo" in room "FBI", as described above. The behavior of the actor analyzes patches from the dataspace, projecting out the first slot from `in-room` assertions. The actor creates a `user-notification` for each arriving and departing peer. Another actor, not depicted, transforms such notifications into suitable network messages.

Highlighting Dataspaces The chat example illustrates several benefits of dataspace coordination over pure message-passing. Garnock-Jones (2017) surveys the comparative advantage of dataspaces versus other concurrency paradigms in greater depth, using different implementations of a similar chat example.

- The *temporal duration of assertions*, i.e. the fact that an assertion remains in the dataspace until the originating actor withdraws it or terminates, allows new actors to seamlessly join an ongoing conversation. When a new end user connects to the chat server, the program spawns a new user-agent actor on its behalf. This new actor expresses an interest in other members present in the current room. The interest effectively acts as a query of existing matching assertions. In response, the dataspace provides a description of all members present in the room.
- After joining a room and learning of existing members, the *same* assertion of interest provides the user-agent with ongoing *incremental knowledge updates*. The actor receives notification each time a current user leaves or a new user joins the room.
- *Group communication* simplifies the protocol and implementation of presence communication. The user-agent actor makes a single assertion that informs every peer to whom it is relevant. Assertions decouple actor identities. They allow communication without knowing exactly which actor(s), or even how many, to address. In traditional point-to-point models, a designated actor must maintain a collection of addresses or channels to enable communication among the other actors. As Joe Armstrong, designer of Erlang points out (Armstrong, 1994) (chapter 8.7), "If we want to send a message to a process, then we need to know its PID. This is often inconvenient since the PID has to be sent to all processes in the system that want to communicate with this process." In the same book (chapter 11), a similar multi-room chat example requires an additional actor for each possible room to keep track of this information.
- The *link between failure and communication* provides a uniform and convenient method for resource management in the presence of faulty actors. In this example,

Untyped

```
(define (main)
  (dataspace
   (list
    (create-user-agent "Mo" "FBI")
    (create-room-list ...)
    ...)))

(define (create-user-agent name initial-room)
  (actor
   ;; behavior: notify the user as peers enter and
   ;; leave the current room.
   (lambda (event current-room)
    (define arrivals
      (announce (patch-added event) current-room " arrived."))
    (define departures
      (announce (patch-removed event) current-room " departed."))
    (transition current-room
                 (append arrivals departures)))
   initial-room
   {(in-room name initial-room)
    (observe (in-room * initial-room))}))

;; Create user-notification patches based on
;; presence assertions in an event.
(define (announce assertions current-room message)
  (project [(in-room $who current-room) assertions]
           (patch {(user-notification (string-append who message))} {})))
```

Fig. 3. A user-agent actor

we may consider the ability to communicate with a particular chat user as a resource. As the availability of this resource is announced via assertion, it is automatically tied to the lifetime of the corresponding actor. An exception in the actor results in the withdrawal of the assertion, notifying interested peers. From the perspective of other user-agents, both orderly and exceptional termination convey the same information—that a certain peer is no longer communicable. Other actors in the system, meanwhile, can react in different ways, perhaps to clean up any associated network connection(s).

2.3 The Role of Recursive Interest in Dataspace Protocols

Dataspace protocols benefit from the use of recursive subscriptions, that is, assertions of interest in assertions of interest (and so on).

To illustrate the point, consider an “arithmetic service” dataspace. Assertions take the form (sum x y z) where $z = x + y$. Clients of the service make requests by asserting interest in sum assertions with particular numbers for x and y , such as

```
(observe (sum 4 5 *))
```

The actor implementing the service listens for requests with

```
(observe (observe (sum * * *)))
```

in this case responding (sum 4 5 9). The arithmetic service answers requests by analyzing the x and y values of sum requests and ignoring z :

```
(project [(observe (sum $x $y _)) e] (sum x y (+ x y)))
```

The example illustrates how interpreting the assertion of interest as a request yields the simplest possible protocol. Abstractly, this protocol for request/response situations combines the request assertion with the interest in the result. It thus reduces both the number of assertions an actor must state to make a request and the different types of assertions comprising the conversation. To appreciate the protocol's simplicity, consider the next-best alternative, which would be to use two distinct structs. One struct would represent requests separately via a (sum-request x y) assertion. The other would describe the results of requests as a (sum x y z) assertion. Assertions of interest in results must include the inputs x and y , so that there is no confusion with answers to other requests. This requirement is equivalent to (observe (sum x y *)), which is the request from above and suggests the simplification to a single struct.

2.4 The Problems of Programming Actors in Untyped Languages

Not surprisingly, using an untyped functional language as the implementation substrate of dataspace actors leads to the serious problems mentioned in the introduction.

(1) Productive communication between actors requires an agreement on the kinds of data to be exchanged; programming in such a manner is analogous to following a type discipline. Untyped languages like our Racket prototype provide little help with finding, and no help adhering to, a type structure that governs the communication between components. Consequently, the potential for mismatched messages undermines the principle of failure isolation. A buggy actor may make utterances that violate these expectations, but it is the actors that interpret the faulty messages that suffer the consequences. Worse, some of these mistakes correspond to simple type errors. A bug in a user-agent actor may turn a numeric user input into an assertion about its name, as a number rather than a string, causing other actors to crash while merrily continuing its own existence.

(2) The recursive nature of (interest in) assertions complicates the matter further. In the arithmetic dataspace from section 2.3, a logic programmer may scan the description of sum assertions and surmises that the protocol lends itself to computing differences as well as sums, as with (observe (sum 4 * 9)). Such a request is incompatible with the implementation of the arithmetic service described above; it violates the assumption that the y slot is a number, even though it matches the service's assertion of interest (observe (observe (sum * * *))). As discussed in section 2.1, project errors when a pattern variable has an infinite number of matches, so delivering such an assertion to the service actor causes it to crash—another failure to isolate misbehaving actors.

(3) The principle of scalability requires a certain amount of cooperation between actors in terms of execution time. Indeed, Garnock-Jones et al. (2016) show several key properties of the dataspace model under the assumption that the behavior function of every actor is total (including exceptions). However, programming actors in an untyped functional language such as Racket or ECMAScript allows the creation of both total and divergent behavior functions, with no way to distinguish the two. Moreover, divergence is a form

of failure, but it bypasses the reasoning mechanism provided by the language for these situations—automatic withdrawal of assertions. The termination guarantees of a simple type system aid in isolating and responding to failures as expected in the actor model.

(4) Additionally, developers may wish to place specific restrictions on the behavior of certain actors. For instance, in the case that some rooms are private, the developer of the chat server may wish to ensure that user-agent actors do not make overly broad queries. Permitting such queries would allow a curious user to discover the existence of, and join, private channels without an invitation. Types are one mechanism for imposing such constraints. If the type of user-agent actors indicates that no such subscriptions are made, the protocol is safe from intrusions.

3 Types for the Dataspace Model

Designing a type system for dataspace requires consideration of the semantics of both communication and computation, with special attention to the intersection of the two concerns. Our design meets this high-level criterion as follows:

- The language of types accounts for the flexible nature of dataspace communication with “true” (set) unions (Pierce, 1991) for describing sets of assertions. Dataspace are about commingling actors: each actor partakes in, and each assertion potentially pertains to, several conversations. Union types mirror the overlapping conversations, making them suitable for describing the (sets of) assertions in a dataspace.
- Communication and computation coincide in the functions used to express actor behavior. Dataspace event dispatch (communication) becomes function application (computation). The return values of behavior functions give rise to routing events. These functions warrant additional checking on inputs and outputs, accounting for their dual purpose.
- Additionally, behavior functions must terminate on all inputs. But, writing actors in a language with recursion permits diverging programs. To address this issue, we employ the standard method of basing the type system on the simply-typed λ -calculus plus an induction schema per recursive type (Jearing, 1995), which is known to make for a terminating language.
- Finally, the computational constructs for creating and accessing the sets of assertions used for communication introduce new possibilities for computational errors. Hence, the type system prohibits creating assertion sets with higher-order data or selecting an infinite branch of an assertion set with `project`.

The rest of this section is an informal introduction to this somewhat unusual combination of ideas.

3.1 Typed Dataspaces By Example

Unions. Union types are the basic building block of typed dataspace actors. In dataspace, groups of actors participate in conversations where each utterance takes the form of an

assertion. Describing these conversations with types means grouping together related assertions. Unions capture the multi-party, overlapping nature of dataspace conversations.⁴

For instance, in the chat room example from section 2, actors communicate presence information through `in-room` assertions and interest in those assertions. Each `in-room` assertion has type `(InRoom String String)`. A value of type `(InRoom String String)` is an assertion `(in-room v1 v2)` where both v_1 and v_2 are strings. By convention, type constructors such as `InRoom` are the camel-cased name of the described structure.

An assertion of interest in other participants, such as `(observe (in-room * "FBI"))`, has type `(Observe (InRoom * String))`. The type uses the parameterization of the `InRoom` constructor to allow wildcard interest in usernames but limit interest to named rooms. As a type, `*` stands for all possible assertions, including concrete strings, numbers, or interests, as well as any infinite set of assertions arising from `{...}` (set creation) expressions containing `*`.

The union type `PresenceAssertions` describes the “Room Presence” conversation, including both the assertion made by a user-agent to signal its presence in a particular room as well as the assertion of interest used to monitor its peers:

```
(define-type PresenceAssertions
  (U (InRoom String String)
     (Observe (InRoom * String))))
```

Together, they describe the possible utterances in the conversation between user-agent actors concerning chat-room presence.

Grouping together the types of each assertion pertaining to a conversation in a union provides useful documentation and aids in separate, modular programming of actors. While a single actor may express only a subset of the assertions, the flexible nature of union types allows the composition of overlapping conversations. For example, the room-list actor monitors `in-room` assertions and publishes a list of results. The `InRoom` assertions, the room-list actor’s interest in them,⁵ the published `RoomList`, and the interest used by other actors to learn the list forms an “Available Rooms” conversation:

```
(define-type RoomAssertions
  (U (InRoom String String)
     (Observe (InRoom * *))
     (RoomList (List String))
     (Observe (RoomList *))))
```

The “Room Presence” and “Available Rooms” conversations overlap, with each `in-room` assertion playing a role in both. In one conversation an assertion signifies the existence of a particular user, while in the other it signifies the existence of a particular room.

As figure 3 shows, a user-agent actor performs additional communication to notify the connected user of specific events. This conversation comprises of `user-notification`

⁴ Such types do not account for temporal or substructural constraints on exchanges; see section 6.

⁵ The type of interest in presence employed by the room list actor, `(Observe (InRoom * *))`, reflects the potential for wildcard interest in all possible rooms, unlike that of the user agent.

12 *S. Caldwell and T. Garnock-Jones and M. Felleisen*

structures, described by the type

```
(define-type NotificationAssertions
  (U (UserNotification String)
      (Observe (UserNotification *))))
```

The full implementation includes other conversations for network connectivity, sending chat messages, and changing rooms, each with a similar type: `NetworkAssertions`, `ChatAssertions`, `RoomAssertions`, and so on.

By describing the conversations in isolation, each actor can be implemented in terms of only those conversations in which it participates. In the case of the user-agent described here, we may elide `NetworkAssertions`, among others:

```
(define-type UserAgentAssertions
  (U PresenceAssertions
      RoomAssertions
      NotificationAssertions))
```

```
(define-type ChatDataspace
  (U PresenceAssertions
      NotificationAssertions
      NetworkAssertions
      ChatAssertions
      RoomAssertions))
```

Fig. 4. Chat Server Communication Type

Dataspaces. Taking the union of all the conversations, as shown in figure 4, yields a type that describes the entire dataspace. Supplying that type to a typed dataspace constructor,

```
(dataspace ChatDataspace
  (actor ...
    ...))
```

demands that all assertions have type `ChatDataspace`. The annotation is referred to as the *communication type* of the dataspace. The communication type is an agreement among the actors, both limiting each individual's actions as well as enabling typed reasoning about the behavior of peers. In the chat dataspace, the type permits the user-agent actor's presence assertion (`in-room "Mo" "FBI"`), of type `(InRoom String String)`, and interest (`observe (in-room * "FBI")`), with type `(Observe (InRoom * String))`, because of their inclusion in the `ChatDataspace` type via `PresenceAssertions`. By contrast, the `ChatDataspace` type prohibits actors that assert (`in-room "Marvin" 42`), because the second slot is not a string, or overly-broad queries such as (`observe *`), because `*` is not a subtype of any of the constructors that appear under `observe` in `ChatDataspace`.

While the communication type restricts the assertions an actor may make, it also enables reasoning about the shape of assertions that may match an expressed interest. The key assurance of the `ChatDataspace` type is that any assertion matching an expressed interest in `in-room` assertions is an `in-room` struct where both slots contain strings. That is, in any set of assertions a sent to an actor, the set corresponding to the names of connected users,

$\{v_1 \mid (\text{in-room } v_1 \ v_2) \in a\}$, is a finite set of strings, and similarly for the set of values in the second slot, the names of inhabited rooms. Consequently, the behavior functions of actors in the dataspace may use `project` to analyze the names of individual users and rooms without triggering an error.

Actors and Simple Behavior Functions. Every individual actor operates within a dataspace, making and withdrawing assertions, processing assertions, and spawning further actors, each of which also resides in the dataspace. In typed programs, the dataspace contains only assertions belonging to a specific type—the communication type τ_c . Hence, τ_c plays a central role in checking individual actor specifications.

The type for actors must then reflect the type τ_c of the dataspace in which it is going to run. A developer can use this type to rationalize the initial assertions and, most importantly, the code for the behavior function. Specifically, the developer must ensure that the behavior function (1) produces only actions—assertions as well as spawned actors—that the dataspace’s type permits; (2) is prepared to deal with any event to which the dataspace may apply it; and (3) spawns only actors that recursively obey these constraints, too.

A translation of these insights calls for equipping an actor term

$$(\text{actor } \textit{behavior} \ \textit{state} \ \textit{assertions})$$

with an interface type. This step resembles the addition of a parameter type to lambda terms during the design of a simple type system for languages based on the lambda calculus. Since an actor communicates within a dataspace of some communication type, the best way to signal this assumption is with an annotation τ_c that requests this match:

$$(\text{actor } \tau_c \ \textit{behavior} \ \textit{state} \ \textit{assertions}).$$

In contrast to function application, actor application is *implicit*. An `actor` term is used by submitting it to the dataspace as an action. In dispatching events, the dataspace applies the actor’s *behavior* function and interprets the resulting actions. These semantics—dispatching events and interpreting actions—are not represented in the surface syntax of the program, meaning it is not possible to check the use of an actor independently from its use in a dataspace. We can express this insight by assigning the type

$$(\text{Actor } \tau_c)$$

to an actor term. By implication, it becomes straightforward to check a dataspace term. All initial actors must have the type $(\text{Actor } \tau_c)$ if τ_c is the type of the dataspace.

Validating that an actor has a given type proceeds according to the three steps above:

1. The assertions of an actor come from two sources: either as part of its initial assertions or as an action produced by its behavior function. For the former, we check that the type of the initial assertion set is included by the assumed communication type τ_c . For the latter, recall the informal signature of behavior functions from section 2.1:

$$\textit{Event} \times \textit{State} \rightarrow \textit{State} \times \textit{Instructions}$$

For the moment, let us simplify things even further by considering the behavior function as taking in assertions of some type τ_{in} and outputting assertions of another

14 *S. Caldwell and T. Garnock-Jones and M. Felleisen*type, τ_{out} :

$$\tau_{in} \rightarrow \tau_{out}$$

Conceptually, τ_{in} and τ_{out} are the essence of the *Event* and *Instructions* types, respectively. Validating the assertions stated by the actor then entails checking that τ_c includes each type of assertion from τ_{out} .

2. Dataspaces compute and route events according to expressed interests. Hence, the assertions constructed with `Observe` in τ_{out} define a bound on possible events. Type checking uses this bound together with the rest of the assertions in the dataspace, represented by τ_c , to predict the types of events produced by routing when the program runs. To make this prediction, the type checker takes the intersection between assertions of interest in τ_{out} with τ_c . The result is a type describing all potential events the behavior function may be applied to, which must be a subtype of τ_{in} .

Concretely, in the case of the user-agent actor τ_{out} is

```
(U (Observe (InRoom * String))
  (InRoom String String)
  (UserNotification String))
```

According to the type, the actor might assert `(observe (in-room * "FBI"))`. It would then receive a notification containing all `in-room` assertions with the string "FBI" in the dataspace. Inspecting the `ChatDataspace` type of figure 4, the type of potentially overlapping assertions is `(InRoom String String)`. Consequently, the user-agent actor's behavior function input type τ_{in} must accommodate such assertions.

3. Finally, an actor may spawn other actors into its dataspace. If every one of these actors has type `(Actor τ_c)`, they all obey the communication discipline of the surrounding dataspace.

Manipulating Assertion Sets. The type system prevents the creation of illegal assertion sets by stratifying types into two levels. The first level, *flat types*, correspond to the plain data suitable for sharing in the dataspace: `Int`, `String`, and so on, as well as type constructors such as lists, structs, and unions when applied to other flat types. The second level is everything else—values that cannot easily be compared for equality, such as functions, objects, and actors. Typing a set-creation form `{e ...}` then checks that each element `e` has a flat type, i.e., *not* a function, object, or actor.

The type `(AssertionSet τ)` describes a set of assertions of type τ arising from an expression `{e ...}`. A patch `(patch e e)` is essentially a pair of assertion sets, thus the type `(Patch $\tau \sigma$)` records the type of assertions to add, τ , and the type of assertions to withdraw, σ .

We use the abbreviation `(Event τ)` to stand for `(Patch $\tau \tau$)`, signifying that dataspace notify actors of both the appearance and disappearance of assertions matching an interest. An actor may request the retraction of an assertion it is not presently making, which is a no-op. An actor may make use of this fact by issuing an overly-broad retraction, alleviating some responsibility for tracking currently made assertions. For example, an actor may submit the patch `(patch {} {*})` to withdraw all of its current assertions. To

account for this fact, we use the abbreviation

$$(\text{Action } \tau \ \sigma) \stackrel{df}{=} (\text{U } (\text{Patch } \tau \ \star) \ (\text{Actor } \sigma))$$

to describe actor actions. The type allows any set of assertions to be withdrawn, as well as the potential to require spawned actors to operate at the communication type σ .

Typed assertion-set projection differs from the untyped version in several ways. The first is a syntactic change to patterns. In order to facilitate type checking, pattern variables come with a type annotation, as in `$name:String`.⁶ The type of this pattern variable is then `(Bind String)`, while the `_` pattern has type `Discard`. Type-checking additionally verifies that the expressions within patterns are well-typed.

Detecting erroneous uses of `project` involves considering paths to binding patterns and `*` in potentially matching types. Recall that `project` raises an error when a binding pattern, such as `$name:String`, has an infinite number of matches in the given set, as in the following expression:

```
(project [(in-room $name:String "FBI") {(in-room * *)}] name)
```

Conceptually, performing the operation requires iterating over the infinite set

$$\{v \mid (\text{in-room } v \ \text{"FBI"}) \in \{(\text{in-room } \star \ \star)\}\}$$

The cause of the error is not simply because the given set is conceptually infinite. Often, the structure of the pattern provides enough information to discriminate most elements of the set. In the expression

```
(project [(in-room $name:String "FBI") {(in-room "Mo" *)}] name)
```

the matched set,

$$\{v \mid (\text{in-room } v \ \text{"FBI"}) \in \{(\text{in-room } \text{"Mo"} \ \star)\}\} = \{\text{"Mo"}\},$$

is finite, thus evaluation poses no issues. An error occurs exactly when a binding variable in the pattern corresponds to a wildcard `*` in the assertion set. The type system therefore tracks both uses—the latter by assigning `*` the type `*`—and analyzes the type of the pattern against the type of the contents of the set. The potential for the pattern to match assertions is determined by computing the overlapping elements of the types. A type error arises only if there is a common path through both types that may potentially match, leading to `(Bind τ)` in the pattern and `*` in the set.

Termination. In order for dataspace programs to make progress, individual actors must terminate—either normally or via an exception—in response to every event. This assumption is easily violated when programming with general-purpose constructs such as functions. The type system disallows recursive functions, but still allows for using recursive data structures such as lists via inductive schemas. Recursive data structures must be used via an inductive eliminator in the shape of a folding loop.

⁶ Our prototype implementation can infer these annotations, but we include them in examples here to clearly separate typed and untyped code.

Actor Subtyping. Subtyping for actor actions aids modular development of actors and permits type-level constraints to be imposed on individual actors, rather than the entire dataspace. For example, the user-agent actor can be developed using a communication type that describes only those conversations in which it participates:

```
(actor UserAgentAssertions ...)
```

yielding a term of type `(Actor UserAgentAssertions)`. Ultimately, however, user-agent actors operate in a dataspace with communication type `ChatDataspace`. Hence, we must check that the difference between the two communication types does not invalidate the reasoning by which type checking the user-agent actor first succeeded.

Checking the actor creation action computes an intersection between the assertions of interest made by the user-agent actor with `UserAgentAssertions` to determine the events it might receive. In a dataspace with a different type, such as `ChatDataspace`, the actor’s interests might match different types of assertions, which may potentially be incompatible with its behavior function type. The question, then, is whether the actor’s interests lead to “surprising” events when operating in a different type of dataspace.

Actor subtyping answers this question. The first item to check is that all of the user-agent’s assertions are allowed in the greater chat dataspace, that is, `UserAgentAssertions` must be a subtype of `ChatDataspace`. Next, we must make sure that running in a `ChatDataspace` context does not yield surprising events, i.e., events not considered the first time we checked the user-agent. The `UserAgentAssertions` type permits one type of interest,

$$\tau_{obs} = (\text{Observe } (\text{InRoom } \star \star))$$

Intersecting τ_{obs} with `ChatDataspace` yields the possible type of events the actor receives when run in the dataspace, `(InRoom String String)`.

Since `(InRoom String String)` is in `UserAgentAssertions`, it is also in the intersection of τ_{obs} and `UserAgentAssertions`. Consequently, `(InRoom String String)` events have already been considered, and deemed safe, against the input type of the user-agent actor’s behavior function during the checking of the corresponding actor form.

3.2 Typing the Chat Room

Figure 5 displays the typed version of the code from figure 3, highlighting the changes. The primary difference between the two figures is the addition of the `UserAgentAssertions` and `ChatDataspace` type abbreviations (figure 4). Otherwise, there are minor changes to insert type annotations on function parameters, binding patterns, and actor-spawning expressions. The example also makes use of the abbreviation \perp for the empty union.

The behavior function comes with a narrow type of input event,

```
(Event (InRoom String String))
```

containing only one of the many forms of assertions with which actors communicate in the dataspace. Because the actor states only one type of interest, the type system is able to verify that incoming events do indeed have such a refined type. The example illustrates how a highly expressive type system can easily validate complex confluences of communication and computation.

```

Typed
(define-type UserAgentAssertions
  ... defined on page 12 ...)

(define-type ChatDataspace
  ... defined in figure 4 ...)

(define (main)
  (dataspace ChatDataspace
    (list
      (create-user-agent "Mo" "FBI")
      (create-room-list ...)
      ...)))

(define (create-user-agent [name : String]
                          [initial-room : String]
                          -> (Actor UserAgentAssertions))

  (actor UserAgentAssertions
    ;; behavior: notify the user as peers enter and
    ;; leave the current room.
    (lambda ([event : (Event (InRoom String String))]
            [current-room : String])
      (define arrivals
        (announce (patch-added event) current-room " arrived."))
      (define departures
        (announce (patch-removed event) current-room " departed."))
      (transition current-room
        (append arrivals departures)))
      initial-room
      {(in-room name initial-room)
       (observe (in-room * initial-room))}))

    ;; Create user-notification patches based on
    ;; presence assertions in an event.
    (define (announce [assertions : (AssertionSet (InRoom String String))]
                [current-room : String]
                [message : String]
                -> (List (Patch (UserNotification String) ⊥)))
      (project [(in-room $who:String current-room) assertions]
        (patch {(user-notification (string-append who message))} {})))

```

Fig. 5. The typed user agent actor

The typed chat dataspace above rules out simple mistakes such as using a number instead of a string for a username. These mistakes can result in non-local actor failure: faults in actors that consume, rather than produce, bad presence information. The `ChatDataspace` communication type prevents such an actor from being introduced into the typed dataspace by specifying that room names are only ever finite sets of strings.

Constraining the User-Agent. The typed chat dataspace permits queries over both the presence of users in a specific room, (`Observe (InRoom * String)`), and over all rooms, (`Observe (InRoom * *)`). The former are used by user-agent actors to monitor the presence of peers in a room while the latter are used to aggregate information about which rooms exist. Since `UserAgentAssertions` includes both `PresenceAssertions` and `RoomAssertions`, the type permits user-agent actors to express both interest in specific room names and wildcard interest in every room. As discussed in section 2.4, the developer may wish to enforce that user-agents do not make overly broad queries. For example, private channel names may leak to an actor that asserts (`observe (in-room * *)`). To address this, the developer may ascribe a refined communication type to user-agent actors:

```
(define-type RestrictedUserAgent
  (U (InRoom String String)
     (Observe (InRoom * String))
     NotificationAssertions))
```

This restrictive type enforces that queries are only over *specific* room names. Actor subtyping permits using the refined user-agent actor in the chat server dataspace.

3.3 Revisiting the Arithmetic Service

Recall the arithmetic service dataspace (section 2.3) and the error that arises from `*` assertions. Types prevent the scenario where sum assertions are abused to request difference calculations. The program may use the communication type

```
(define-type ArithmeticAssertions
  (U (Sum Int Int Int)
     (Observe (Sum Int Int *))
     (Observe (Observe (Sum * * *))))))
```

and this type rules out difference-calculation requests.

Such a difference request assertion, say (`observe (sum 4 * 9)`), would have type (`Observe (Sum Int * Int)`), which is *not* subsumed by `ArithmeticAssertions`. In particular, because `*` stands for *all* types of assertions, including strings and structures, it is incompatible with the occurrence of `Int` in the corresponding position in `ArithmeticAssertions`. An actor with such an output fails to type check in the `ArithmeticAssertions` dataspace, thus preventing the dynamic error explained in section 2.4.

Alternatively, we could have chosen a communication type that is more permissive with regard to interest in sum assertions:

```
(define-type PermissiveInterests
  (U (Sum Int Int Int)
     (Observe (Sum * * *))
     (Observe (Observe (Sum * * *))))))
```

A dataspace of this type allows difference requests. However, now a type error arises for the actor implementing the arithmetic service. Recall that the service iterates over incoming assertion sets e ,

```
(project [(observe (sum $x:Int $y:Int _)) e] (sum x y (+ x y)))
```

This projection occurs inside the body of a behavior function, which assumes some type for the incoming assertions contained by e . If that e has type `PermissiveInterests`, the type system signals an error, since the path to $\$x:\text{Int}$ in the pattern leads to \star in the patch set. By contrast, assuming the contents of the set e have type `ArithmeticAssertions` allows the projection to type check.

When we try to create an actor action with the behavior function for a dataspace with communication type `PermissiveInterests`, the type system detects a mismatch between the assumption and reality. Concretely, the actor makes an assertion of interest (`observe (observe (sum \star \star \star))`) to learn about sum requests, with type

$$(\text{Observe } (\text{Observe } (\text{Sum } \star \star \star)))$$

Potentially matching assertions are the intersection with `PermissiveInterests`,

$$(\text{Observe } (\text{Sum } \star \star \star))$$

which is *not* a subtype of `ArithmeticAssertions`. Finally, checking the actor action succeeds if the supplied communication type annotation is `ArithmeticAssertions`. Actor subtyping prevents instantiating such an actor in the `PermissiveInterests` dataspace through a similar failed type-checking attempt.

The arithmetic service actor typifies the interplay between projection, function, and actor typing. When writing a behavior function, the developer analyzes incoming assertion sets using `project`. In order for the function itself to type check, the set must have a type compatible with the supplied patterns. These constraints flow outward, to the function's parameter for incoming events, and they then become part of the domain of the function's type. When the function reaches an actor action, the type system finds the assumptions in the domain of the function type and compares them with the reality of the surrounding dataspace. Only when all of these elements agree are programs well-typed.

4 The Semantics of Actors

4.1 Specifying an Actor

Figure 6 introduces λ_{ds} , a model language for articulating untyped dataspace actors. It is representative of our ECMAScript and Racket prototypes, i.e., functional languages that extend the λ -calculus syntax with means to interface with dataspace. In our examples we also use `let` and `if` with their usual meaning.

A complete program in λ_{ds} is a description of a dataspace, dataspace M . The expression M computes a list of actors to launch at the start of the program. Over the course of execution, additional actors may be dynamically spawned through actor actions as well as removed due to failure or termination. Section 5 describes how this initial description yields a running actor system.

Extensions to the base functional model compute values that, when sent to the surrounding dataspace, trigger certain actions: an assertion of interest (`observe M`); an assertion of fact (`$m(\vec{M})$`); the spawning of a new actor (`actor $M_b M_s M_a$`). The three parts of actor correspond to the behavior function (M_b), private state (M_s), and initial actions (M_a) of an actor. Additionally, the extensions also include expression forms to create sets of assertions ($\{\vec{SK}\}$) and compose patches (M^+/M^-).

Behavior functions in λ_{ds} return a pair of a new state value and a list of actions, as opposed to the `quit` and `transition` records of our prototypes described in section 2.1. The capabilities of the two interfaces are the same; rather than issue an explicit `quit`, actors in λ_{ds} signal termination by raising an `error`, triggering their removal from the dataspace.

A set constructor `SK` describes assertions, ranging from singletons to infinite sets (\star); $\{\overrightarrow{SK}\}$ translates to an assertion set π (from figure 7) for placement in the dataspace. When an `SK` expression yields a value not suitable for dataspace assertions, such as a function, the reduction ends in an `error`. The different variants of assertions employed in a dataspace program correspond to a set of message constructors m , which otherwise behave like tuples.

The `project` form is the key mechanism for de-structuring incoming assertion sets in a behavior function. Specifically,

$$\text{project } \pi \text{ with PAT in } M$$

instantiates M with the bindings of `PAT` for each matching assertion in π and assembles the results into a list. A pattern's bindings may match an infinite number of values, as in

$$\text{project } \{\star\} \text{ with } \$x \text{ in } x$$

Our semantics interprets such expressions as errors.

4.1.1 The Reduction Semantics of λ_{ds}

The reduction semantics of λ_{ds} is mostly conventional (Felleisen *et al.*, 2009); see figure 7. The addition of assertions c and assertion sets π requires only a small extension over the standard call-by-value semantics. Assertion sets are created from set constructors with the metafunction `make-set`. The `project` metafunction implements projection, which eliminates assertion sets. Appendix A provides the full definition of each metafunction.

The semantics distinguishes among three sources of errors in order to characterize the soundness of our model precisely:

1. `errorprim` arises from application of partial primitive operations;
2. `errorh-o` arises from assertion sets containing functions or actors;
3. `errorinf` arises when `project` selects an infinite subset of assertions.

$I \in \mathbf{Init}$ = dataspace M	$\text{PAT} \in \mathbf{Pat}$ = $\$x \mid -$ $\mid M$ $\mid (\overrightarrow{\text{PAT}})$ $\mid \text{observe PAT} \mid m(\overrightarrow{\text{PAT}})$
$M \in \mathbf{Expr}$ = $\lambda x.M \mid MM \mid x \mid p \overrightarrow{M} \mid b$ $\mid \text{error}_\eta$ $\mid (\overrightarrow{M}) \mid \text{cons } M M$ $\mid \text{observe } M \mid m(\overrightarrow{M})$ $\mid \{\overrightarrow{SK}\} \mid M^+ / M^-$ $\mid \text{actor } M_b M_s M_a$ $\mid \text{project } M \text{ with PAT in } M$	$\text{SK} \in \mathbf{SetCons}$ = \star $\mid M$ $\mid (\overrightarrow{SK})$ $\mid \text{observe SK} \mid m(\overrightarrow{SK})$
$m \in \mathbf{Msg}$ = ... message constructors $p \in \mathbf{Prim}$ = ... primitives $b \in \mathbf{BasicVal}$ = ... base values	$\eta \in \mathbf{ErrorTag}$ = <code>prim</code> \mid <code>h-o</code> \mid <code>inf</code>

Fig. 6. The Syntax of λ_{ds}

Evaluation Syntax

$M = \dots \mid \pi$	$c \in \mathbf{Assertion} = b \mid (\vec{c}) \mid \mathbf{cons} \ c \ c \mid \mathbf{observe} \ c \mid m(\vec{c})$
$v \in \mathbf{Val} = \lambda x. M \mid b$	$\pi \in \mathbf{ASet} = \mathcal{P}(c)$
$\begin{array}{l} (\vec{v}) \\ \mathbf{cons} \ v \ v \\ \mathbf{observe} \ v \mid m(\vec{v}) \\ \pi \\ v^+ / v^- \\ \mathbf{actor} \ v_b \ v_s \ v_a \\ * \mid \$x \mid - \end{array}$	$E \in \mathbf{Ctx} = \square \mid EM \mid vE \mid p \vec{v} E \vec{M} \mid (\vec{v} E \vec{M})$ $\begin{array}{l} \mathbf{cons} \ E \ M \mid \mathbf{cons} \ v \ E \\ \mathbf{observe} \ E \mid m(\vec{v} E \vec{M}) \\ \{ \vec{v} E \vec{SK} \} \mid E/M \mid v/E \\ \mathbf{actor} \ E \ M \ M \mid \mathbf{actor} \ v \ E \ M \\ \mathbf{actor} \ v \ v \ E \\ \mathbf{project} \ E \ \text{with PAT in } M \\ \mathbf{project} \ v \ \text{with } E \ \text{in } M \end{array}$

Evaluation

$$eval(M) = \begin{cases} v & \text{if } M \longrightarrow^* v \\ \mathbf{error}_\eta & \text{if } M \longrightarrow^* \mathbf{error}_\eta \end{cases}$$

Notions of Reduction

$E[(\lambda x. M) v] \longrightarrow E[M[x \leftarrow v]]$	(β_v)
$E[p \vec{v}] \longrightarrow E[v']$	where $\delta(p, \vec{v}) = v'$ (δ)
$E[p \vec{v}] \longrightarrow \mathbf{error}_{\text{prim}}$	where $\delta(p, \vec{v})$ is undefined (δ -error)
$E[\{\vec{v}\}] \longrightarrow E[\pi]$	where $make\text{-set}(\vec{v}) = \pi$ (make-set)
$E[\{\vec{v}\}] \longrightarrow \mathbf{error}_{\text{h-o}}$	where $make\text{-set}(\vec{v})$ is undefined (make-set-error)
$E[\mathbf{project} \ \pi \ \text{with } v \ \text{in } M] \longrightarrow E[M']$	where $project(\pi, v, M) = M'$ (project)
$E[\mathbf{project} \ \pi \ \text{with } v \ \text{in } M] \longrightarrow \mathbf{error}_{\text{inf}}$	where $project(\pi, v, M)$ is undefined (project-error)

Metafunctions

$make\text{-set}$	$\vec{\mathbf{Val}} \longrightarrow_{\text{partial}} \mathbf{ASet}$
$make\text{-set}(\vec{v})$	translates a vector of values into a dataspace representation undefined if given higher-order values
$project$	$\mathbf{ASet} \times \mathbf{Val} \times \mathbf{Expr} \longrightarrow_{\text{partial}} \mathbf{Expr}$
$project(\pi, v_p, M)$	creates a list by replacing pattern variables from v_p , an evaluated pattern, in M with values from matching assertions in π undefined if there is an infinite number of different matches
δ	$\mathbf{Prim} \times \vec{\mathbf{Val}} \longrightarrow_{\text{partial}} \mathbf{Val}$
	applies a primitive; undefined in cases due to partial primitives

Fig. 7. The Essence of the Formal Semantics of λ_{ds} **4.2 Types for Dataspaces and Actors**

Figure 8 extends the syntax of λ_{ds} with simple types, giving rise to λ_{ds}^U . The primary difference is the presence of type annotations in functions, actors, dataspace, and patterns; additionally, expressions now include $\mathbf{fold} \ M_c \ M_n \ M_t$, a representative induction scheme

$I = \text{dataspace } \tau_c M$ $M = \dots$ $\quad \lambda x : \tau. M$ $\quad \text{fold } M M M$ $\quad \text{actor } \tau_c M_b M_s M_a$ $\text{PAT} = \dots$ $\quad \$x : \tau$ $v = \dots$ $\quad \lambda x : \tau. M$ $\quad \text{actor } \tau_c v_b v_s v_a$ $\quad \text{dataspace } \tau_c v$ $\quad \$x : \tau$	$\tau, \sigma \in \mathbf{Type} = \tau \rightarrow \tau \mid B$ $\quad \text{Observe } \tau \mid m(\vec{\tau})$ $\quad (\vec{\tau}) \mid \text{List } \tau$ $\quad \text{AssertionSet } \tau \mid \text{Patch } \tau \tau$ $\quad \bigcup \vec{\tau}$ $\quad \text{Actor } \tau$ $\quad \star$ $\quad \$: \tau \mid \text{Discard}$ $B \in \mathbf{BaseTy} = \text{base types: String, Int, etc.}$ $\Gamma \in \mathbf{Env} = \overline{x : \tau}$ $\perp \stackrel{df}{=} \bigcup.$ $\text{Action } \tau \sigma \stackrel{df}{=} (\text{Patch } \tau \star) \cup (\text{Actor } \sigma)$
---	--

Fig. 8. Typed Syntax of λ_{ds}

for iterating over lists. Reduction of `fold` terms is standard; lists unfold to applications of M_c with base case M_n .

The language of types reflects the underlying expression language and adds union types for dealing with dataspace communication. A union of types is written $\bigcup \vec{\tau}$. When convenient, we use the infix notation $\tau \cup \tau$. We do not provide any elimination forms for union types. Rather, we use unions primarily for describing the contents of assertion sets and rely on the pattern supplied to `project` to discriminate the branches of a union. The definition of the *project* metafunction is the same as the one in the untyped model, relying on the structure of the pattern to find matching assertions. Therefore, the type of the pattern must predict how matching will proceed during evaluation. For example, let M describe a set of assertions τ , where

$$\tau = m_1(\text{String}) \cup m_2(\text{Int}) \cup m_3()$$

Using a pattern that describes only one of the message constructors, m_2 , as in

$$\text{project } M \text{ with } m_2(\$x : \text{Int}) \text{ in } x$$

focuses type checking on only those members of the union that might match the pattern, allowing the projection to ignore m_1 and m_3 . In a full implementation, we also expect a discipline such as occurrence typing (Tobin-Hochstadt & Felleisen, 2008) to facilitate programming with unions.

The type `Actor` τ_c is that of actors in a dataspace with communication type τ_c .

The data constructor `observe` is directly reflected as the type constructor `Observe`. Similarly, $m(\vec{\tau})$ is the type of a message constructed by m with fields $\vec{\tau}$. The type \star is the type of assertions created with \star . Thus, the expression `{observe in-room(\star , "FBI")}` has type

$$\text{AssertionSet } (\text{Observe in-room}(\star, \text{String}))$$

The type `Patch` $\tau^+ \tau^-$ describes patches M^+ / M^- where M^+ has type `AssertionSet` τ^+ and M^- type `AssertionSet` τ^- . Finally, $\$: \tau$ is the type of patterns $\$x : \tau$, `Discard` is the type of `_`, and `Action` $\tau \sigma$ abbreviates $(\text{Patch } \tau \star) \cup (\text{Actor } \sigma)$.

4.3 Static Semantics

The purpose of the typing rules for λ_{ds}^U in figure 9 is to establish three key invariants using a standard judgment. First, actors place only valid data in their dataspace, i.e., first-order sets of assertions at the specified communication type. Second, actors extract only finite subsets of assertions via `project`. Third, well-typed actors are terminating; that is, they either signal an error or terminate with a list of actions and a new private state value. This last invariant discharges an assumption of the universal soundness theorem of Garnock-Jones et al (2016).

$$\begin{array}{c}
 \boxed{\vdash I} \\
 \frac{\text{flat}(\tau_c) \quad \vdash M_{boot} : \text{List}(\text{Actor } \tau_c)}{\vdash \text{dataspace } \tau_c M_{boot}} \boxed{\text{T-DATASPACE}} \\
 \boxed{\Gamma \vdash M : \tau} \\
 \frac{\Gamma \vdash x : \tau \quad \tau <: \sigma}{\Gamma \vdash x : \tau} \boxed{\text{T-SUB}} \\
 \frac{\Gamma \vdash M_{behavior} : (\text{Patch } \tau_{in} \tau_{in}, \tau_{state}) \rightarrow (\text{List}(\text{Action } \tau_{out} \tau_c), \tau_{state}) \quad \Gamma \vdash M_{state} : \tau_{state} \quad \Gamma \vdash M_{assertions} : \text{AssertionSet } \tau_{out} \quad \tau_{out} <: \tau_c \quad \text{predict-routing}(\tau_{out}, \tau_c) <: \tau_{in} \quad \text{flat}(\tau_c)}{\Gamma \vdash \text{actor } \tau_c M_{behavior} M_{state} M_{assertions} : \text{Actor } \tau_c} \boxed{\text{T-ACTOR}} \\
 \frac{\Gamma \vdash M_s : \text{AssertionSet } \tau_s \quad \Gamma \vdash \text{PAT} : \tau_p \quad \text{safe}(\tau_s, \tau_p) \quad \text{bindings}(\text{PAT}) = \Gamma' \quad \Gamma, \Gamma' \vdash M_b : \tau_b}{\Gamma \vdash \text{project } M_s \text{ with PAT in } M_b : \text{List } \tau_b} \boxed{\text{T-PROJECT}} \\
 \frac{\overrightarrow{\Gamma \vdash \text{SK} : \tau}}{\Gamma \vdash \{\overrightarrow{\text{SK}}\} : \text{AssertionSet } \bigcup \overrightarrow{\tau}} \boxed{\text{T-SET}} \quad \frac{\pi \models \tau}{\Gamma \vdash \pi : \text{AssertionSet } \tau} \boxed{\text{T-}\pi} \\
 \frac{}{\Gamma \vdash \text{error}_\eta : \tau} \boxed{\text{T-ERROR}}
 \end{array}$$

Fig. 9. Selected Elements of an Actor Type System

The typing rule for dataspace programs, T-DATASPACE, ensures that the communication type τ_c describes assertions with the premise $\text{flat}(\tau_c)$. The flat judgment identifies the types of basic, first-order data such as numbers, strings, tuples of numbers, and so on suitable for dataspace assertions. Its complete definition is in appendix A. The rule also checks that M_{boot} , the boot actors, all safely operate with communication type τ_c .

According to T-ACTOR, an actor of shape `actor` $\tau_c M_b M_s M_a$ may participate in a dataspace of type τ_c when the type of its behavior function, M_b , fits the template

$$(\text{Patch } \tau_{in} \tau_{in}, \tau_{state}) \rightarrow (\text{List}(\text{Action } \tau_{out} \tau_c), \tau_{state})$$

and satisfies certain conditions concerning τ_{in} , τ_{out} , and τ_c . Specifically, all assertions produced by the behavior function, τ_{out} , must be valid utterances in τ_c . Next, τ_{in} , the type of assertions the actor is prepared to handle, must account for the actor's interests. Intuitively, an actor must be prepared to receive all of the assertions it asks for. The type τ_{in} describes

$$\begin{array}{l}
\textit{predict-routing} \quad : \quad \mathbf{Type} \times \mathbf{Type} \longrightarrow \mathbf{Type} \\
\textit{predict-routing}(\tau_o, \tau_c) = \textit{strip-obs}(\tau_o) \tilde{\cap} \tau_c \\
\hline
\textit{strip-obs} \quad : \quad \mathbf{Type} \longrightarrow \mathbf{Type} \\
\textit{strip-obs}(\mathbf{Observe} \tau) = \tau \\
\textit{strip-obs}(\star) = \star \\
\textit{strip-obs}(\bigcup \vec{\tau}) = \bigcup \overrightarrow{\textit{strip-obs}(\tau)} \\
\textit{strip-obs}(\tau) = \perp \quad \text{otherwise} \\
\hline
\tilde{\cap} \quad : \quad \mathbf{Type} \times \mathbf{Type} \longrightarrow \mathbf{Type} \\
\tau \tilde{\cap} \tau = \tau \\
\bigcup \vec{\tau} \tilde{\cap} \sigma = \bigcup \overrightarrow{\tau \tilde{\cap} \sigma} \\
\tau \tilde{\cap} \bigcup \vec{\sigma} = \bigcup \overrightarrow{\tau \tilde{\cap} \sigma} \\
\star \tilde{\cap} \tau = \tau \\
\tau \tilde{\cap} \star = \tau \\
\mathbf{List} \tau \tilde{\cap} \mathbf{List} \sigma = \mathbf{List} (\tau \tilde{\cap} \sigma) \\
() \tilde{\cap} () = () \\
(\tau_1, \vec{\tau}_2) \tilde{\cap} (\sigma_1, \vec{\sigma}_2) = \begin{cases} \perp & \text{if } \tau_{11} <: \perp \text{ or } \tau_{22} <: \perp \\ (\tau_{11}, \vec{\sigma}_{22}) & \text{if } \tau_{22} = (\vec{\sigma}_{22}) \end{cases} \\
\text{where } \tau_{11} = \tau_1 \tilde{\cap} \sigma_1, \tau_{22} = (\vec{\tau}_2) \tilde{\cap} (\vec{\sigma}_2) \\
m(\vec{\tau}_n) \tilde{\cap} m(\vec{\sigma}_n) = \begin{cases} m(\vec{\sigma}) & \text{if } \tau = (\vec{\sigma}) \\ \perp & \text{otherwise} \end{cases} \\
\text{where } \tau = (\vec{\tau}_n) \tilde{\cap} (\vec{\sigma}_n) \\
\mathbf{Observe} \tau \tilde{\cap} \mathbf{Observe} \sigma = \begin{cases} \perp & \text{if } \tau' <: \perp \\ \mathbf{Observe} \tau' & \text{otherwise} \end{cases} \\
\text{where } \tau' = \tau \tilde{\cap} \sigma \\
\tau \tilde{\cap} \sigma = \perp \quad \text{otherwise}
\end{array}$$

Fig. 10. Key Support Metafunctions

the actor's assumptions about the sets of assertions it receives, including which subsets may be infinite. The actual assertions it receives arise from dataspace routing, which matches stated interests against all current assertions. The *predict-routing* metafunction, defined in figure 10, mirrors run-time routing in order to predict the types of assertions received by an actor. It computes the overlap between the types of interests expressed by an actor and the types of possible assertions in the dataspace. The assertions the actor may express interest in are exactly those prefixed by `Observe` in τ_{out} ; the metafunction *strip-obs* finds all such types, while the $\tilde{\cap}$ metafunction determines the type representation of the overlap between such interests and the potential assertions in the dataspace, τ_c (figure 10). Since \star stands for all possible assertions, including `observe`-prefixed ones, *strip-obs* treats \star as if it were the unfolding `Observe` \star . Finally, the second parameter of `Action` $\tau_{out} \tau_c$ in T-ACTOR requires the type of any spawned actor to conform to communication type τ_c .

Rule T-PROJECT eliminates two potential problems from matching a pattern with type τ_p against a set of type `AssertionSet` τ_s . First, binding variables in the pattern may have an infinite number of matches. Second, matching assertions may be incompatible with the

$$\begin{array}{c}
\frac{\text{flat}(\tau)}{\tau <: \star} \boxed{\text{S-}\star} \qquad \frac{\tau <: \sigma \quad \text{predict-routing}(\tau, \sigma) <: \tau}{\text{Actor } \tau <: \text{Actor } \sigma} \boxed{\text{S-ACTOR}} \qquad \boxed{\tau <: \sigma} \\
\\
\frac{\text{finite}(\tau) \quad \tau <: \sigma}{\text{safe}(\tau, \$: \sigma)} \boxed{\text{PS-CAPTURE}} \qquad \frac{}{\text{safe}(\tau, \text{Discard})} \boxed{\text{PS-DISCARD}} \qquad \boxed{\text{safe}(\tau, \sigma)} \\
\\
\frac{}{\pi \models \star} \boxed{\text{M-}\star} \quad \frac{\pi \subseteq \mathbb{Z} \quad \exists n. |\pi| = n}{\pi \models \text{Int}} \boxed{\text{M-INT}} \quad \frac{\pi' \models \tau}{\pi \models \text{Observe } \tau} \quad \frac{\pi \subseteq \{\text{observe } v \mid v \in \pi'\}}{\pi \models \text{Observe } \tau} \boxed{\text{M-OBSERVE}} \quad \boxed{\pi \models \tau}
\end{array}$$

Fig. 11. Key Support Judgments

type associated with binding variables, $\$: \tau$. The $\text{safe}(\tau_s, \tau_p)$ judgment enforces this constraint by finding the portion of the analyzed assertion set that may match and bind variables in the pattern. Both aspects are checked by rule PS-CAPTURE (figure 11), which corresponds to an identified match. The $\text{finite}(\tau)$ premise ensures matched assertion types do not contain \star , while $\tau <: \sigma$ checks that they meet the pattern's expectations. Appendix A provides the full definition of each judgment.

The T- π rule describes assertion sets π , which are not a part of the surface syntax of λ_{ds}^U . They arise through the evaluation of $\{\vec{SK}\}$ forms and via dataspace event dispatch. Since we wish to employ a standard progress-and-preservation proof technique, we must assign them types. The $\pi \models \tau$ judgment (figure 11; appendix A) checks that the set π corresponds to the structure of type τ in a way that allows it to be used by λ_{ds}^U actors.

For example, consider the type for interest in presence of a chat room:

$$\tau_q = \text{Observe}(\text{in-room}(\star, \text{String}))$$

A set π has type $\text{AssertionSet } \tau_q$ under these conditions:

- Every element of π must be a message constructed with `in-room` holding two fields, prefixed by the `observe` constructor.
- The set of values appearing in the second slot of each message must be a finite set of strings. That is, $\{v_2 \mid \text{observe in-room}(v_1, v_2) \in \pi\}$ is a finite set of `Strings`.
- The set of values appearing in the first slot of each message may be any set of assertions, including infinite sets.

The subtyping judgment $\text{Actor } \tau <: \text{Actor } \sigma$ ensures that every utterance in τ is a valid utterance in σ . Furthermore, the judgment must also check for the possibility of interference. That is, transplanting the τ -typed actor's subscriptions to the new σ -typed dataspace must not result in assertions that the actor is unprepared to handle. As in T-ACTOR, this condition is checked by computing $\text{predict-routing}(\tau, \sigma)$, a type that reflects routing in dataspace.

4.4 Properties

The typed language of dataspace actors satisfies a number of critical properties, most importantly, type soundness and termination.

Lemma 1 (Progress). *If $\vdash M : \tau$ and $M \neq \text{error}_\eta$, then either $M \in \mathbf{Val}$ or there exists an M' such that $M \longrightarrow M'$.*

Proof By case analysis on the shape of M . \square

Lemma 2 (Preservation). *If $\vdash M : \tau$ and $M \longrightarrow M'$, then $\vdash M' : \tau$.*

Proof By case analysis on possible reductions $M \longrightarrow M'$. The proof relies on several auxiliary lemmas about operations on assertion sets. In particular, lemmas 3 and 4 show that creating assertion sets never causes an error and uses of `project` always select finite subsets of assertions. \square

Lemma 3 (Soundness of Creating Assertion Sets). *If for some \vec{v} and τ , $\vdash \{\vec{v}\} : \text{AssertionSet } \tau$, then there exists π such that $\text{make-set}(\vec{v}) = \pi$ and $\pi \models \tau$.*

Proof By induction on the typing derivation. \square

Lemma 4 (Soundness of `project`). *If $\Gamma \vdash \text{project } \pi$ with v_p in $M : \tau$, there exists an M' such that $\text{project}(\pi, v_p, M) = M'$, $\Gamma \vdash M' : \tau$.*

Proof By induction on the typing derivation. \square

Theorem 5 (Soundness & Termination). *If $\vdash M : \tau$ and $\text{error}_\eta \notin M$, then either*

1. $M \longrightarrow^* v$ and $\vdash v : \tau$; or
2. $M \longrightarrow^* \text{error}_{\text{prim}}$.

Interpretation The second case of theorem 5 implies that errors may only arise due to the application of partial primitives, never through malformed communication (`errorh-o`) or touching infinite sets of assertions (`errorinf`). \square

Proof Starting from a well-typed term, we employ the usual progress (lemma 1) and preservation (lemma 2) lemmas (Wright & Felleisen, 1994) to show soundness.

To show that reduction sequences terminate with a value, we use the standard “candidate” technique (Girard, 1971). One salient detail of the proof is that, by the nature of assertions, the pattern variables in `project` are always instantiated with first-order values.

\square

The reduction semantics of λ_{ds}^{\cup} programs suffices to describe how an individual actor’s behavior function computes a response to an event. It does not explain, however, any properties of *dataspace actors* specified with λ_{ds}^{\cup} , that is, the meaning of a complete program `dataspace` $\tau_c M$. In order to understand how systems of actors interact, we must examine the specifics of dataspace coordination, allowing us to pose and answer questions concerning actor behavior with the proper context.

5 The Complete Dataspace Model

Dataspace coordination permits groups of actors to share knowledge via a common repository, events, and actions. Events arise when new information appears and is relevant to a stated interest of an actor. Section 5.1 gives a formal model of dataspace coordination. The model abstracts over the details of how an individual actor responds to an event, calling only for a total function that transforms a current state value to a new state and some actions. Implementing actor behavior with a specific language, such as λ_{ds}^U , requires showing that it lives up to the expected interface. It is imperative that the internal actor language and dataspace coordination model share a common ontology, enabling them to exchange (representations of) events and actions. Section 5.2 illustrates how to link the two parts of the model. Section 5.3 shows that both models meet each others' expectations, thus enjoying the desired soundness property.

5.1 Recap of the Dataspace Syntax and Semantics

Figure 12 provides a syntax for the dataspace part of our model. The internal behavior of each actor is a mathematical function. This interface is an abstraction point; it separates the concerns of an individual actor implementation from the details of dataspace concurrency and coordination. Consequently, we can just as easily use the dataspace configuration syntax and semantics with λ_{ds}^U , a model of a Racket-like functional language, as with an object-oriented model based on our ECMAScript prototype. To minimize bookkeeping, the grammar reuses the syntax of events, values, and actions from section 4.

Programs	$P \in \mathbf{Prog}$	=	dataspace \vec{S}
Actor Specifications	$S \in \mathbf{Specs}$	=	actor $f \nu \pi$
Behavior Functions	$f \in \mathbf{BehFun}$	=	$\mathbf{Event} \times \mathbf{Val} \xrightarrow{\text{total}} \mathbf{Action} \times \mathbf{Val} + \mathbf{Error}$
Events	$e \in \mathbf{Event}$	=	Δ
Actions	$a \in \mathbf{Action}$	=	$\Delta \mid S$
Patches	$\Delta \in \mathbf{Patch}$	=	π^+ / π^- where $\pi^+ \cap \pi^- = \emptyset$
Errors	$\dagger \in \mathbf{Error}$	=	\mathbf{error}_{ds}
Assertion sets	$\pi \in \mathbf{ASet}$	(defined in fig. 7)	
Assertions	$c \in \mathbf{Assertion}$	(defined in fig. 7)	
Values	$v \in \mathbf{Val}$	(defined in fig. 7)	

Fig. 12. Syntax of Dataspaces

Writing down the reduction semantics for dataspace programs—modulo the semantics for internal actor behavior—requires three elements: (1) a generalized syntax to specify intermediate configurations; (2) several basic notions of reduction; and (3) metafunctions to keep the formulation of the reductions concise.

Evaluation Syntax. Figure 13 defines the evaluation syntax of dataspace programs. A running dataspace configuration C contains a pending action queue, the set of all current assertions, and the contained actors. Each actor is represented by an internal name ℓ and a state Σ . The state of an actor consists of a queue of events to handle, a behavior B comprised of a function and current private state value (f, v) , and a queue of actions that need to

be processed by the surrounding dataspace. Quiescent terms are those without pending actions; inert terms have neither pending actions nor events to handle.

Configurations	$C \in \mathbf{Config} = \overrightarrow{[(k, a)]; R; \vec{A}}$	Quiescent Terms
Behaviors	$B \in \mathbf{ABehav} = (f, v)$	$A_Q \in \mathbf{Actor}_Q = \ell \mapsto \Sigma_Q$
Actors	$A \in \mathbf{Actor} = \ell \mapsto \Sigma$	$\Sigma_Q \in \mathbf{AState}_Q = \vec{e} \triangleright B \triangleright \cdot$
Actor States	$\Sigma \in \mathbf{AState} = \vec{e} \triangleright B \triangleright \vec{a}$	
Assertion Tables	$R \in \mathbf{DS} = \mathcal{P}(\mathit{Lift}(\mathbf{Loc}) \times \mathbf{ASet})$	
Local Labels	$j, k, \ell \in \mathbf{Loc} = \mathcal{N}$	
boot : Specs + Prog \rightarrow AState + Config		
	$\mathit{boot}(\mathit{actor} f v \pi) = \cdot \triangleright (f, v) \triangleright \pi / \emptyset$	Inert Terms
	$\mathit{boot}(\mathit{dataspace} \vec{S}) = \overrightarrow{[(0, \vec{S})]; \emptyset; \cdot}$	$C_I \in \mathbf{Config}_I = [\cdot; R; \vec{A}_I]$
		$A_I \in \mathbf{Actor}_I = \ell \mapsto \Sigma_I$
		$\Sigma_I \in \mathbf{AState}_I = \cdot \triangleright B \triangleright \cdot$
Evaluation Contexts	$E^\Sigma = [\cdot; R; \vec{A}_I(\ell \mapsto \square) \vec{A}_Q]$	

Fig. 13. Evaluation Syntax and Inert and Quiescent Terms of Dataspaces

Reduction Relation. Figure 14 presents the reduction semantics of dataspace. The \longrightarrow_Σ relation operates on individual actor states Σ , while the \longrightarrow_{ds} relation describes the reduction of dataspace configurations, with the aim of reaching quiescent or even inert states:

- notify hands an event to the behavior function of an actor and records the new state plus ensuing actions;

$$\begin{aligned} \vec{e} e_0 \triangleright (f, v) \triangleright \vec{a} &\longrightarrow_\Sigma \vec{e} \triangleright (f, v') \triangleright \vec{a}' \vec{a} && \boxed{\text{when } f(e_0, v) = (\vec{a}', v')} && \text{(notify)} \\ \vec{e} e_0 \triangleright (f, v) \triangleright \vec{a} &\longrightarrow_\Sigma \cdot \triangleright (\lambda eu. (\cdot, v), v) \triangleright \emptyset \vec{a} && \boxed{\text{when } f(e_0, v) \in \mathbf{Error}} && \text{(exception)} \end{aligned}$$

$$\overrightarrow{[(k', a)](k, S); R; \vec{A}_Q} \longrightarrow_{ds} \overrightarrow{[(k', a)]; R; \vec{A}_Q}(\ell \mapsto \mathit{boot}(S)) \quad \text{(spawn)}$$

where ℓ distinct from k , every k' , and the labels of every A_Q

$$\overrightarrow{[(k, a)]; R; \vec{A}_Q} A_{out}(a'') \vec{A} \longrightarrow_{ds} \overrightarrow{[(\ell, a'')](k, a); R; \vec{A}_Q} A_{out}(\cdot) \vec{A} \quad \text{(gather)}$$

where $A_{out}(a) = \ell \mapsto \vec{e}' \triangleright B \triangleright \vec{a}' a$

$$\overrightarrow{[(k', a)](k, \pi^+ / \pi^-); R; \vec{A}_Q} \longrightarrow_{ds} \overrightarrow{[(k', a)]; R \oplus (k, \Delta'); \overline{\mathit{bc}_\Delta(k, \Delta', R, A_Q)}} \quad \text{(patch)}$$

where $\Delta' = (\pi^+ - \{c \mid (k, c) \in R\}) / (\pi^- \cap \{c \mid (k, c) \in R\})$

$$E^\Sigma[\Sigma_Q] \longrightarrow_{ds} \mathit{round-robin}(E^\Sigma, \Sigma') \quad \text{if } \Sigma_Q \longrightarrow_\Sigma \Sigma' \quad \text{(schedule)}$$

Fig. 14. Reduction Semantics of Dataspaces

$$\begin{aligned}
bc_{\Delta} &: \mathbf{Loc} \times \mathbf{Event} \times \mathbf{DS} \times \mathbf{Actor}_Q \longrightarrow \mathbf{Actor} \\
bc_{\Delta}(k, \pi_{add}/\pi_{del}, R^{old}, \ell \mapsto \vec{e} \triangleright B \triangleright \cdot) &= \begin{cases} \ell \mapsto \Delta_{fb} \vec{e} \triangleright B \triangleright \cdot & \text{if } \ell = k \text{ and } \Delta_{fb} \neq \emptyset/\emptyset \\ \ell \mapsto \Delta_{other} \vec{e} \triangleright B \triangleright \cdot & \text{if } \ell \neq k \text{ and } \Delta_{other} \neq \emptyset/\emptyset \\ \ell \mapsto \vec{e} \triangleright B \triangleright \cdot & \text{otherwise} \end{cases} \\
\text{where} & \\
\Delta_{fb} &= \{c \mid c \in \pi_{\bullet, add}, (\ell, \mathbf{observe } c) \in R^{new}\} \cup \{c \mid c \in (\pi_{\circ} \cup \pi_{\bullet, add} - \pi_{\bullet, del}), \mathbf{observe } c \in \pi_{add}\} / \\
&\quad \{c \mid c \in \pi_{\bullet, del}, (\ell, \mathbf{observe } c) \in R^{old}\} \cup \{c \mid c \in \pi_{\circ}, \mathbf{observe } c \in \pi_{del}\} \\
\Delta_{other} &= \{c \mid c \in \pi_{\bullet, add}, (\ell, \mathbf{observe } c) \in R^{old}\} / \{c \mid c \in \pi_{\bullet, del}, (\ell, \mathbf{observe } c) \in R^{old}\} \\
\pi_{\bullet} &= \{c \mid (j, c) \in R^{old}, j \neq k\} \\
R^{new} &= R^{old} \oplus (\ell, \pi_{add}/\pi_{del}) \\
\pi_{\bullet, add} &= \pi_{add} - \pi_{\bullet} \\
\pi_{\circ} &= \{c \mid (j, c) \in R^{old}\} \\
\pi_{\bullet, del} &= \pi_{del} - \pi_{\bullet}
\end{aligned}$$

Fig. 15. Dataspace Routing

- exception terminates an actor that raises an exception;
- spawn creates a new actor;
- gather enqueues an action for the dataspace;
- patch realizes a state-change notification; and
- schedule selects which actor to run.

Metafunctions. Figure 15 defines bc_{Δ} , the metafunction that implements routing and the semantics of `observe`. It computes the relevant changes to an actor with label ℓ based on a patch made by actor with label k . In the case that $\ell = k$, it must also consider the possibility that the actor’s interests change as a result of the patch. In that case, the constructed patch Δ_{fb} reflects the most up-to-date interests of the actor.

Appendix A provides the definitions of two additional metafunctions:

- \oplus updates an actor’s assertions in a dataspace;
- *round-robin* rotates the actors in the configuration.

5.2 Connecting the Dots

There is still one divide left to bridge: the models of dataspaces from section 5.1 and the language for programming individual actors from section 4 employ different notions of actor behavior functions. On one side, dataspaces appeal to abstract mathematical functions, while on the other side λ_{ds}^{\cup} provides concrete λ -terms, symbolic descriptions of functions given meaning through a reduction relation. This difference must be reconciled in order to combine the two semantics into a complete understanding of dataspace programs. The two boxes in figure 14 point out the two positions where a base language of computation interfaces with the dataspace part of the overall language.

The *eval* function for λ_{ds}^{\cup} (figure 7) provides a first step for reconciling the difference between the two parts of the model. It is a mechanism for accessing λ_{ds}^{\cup} terms as if they were mathematical functions. Its signature, though,

$$eval : \mathbf{Expr} \longrightarrow \mathbf{Val}$$

$$\begin{aligned}
\mathit{behavior}_{\lambda_{ds}^{\cup}} & : \mathbf{Event} \times \mathbf{Val} \longrightarrow \overrightarrow{\mathbf{Action}} \times \mathbf{Val} + \mathbf{Error} \\
\mathit{behavior}_{\lambda_{ds}^{\cup}}(e, (v_f, v_s)) & = \begin{cases} (\overrightarrow{[\overline{v_a}]}, (v_f, v'_s)) & \text{if } v = (v_{list}, v'_s) \\ \mathbf{error}_{ds} & \text{if } v = \mathbf{error}_{\eta} \end{cases} \\
& \text{where } M_{dispatch} = v_f([e], v_s) \\
& \quad v = \mathit{eval}(M_{dispatch}) \\
& \quad \overrightarrow{v_a} = \mathit{seq}(v_{list})
\end{aligned}$$

$$\begin{aligned}
[\] & : \mathbf{Event} \longrightarrow \mathbf{Val} \\
[\pi^+/\pi^-] & = \pi^+/\pi^-
\end{aligned}$$

$$\begin{aligned}
[\] & : \mathbf{Val} \longrightarrow \mathbf{Action} \\
[\pi^+/\pi^-] & = \pi^+ / (\pi^- - \pi^+) \\
[\mathbf{actor} \ \tau \ v_f \ v_s \ \pi] & = \mathbf{actor} \ \mathit{behavior}_{\lambda_{ds}^{\cup}}(v_f, v_s) \ \pi
\end{aligned}$$

$$\begin{aligned}
\mathit{seq} & : \mathbf{Val} \longrightarrow \overrightarrow{\mathbf{Val}} \\
\mathit{seq}(\mathbf{nil}) & = \cdot \\
\mathit{seq}(\mathbf{cons} \ v_h \ v_t) & = v_h, \mathit{seq}(v_t)
\end{aligned}$$

Fig. 16. Mind the Gap

is incompatible with **BehFun** (figure 12). However, we can use it to create a behavior function. Figure 16 defines $\mathit{behavior}_{\lambda_{ds}^{\cup}}$, which is suitable for λ_{ds}^{\cup} actors. The idea is to store each actor's private state as a pair of λ_{ds}^{\cup} values: one for the λ -term implementing the behavior and one for the actual private state. Invoking eval on a function application term comprised of the behavior function, the incoming event (as a λ_{ds}^{\cup} term⁷), and the current state value yields the actor's response to the event. To reconcile the response, a λ_{ds}^{\cup} value, and a dataspace action, figure 16 also defines $[\]$ (pronounced “liff”), a translation that relates λ_{ds}^{\cup} values to dataspace actions. Because the syntax of events and actions in λ_{ds}^{\cup} is largely the same as that of dataspaces, the translation is mostly straightforward. It ensures that the two sets in a patch are disjoint (n.b. figure 12), and performs cosmetic surgery on actor actions, inserting $\mathit{behavior}_{\lambda_{ds}^{\cup}}$ and storing the function term in the initial state. The final metafunction from the figure, seq , reconciles the \mathbf{cons} lists computed in λ_{ds}^{\cup} with the syntactic sequences manipulated by the dataspace coordination model.

Running λ_{ds}^{\cup} Programs. We finally have all the tools to define the meaning of a λ_{ds}^{\cup} program, dataspace $\tau_c M$, with the metafunction $\mathit{initialize}$ (figure 17).

The function evaluates the given expression, yielding a list of actor descriptions. The seq metafunction translates the \mathbf{cons} list to a syntactic sequence, $(\overrightarrow{\cdot})$. Each actor value in

⁷ Though the syntax of patches is the same in both models, we include $[\]$ (pronounced “lower”) to point out where such a translation would need to happen in an implementation.

$$\begin{aligned}
\text{initialize} & : \mathbf{Init} \longrightarrow \mathbf{Config} + \mathbf{Error} \\
\text{initialize}(\text{dataspace } \tau_c M) & = \begin{cases} \mathbf{error}_{\text{ds}} & \text{if } v = \mathbf{error}_\eta \\ \mathbf{boot}(\text{dataspace } \vec{S}) & \text{otherwise} \end{cases} \\
& \text{where } v = \text{eval}(M) \\
& \quad \vec{v}_a = \text{seq}(v) \\
& \quad \vec{S} = \lceil v_a \rceil
\end{aligned}$$

Fig. 17. λ_{ds}^{\cup} Programs as Dataspaces

the resulting sequence is lifted ($\lceil \cdot \rceil$) to yield an initial actor action; boot-ing the dataspace containing the actor actions yields the program's starting configuration. At this point, reduction proceeds via \longrightarrow_{ds} .

5.3 Properties

With a complete operational description of λ_{ds}^{\cup} dataspace systems in hand, we may now show that the soundness theorem of section 4.4 generalizes to complete programs. Critically, λ_{ds}^{\cup} lives up to the interface imposed on actors by the semantics of section 5.1, and dataspace routing matches the expectations encoded in the type rules of λ_{ds}^{\cup} from section 4.3. In conjunction with a fairness property, the communication structure of a λ_{ds}^{\cup} program matches static expectations.

Rule T-ACTOR (figure 9) models dataspace routing by relating the assertions of interest made by an actor to the events it receives. Hence, we need to relate an individual actor's received events and performed actions across temporally distant reductions.

The relationship crucially relies on the fundamental theorem of dataspace event dispatch: a dataspace applies an actor's behavior function to only those assertions in which the actor has expressed a prior interest.

Theorem 6 (Soundness of Routing (Garnock-Jones, 2017, p. 64)). *If $C_0 \longrightarrow_{ds}^+ C_n$ where C_n is a configuration that is about to dispatch an event to actor ℓ :*

$$[\cdot; R_n; \vec{A}] (\ell \mapsto \vec{e}'_n(\pi^+/\pi^-) \triangleright (f, u) \triangleright \vec{a}'_n) \vec{A}_Q]$$

then there is some $C_i, i < n$, with $C_i = [(k, a); R_i; \vec{A}]$ such that the actor has a stated interest in each delivered assertion:

$$\{(\ell, \text{observe } c) \mid c \in (\pi^+ \cup \pi^-)\} \subseteq R_i$$

□

Next, we show that dataspace routing preserves the type associated with assertion sets.

Lemma 7 (Preservation of Types Across Dataspace Reductions). *The set of assertions π held by a dataspace formed with communication type τ_c have the property $\pi \models \tau_c$ at each reduction step.*

Proof The dataspace operations on assertion sets are set union, intersection, and subtraction, all of which preserve types. □

The validity of T-ACTOR follows.

Lemma 8 (Safe Event Dispatch). *If the dataspace routes a patch π^+/π^- to an actor with a behavior function of type $(\text{Patch } \tau_{in} \tau_{in}, \tau_{state}) \rightarrow (\text{List } (\text{Action } \tau_{out} \tau_c), \tau_{state})$ then $\vdash \pi^+/\pi^- : \text{Patch } \tau_{in} \tau_{in}$.*

Proof Dataspaces compute the intersection of assertions to determine which actors to invoke on which events. The type system accounts for this intersection in rule T-ACTOR, which predicts that events delivered to the actor correspond to a type covered by the domain of the behavior function. Theorem 6 plus lemma 7 jointly verify that the semantics of dataspace routing matches this expectation. \square

The function $behavior_{\lambda_{ds}^U}$ is the main link connecting the two parts of the dataspace model; it relies on correctness lemmas for several helper functions.

Lemma 9 (Typed Lists Translate to Typed Sequences). *If $\vdash v : \text{List } \tau$ then $seq(v) = \vec{v}_i$ and $\vdash v_i : \tau$.*

Proof By routine induction. \square

Lemma 10 (Typed Action Translation). *If $\vdash v : \text{Action } \tau \sigma$ then there exists an a such that $\lceil v \rceil = a$.*

Proof Immediate from $behavior_{\lambda_{ds}^U} \in \mathbf{BehFun}$, which lemma 11 proves. \square

We may now show that typed actors are well-behaved, always yielding a suitable answer in response to a dispatched event.

Lemma 11 (Linking). $behavior_{\lambda_{ds}^U} \in \mathbf{BehFun}$.

Proof First, we show that the term $M_{dispatch} = (v_f([e], v_s))$ is always well-typed. Though we do not formally state this property, it is self-evident that dataspace reductions handle an actor's private state appropriately; the behavior function is always invoked with the most recently returned, or initial, value. Therefore, the second argument of $behavior_{\lambda_{ds}^U}$ is always a pair (v_f, v_s) . Moreover, these pairs originate from a well-typed actor term passed to $\lceil \cdot \rceil$. Also, $behavior_{\lambda_{ds}^U}$ only ever updates the second element of the pair, leaving the function term v_f unchanged. We may then conclude that v_f came from a well-typed actor term, and by inversion of T-ACTOR it has type

$$(\text{Patch } \tau_{in} \tau_{in}, \tau_{state}) \rightarrow (\text{List } (\text{Action } \tau_{out} \tau_c), \tau_{state})$$

Since v_s is either the initial state from the actor term or the second element of a pair returned by v_f , it must have type τ_{state} . By Lemma 8, each delivered patch π^+/π^- , and consequently $[e]$, has type $\text{Patch } \tau_{in} \tau_{in}$. Therefore

$$\vdash M_{dispatch} : (\text{List } (\text{Action } \tau_{out} \tau_c), \tau_{state})$$

Second, by theorem 5, evaluation yields either a value of a suitable type or an error. That is, either $eval(M_{dispatch}) = v$ and $\vdash v : (\text{List } (\text{Action } \tau_{out} \tau_c), \tau_{state})$ or $eval(M_{dispatch}) = \mathbf{error}_{prim}$.

- Case: $eval(M_{dispatch}) = \mathbf{error}_{prim}$. The result of $behavior_{\lambda_{ds}^U}$ is then \mathbf{error}_{ds} . Since $\mathbf{error}_{ds} \in \mathbf{Error}$, $behavior_{\lambda_{ds}^U}$ yields a suitable answer.
- Case: $eval(M_{dispatch}) = v$. By inversion, $v = (v_{list}, v'_s)$ where $\vdash v'_s : \tau_{state}$ and $\vdash v_{list} : \text{List } (\text{Action } \tau_{out} \tau_c)$. By lemmas 9 and 10, the function produces a vector of actions: $\lceil seq(v_{list}) \rceil \in \mathbf{Action}$. \square

Technically, $\text{behavior}_{\lambda_{ds}^{\cup}}$ is *not* a total function matching the signature of **BehFun**; it is only defined on events of the expected type, and private states that are pairs of the expected shape. Figure 12 requires behavior functions be total over the *entire* $\mathbf{Event} \times \mathbf{Val}$ space. However, we have shown that the function is only ever invoked on the portion of the space for which it is defined. Therefore, it meets the actual requirement of the dataspace semantics of always terminating with a suitable answer in response to a dispatched event.

When all individual actors meet the interface, dataspace enjoy a progress property.

Lemma 12 (Progress of Dataspace Configurations (Garnock-Jones, 2017, p. 61)). *Dataspace configurations are either inert or may further reduce.* \square

We can now show that the soundness of λ_{ds}^{\cup} extends to complete dataspace programs: actors fail only due to partial primitives. That is, during communication, errors arise only from the interpretation of messages, not their shape.⁸

Theorem 13 (System Soundness). *If $\vdash \text{dataspace } \tau_c M$, then either*

- $\text{initialize}(\text{dataspace } \tau_c M) = \text{error}_{ds}$; or
- $\text{initialize}(\text{dataspace } \tau_c M) = \Sigma$, where
 - $\Sigma \xrightarrow{ds^*} \Sigma_1$; or
 - for all $\Sigma', \Sigma \xrightarrow{ds^*} \Sigma'$ implies there exists Σ'' such that $\Sigma' \xrightarrow{ds} \Sigma''$

Proof By inversion of $\vdash \text{dataspace } \tau_c M$, it must be the case that

$$\vdash M : \text{List} (\text{Actor } \tau_c)$$

Since λ_{ds}^{\cup} is sound (theorem 5), either

- $\text{eval}(M) = \text{error}_{\text{prim}}$; or
- $\text{eval}(M) = v$ and $\vdash v : \text{List} (\text{Actor } \tau_c)$.

In the first case, $\text{initialize}(\text{dataspace } \tau_c M) = \text{error}_{ds}$; the program fails during startup, an acceptable outcome. Otherwise, evaluation yields a value v with a list type.

By lemma 9, $\text{seq}(v) = \vec{v}_a$ with $\vdash v_a : \text{Actor } \tau_c$. By inversion of the type derivation, each value in the sequence is an actor action, i.e. $v_a \equiv \text{actor } \tau_c v_f v_s \pi$.

The next step translates each λ_{ds}^{\cup} actor specification to a dataspace process, which according to lemma 11 is compatible with the semantics of section 5.1:

$$[\text{actor } \tau_c v_f v_s \pi] = \text{actor } \text{behavior}_{\lambda_{ds}^{\cup}} (v_f, v_s) \pi$$

and boot produces an initial actor state. Finally, lemma 12 shows that such states either reduce to inertness or without end. \square

Finally, λ_{ds}^{\cup} dataspace systems are fair. Programmers may rely on the fact that a ready actor will eventually run.

Theorem 14 (Fairness). *If an actor is non-quiescent, it will execute within a finite number of reductions.*

⁸ Dataspace programs are also deadlock free; section 6 discusses the usefulness of this fact.

Proof The reduction semantics of dataspaces are deterministic (Garnock-Jones, p.62) given the fixed scheduling rule. Coupled with terminating actor behaviors (lemma 11), and a round-robin scheduling policy (figure 14; appendix A.3, definition 29), any non-quiescent actor eventually finds itself in the hole of an evaluation context. \square

6 Assessment

A typical type system offers a number of attractive benefits: a design guide, documentation, and IDE support. Soundness adds reliability in different ways: compatibility checking of specification (type) and implementation (code), error prevention, debugging, and optimization. Thus far, no type system captured tuple-space-style communication precisely and none with soundness proofs. While the typical type system advantages accrue to ours, the error prevention and detection aspect deserves a special assessment.

In terms of error prevention, the type system provably eliminates many standard problems and three kinds of novel faults specific to the dataspace model of actor computation. (1) Types guarantee the absence of mistakes related to the shape of exchanged data, up to the expressivity of the base type system. (2) The type system tracks the structure of infinite assertion sets and statically detects when an actor selects a possibly infinite subset. Thus, the system prevents cases where data produced by a faulty component leads to the crash of a well-behaved one. (3) The presented types eliminate the possibility of actors that diverge while handling state-change notifications, an issue that thus far had to be handled outside the language.

Our type system also permits the imposition of stringent constraints on specific actors. Technically speaking, actor subtyping allows a developer to place individualized constraints on the assertions of each actor in a program. By the soundness property, we know that every assertion an actor makes falls within the range type of its behavior function. The contrapositive of this fact is useful as well: an assertion can never be made by the actor unless explicitly allowed by the range type. Using this ability, we can verify some aspects of well-behaved actors, such as only engaging in conversations fitting their role.

Some aspects of the dataspace communication model are beyond the expressive power of our type system. While typed actors agree on a vocabulary, λ_{ds}^U does not provide any further guarantee about dataspace exchanges. The type system checks only that assertions are safe to utter, but not that they contribute meaningfully to the conversation. Protocols may come with temporal constraints, requirements on the maximum (or minimum) number of related assertions that may exist, and restrictions on data beyond simple types—that an identifier is globally unique, a sequence number is greater than its predecessor, and so on. All of these properties are beyond the power of our structural type system.

Type systems for concurrency often prove deadlock-freedom as a primary result, but the situation in λ_{ds}^U is a bit murkier. Dataspace communication is completely indirect, so in one sense all programs are deadlock-free. However, deadlock-like scenarios may arise. Actor A waits for an assertion x before responding with an assertion y , while actor B waits for an assertion of y before making an assertion x . The type system of λ_{ds}^U is no help in this situation. The types of behavior functions capture only the entire input/output possibilities of an actor. Therefore, they cannot track the dependencies between particular assertions needed to detect the error in this program. In ongoing work, we are experimenting with a

domain-specific language for dataspace actors that lends itself to tracking the relationship between particular types of assertions.

Finally, we have emphasized the importance of termination for actor systems. In reality, an actor that provably finishes computing after four billion years is just as problematic as one that diverges. The truly desired property for actors is *responsiveness*. Termination is a compromise guarantee, albeit one that is particularly useful for language models. Advances in reasoning about worst-case execution bounds (Leivant, 2001; Hoffmann & Shao, 2015) may provide a solution to this aspect of dataspace actors.

Note Prior work on dataspace actors emphasizes the ability to nest collections of actors. Technically, an actor may not only spawn new actors but also new dataspaces by sending a `dataspace` action. This nested dataspace operates as if it were an actor in the current one. Nested dataspaces offer a means of spatially separating actors and conversations. Distinguished constructors route assertions between parent and child dataspaces. We omit this feature from this presentation, as it neither alters the shape of actor interactions nor poses an additional challenge for the design of our type system. Appendix B provides an overview of what is needed to accommodate arbitrary nesting of dataspaces.

7 Implementation and Experiences

Designing a type system is an exercise in trade-offs. Every guarantee places a burden on the language implementer and a restriction on its programmers; each simplification risks losing desired precision. Implementing the proposed system as an extension to the existing Racket prototype allows us to explore these choices on realistic examples within a tight design feedback loop.

In summary, we found that:

1. The type system lends itself to a straightforward implementation.
2. The restrictions are not burdensome. Typed programs may be written in much the same style as their untyped counterparts, modulo the insertion of type annotations.
3. For the small programs in our design feedback loop, pattern matching on incoming assertions suffices to eliminate union types.

Dataspace Types in Action. The examples and snippets from section 3 are written in the concrete syntax of our implementation, which extends the untyped Racket prototype. As the chat server example illustrates, the typed program in figure 5 is largely the same as the untyped one in figure 3. The major difference is that the programmer must write down the communication type, though even the design of untyped programs requires similar, informal documentation of the shape of assertions in the dataspace.

Dataspace Types as Macros. The type checker is defined across a collection of syntax transformers that employ the types-as-macros technique (Chang *et al.*, 2017). In this scheme, the expansion of a typed term produces both an elaboration—untyped syntax implementing the desired behavior—and its type. Transformers obtain the types of sub-terms through recursive traversals (Flatt *et al.*, 2012) and perform some analysis before finally computing their own type and elaboration. Chang’s Turnstile library manages much

of the required bookkeeping, providing a notation for writing rules that resembles the standard conventions of figure 9.

```

1 (define-typed-syntax (actor  $\tau$ -c beh st0 as0)  $\gg$ 
2   #:fail-unless (flat-type? #' $\tau$ -c)
3     "Communication type must be first-order"
4   [⊢ beh  $\gg$  beh-  $\Rightarrow$  ( $\rightarrow$  (Patch  $\tau$ -in  $\tau$ -in)  $\tau$ -st
5     (Instruction  $\tau$ -st  $\tau$ -out  $\tau$ -act))]
6   [⊢ st0  $\gg$  st0-  $\Leftarrow$   $\tau$ -st]
7   [⊢ as0  $\gg$  as0-  $\Leftarrow$  (AssertionSet  $\tau$ -out)]
8   #:fail-unless (<: #' $\tau$ -out #' $\tau$ -c)
9     "Actor makes assertions not allowed in this dataspace"
10  #:fail-unless (<: #'(Actor  $\tau$ -act)
11    #'(Actor  $\tau$ -c))
12    "Spawned actors not allowed in this dataspace"
13  #:fail-unless (<: ( $\cap$  (strip-? #' $\tau$ -out) #' $\tau$ -c) #' $\tau$ -in)
14    "Not prepared to handle all inputs"
15  -----
16  [⊢ (untyped:actor beh- st0- as0-)  $\Rightarrow$  (Actor  $\tau$ -c)])

```

Fig. 18. T-ACTOR as a Turnstile type-and-elaboration rule

Figure 18 provides a representative sample, lightly edited for presentation, from the implementation. The code defines the typed `(actor τ -c beh st0 as0)` form following the skeleton of the T-ACTOR rule from figure 9. The body (lines 2–14) is a sequence of premises, which come in two forms. The first, `#:fail-unless`, performs error checking. On lines 2–3, the rule checks that the given annotation sensibly describes assertions. The `flat-type?` procedure implements the flat judgment. If the check fails, the rule signals an error using the supplied message. The second type of premise inspects the subterms of the given syntax, utilizing Turnstile’s support of bidirectional checking (Pierce & Turner, 1998) to either infer (\Rightarrow) or check (\Leftarrow) types.⁹ The conclusion of the rule (line 16) specifies the elaboration as the untyped actor form applied to the elaboration of the behavior, state, and assertion terms and synthesizes the type of the term. Neither the conclusion nor the premises mention the environment, Γ , which is created and accessed implicitly and hygienically in this framework.

8 Related Work

The dataspace model of coordination and computation seeks to explore a mechanism for replicating knowledge among actors that sits between the complete isolationist approach of message sending on one end and the fully-shared memory one on the other end. In this spirit, Gelernter’s tuple spaces and the recently derived Fact Spaces have the closest resemblance to dataspaces. Hence this section compares the type systems of these latter two approaches with the one we have imposed on dataspace actors here. In short, integration of these models with typed programming languages has never provided a sound model of communication, as our system does.

⁹ The `Instruction` type constructor is a convenience for accommodating both `transition` and `quit` actions.

We also consider approaches for traditional actor systems and other coordination models, as well as session types, a particularly well-known approach to *behavioral* typing. Behavioral type systems typically track explicit communication actions, such as sending a message along a channel. The structure of such communication actions may then be checked for consistency between components, often resulting in strong static guarantees. The underlying communication model of dataspaces, with anonymous components publishing and subscribing (Eugster *et al.*, 2003) to relevant information, lacks such explicit interaction. We therefore consider behavioral types for dataspaces a direction for future work, and so review developments that have some promise in their application to dataspaces.

Fact Spaces and AmbientTalk. The Fact Spaces model (Mostinckx *et al.*, 2007), and especially its prototype implementation CRIME, is similar to the dataspace model. Programs react to both the appearance and disappearance of facts from a shared repository. Reactions are programmed in a logic coordination language, computing new facts based on current facts, and recording (implicitly) the dependencies between facts and application actions.

The Fact Spaces model has been integrated with the AmbientTalk language (Van Cutsem *et al.*, 2014). AmbientTalk is a traditional Actor model (Hewitt *et al.*, 1973; Agha, 1986) language in the mold of E (Miller *et al.*, 2005). In E and AmbientTalk, objects are organized into *vats*; each vat runs its own event loop, dispatching incoming messages to the contained objects and running the corresponding method to completion. In addition to point-to-point messaging, AmbientTalk provides a publish/subscribe communication mechanism via topic (type) tags. Topic tags form a small nominal type system but provide no structural guarantees.

Fact Spaces and AmbientTalk are based on, and always implemented in, untyped languages. Consequently, little can be guaranteed about the behavior of programs ahead of time. Ambient contracts (Scholliers *et al.*, 2011) explore, among other things, protocol enforcement for AmbientTalk programs. In addition, ambient contracts provide functionality related to service discovery and recovering from peer disconnections.

Tuple Spaces. Linda (Gelernter, 1985; Carriero *et al.*, 1994) introduced the tuple space model of coordination. Linda resembles a blackboard style system, where processes deposit messages in the shared space that are later read and removed by other processes. LIME (Murphy *et al.*, 2006) is an extension of the tuple space model where processes register *reactions*, handler functions that run on matching tuples as they appear but, notably, not as they disappear from the space.

Linda implementations with a typed base language tend to use a single `Tuple` type for describing items retrieved from the tuple space (Picco *et al.*, 2005; Jav, 2017; Wyckoff *et al.*, 1998). Such untyped interfaces require casting to a more specific type after read, with the potential to fail due to type mismatches. Gigaspaces (gig, 2017) parameterize the type of tuple space operations such that reading a tuple does not require a cast, but do not associate a type with the entire space. Consequently, there is no assurance that the type is sensible, i.e., a tuple of that type can ever be put in the space. Blossom (van der Goot, 2000) is a tuple space implementation that fixes the type of the entire tuple space, but, being based on C++, is unsound.

Actor and Actor-like Languages. Point-to-point actor languages have been the subject of a number of different type systems and analyses. Though the differences in the underlying communication model makes direct comparisons uninformative, there are some similarities that may allow the results of one to carry-over to the other in a modified form.

The Conversation Calculus (Vieira *et al.*, 2008) shares notable similarities with dataspace in its focus on multi-party interaction between anonymous components. In the conversation calculus, communication takes place within the context of a distinct, potentially nested, conversation name. A process initiates a conversation by instantiating a service name provided by another process. Much like in dataspace, communication within a conversation is anonymous: routing finds both an attempt to send and to receive a message with the same tag within a conversation. Conversation types (Caires & Vieira, 2009) describe the sequence of messages exchanged during each instantiation of a conversation. The type system provides flexibility in how a conversation type decomposes into multiple process types, as well as how multiple process types may merge into a conversation type, while providing the guarantee of soundness and deadlock-freedom. A significant difference that prevents their immediate application to dataspace is the persistent nature of assertions as well as multicast-by-default communication.

The recently developed mailbox calculus and its type system (de'Liguoro & Padovani, 2018) may also have an application to dataspace and vice versa. Though expressive enough to describe different communication mechanisms, the mailbox calculus is tailored particularly to actor-like communication; the capability to receive from a mailbox is unique, while any number of different components may send to a mailbox. Types describe the potential contents of a mailbox with patterns, which take the form of commutative regular expressions. Like assertions, such patterns do not describe the identity of the underlying components, and grant a degree of agnosticism towards the multiplicity of a message. Checking mailbox types relies on computing the pattern describing a mailbox *after* a message has been received, unlike the persistent nature of assertions in a dataspace.

He *et al.* (2014) proposed a design for Typed Akka, a more traditional message-passing actor framework for Scala. Typed Akka actors specify a particular type for the messages that they receive. While we emphasize union types for describing the assertions communicated between dataspace actors, Typed Akka seeks to work within the Scala type system. Consequently, it cannot utilize unions, even though they naturally express the underlying communication.¹⁰

Others have considered static assurance of actor isolation through types (Srinivasan & Mycroft, 2008; Haller & Loiko, 2016; Clebsch *et al.*, 2015) or verification (Summers & Müller, 2016). These approaches focus on isolating mutable heap references among actors, versus partial failure of components discussed here. The dataspace model is functional, with actors communicating strictly first-order values, obviating the need for tracking ownership of mutable heap values.

Session Types. Session type systems (Honda *et al.*, 1998) provide strong guarantees about communication between concurrent components. Theoretical variants are based on the π -

¹⁰ The next version of the language, Scala 3, is scheduled to include union types (Dotty Compiler Team, 2019), opening the possibility of more expressive types for actors.

calculus (Milner, 1999) and describe the bidirectional flow of values between each pair of processes. Multiparty session types (Honda *et al.*, 2008) cover pairwise communication within a delineated group of processes.

Over the past three decades, the π -calculus community has demonstrated that the calculus can encode many communication and coordination mechanisms. In this spirit, it is possible to encode dataspace programs in the channel-based model of the π -calculus. The most natural encoding represents the dataspace itself as a process and then links each actor via channels to this “dataspace process.” This central process holds the table of current assertions and sends notifications to all connected actors of updates to their interests, receiving back a description of the actors’ next actions, and so on. In short, this encoding would mimic the semantics of dataspace (section 5.1), with notify and gather supplanted by analogous channel send and receive operations.

Furthermore, sessions-type researchers have succeeded in designing suitable type systems for many of these π -calculus encodings. Here is a sketch of a session type system for the dataspace encoding. From the perspective of a dataspace process with communication type τ_c , the protocol along each channel would resemble a function call corresponding to our signature for behavior functions:¹¹

$$!(\text{Event } \tau_c).?(List (\text{Action } \tau_{out} \tau_c))$$

Where we use $!\tau$ to mean sending a message of type τ along the channel, $?\tau$ for receive, and $\tau.\tau$ for sequencing.

A close look reveals that a type system for an encoding is un-enlightening for the designers of typed tuple space languages and developers who use such languages. The clearest distinction concerns the routing process itself, especially how it informs the design of the type system. Our type system refines the communication types to mirror the idea that actors receive only messages that they asked for. Technically, the type of events an actor must handle, τ_{in} , is related to τ_c and the type of interests in τ_{out} . Plain session types obscure this connection, because the dataspace process must treat every connected actor in a uniform manner. In other words, the type system of the encoding lacks the modicum of dependency baked into our own type system.

This observation suggests the use of Dependent session types (Toninho *et al.*, 2011). With dependent session types, a developer can prove facts about dataspace routing on a per-actor basis. We consider this approach less useful than our proposed type system, even if it is equally expressive in the context of a π -calculus encoding. Most critically, the type system can no longer guide the working programmer to a natural type. Instead, this burden is shifted to the programmer, who must articulate these relationships as dependent types on a per-actor basis. We consider this an excessively large burden on the aspiring dataspace programmer, because the programmer would have to mentally shift back and forth between the encoding and the direct design goal. By comparison, our type system directly incorporates an abstraction of the underlying routing mechanism—embodied in the *predict-routing* metafunction—and thus expresses type constraints in the linguistic domain in which the developer actually programs.

¹¹ In this encoding, each actor process would maintain its private state value to obviate the need for an existential quantifier.

Nevertheless we conjecture that translating dataspace into sessions and channels may provide ideas for expressing and checking dataspace protocols. A particularly challenging aspect will be the broadcast-like, time-enduring nature of dataspace assertions. We expect that the recent research on session types for messaging actors (Mostrous & Vasconcelos, 2011; Crafa, 2012) might be a good starting point for this line of investigation.

9 Conclusion

No single coordination model can perfectly express all patterns of communication. The tuple space family of models emphasizes the need for sharing information among a group of actors, a need that is often overlooked by point-to-point models, even though the latter might be able to encode the former with patterns. The dataspace model seeks to revive interest in addressing this need for sharing, and as we have argued (section 2.2) offers a number of advantages for programming group conversations.

Tuple space languages have never featured a type system that captures the structure of exchanged data, however. As a consequence, programmers lose the advantages that types play in design, error prevention, documentation, optimization, and so on. Moreover, a lack of types governing communication can have an especially pernicious effect on actor-like systems. Without types, a simple data-format error in a message can cause the crash of an otherwise well-behaved actor, undermining the principle of fault-isolation.

This paper presents a structural type system for dataspace actors, the first such in the tuple space family. The new type system accounts for the novel communication medium and completely eliminates data-format errors from the model. Practically speaking, this step simplifies the task of actor fault detection and recovery.

Experience with a prototype implementation of the type system suggests that it supports a range of protocols and lends itself to practical use. To cope with complex actors, the type system might need additional tools for programming with unions, such as occurrence typing, and reasoning based on type-state (Strom & Yemini, 1986). Looking beyond the internal structure of actors, dataspace protocols stand to benefit from the checking of behavioral properties as well as structural ones.

Given the similarities between dataspace and Fact Spaces, it should be straightforward to adapt the type system to the Fact Space model and its implementations. Furthermore, the type system, by construction, has a clearly defined, narrow interface to the dataspace mechanism (section 5). Hence, we should also be able to adapt our type system for dataspace to similar coordination systems for actors, such as tuple spaces and pub/sub layers.

Acknowledgments The authors thank Stephen Chang (Northeastern, UMass Boston) for his Turnstile package and comments on early drafts of this paper. The anonymous reviewers helped clarify many passages of the original submission. The research was partially supported by CISCO and NSF grants SHF 1763922 and 1518844.

References

(2017). *Gigaspace*. <https://www.gigaspace.com>. Accessed: 2017-10-19.

- (2017). *JavaspaceTM service specification, version 2.3*.
<https://river.apache.org/release-doc/current/specs/html/js-spec.html>. Accessed: 2017-01-24.
- Agha, Gul. (1986). *Actors: A Model of Concurrent Computation in Distributed Systems*. Massachusetts: MIT Press.
- Armstrong, Joe. (1994). *Programming Erlang*. The Pragmatic Programmers.
- Armstrong, Joe. 2003 (Dec). *Making Reliable Distributed Systems in the Presence of Software Errors*. PhD dissertation, Royal Institute of Technology, Stockholm.
- Caires, Luís, & Vieira, Hugo Torres. (2009). Conversation types. *Pages 285–300 of: ESOP*.
- Carriero, Nicholas J, Gelernter, David, Mattson, Timothy G, & Sherman, Andrew H. (1994). The Linda alternative to message-passing systems. *Parallel computing*, **20**(4), 633–655.
- Chang, Stephen, Knauth, Alex, & Greenman, Ben. (2017). Type systems as macros. *Pages 694–705 of: POPL*.
- Clebsch, Sylvan, Drossopoulou, Sophia, Blessing, Sebastian, & McNeil, Andy. (2015). Deny capabilities for safe, fast actors. *International Workshop on Programming Based on Actors, Agents, and Decentralized Control*. AGERE! 2015.
- Clocksin, W.F., & Mellish, C.S. (1981). *Programming in prolog*. Springer.
- Crafa, Silvia. (2012). Behavioural types for actor systems. *Technical report*.
- de'Liguoro, Ugo, & Padovani, Luca. (2018). Mailbox types for unordered interactions. *ECOOP*.
- Dotty Compiler Team. (2019). *Union types - more details*.
<https://dotty.epfl.ch/docs/reference/new-types/union-types-spec.html>. Accessed: 2019-12-13.
- ECMA. (2015). *ECMA-262: ECMAScript 2015 language specification*. Sixth edn. ECMA International.
- Englemore, Robert, & Morgan, Anthony (Eds). (1988). *Blackboard systems*. Addison-Wesley.
- Eugster, Patrick Th., Felber, Pascal A., Guerraoui, Rachid, & Kermarrec, Anne-Marie. (2003). The many faces of publish/subscribe. *ACM Computing Surveys*, **35**(2), 114–131.
- Felleisen, Matthias, Findler, Robert Bruce, & Flatt, Matthew. (2009). *Semantics engineering with plt redex*. MIT Press.
- Flatt, Matthew, & PLT. (2010). *Reference: Racket*. Tech. rept. PLT-TR-2010-1. PLT Inc. <http://racket-lang.org/tr1/>.
- Flatt, Matthew, Culpepper, Ryan, Darais, David, & Findler, Robert Bruce. (2012). Macros that work together: Compile-time bindings, partial expansion, and definition contexts. *Journal of Functional Programming*, **22**(2), 181–216.
- Garnock-Jones, Tony. 2017 (Dec). *Conversational Concurrency*. PhD dissertation, Northeastern University.
- Garnock-Jones, Tony, & Felleisen, Matthias. (2016). Coordinated concurrent programming in Syndicate. *Pages 310–336 of: ESOP*.
- Garnock-Jones, Tony, Tobin-Hochstadt, Sam, & Felleisen, Matthias. (2014). The network as a language construct. *Pages 473–492 of: ESOP*.
- Gelernter, David. (1985). Generative communication in Linda. *ACM trans. program. lang. syst.*
- Girard, Jean-Yves. (1971). Une extension de l'interprétation de Gödel à l'analyse, et son application à l'élimination des coupures dans l'analyse et la théorie des types. *Studies in logic and the foundations of mathematics*, **63**, 63–92.
- Haller, Philipp, & Loiko, Alex. (2016). LaCasa: Lightweight affinity and object capabilities in Scala. *Pages 272–291 of: OOPSLA*.
- He, Jiansen, Wadler, Philip, & Trinder, Philip. (2014). Typecasting actors: from Akka to Takka. *Pages 23–33 of: Proceedings of the Fifth Annual Scala Workshop*.

- Hewitt, Carl, Bishop, Peter Boehler, Greif, Irene, Smith, Brian Cantwell, Matson, Todd, & Steiger, Richard. (1973). Actor induction and meta-evaluation. *POPL*.
- Hoffmann, Jan, & Shao, Zhong. (2015). Automatic static cost analysis for parallel programs. *Pages 132–157 of: ESOP*.
- Honda, Kohei, Vasconcelos, Vasco T., & Kubo, Makoto. (1998). Language primitives and type discipline for structured communication-based programming. *Pages 122–138 of: ESOP*.
- Honda, Kohei, Yoshida, Nobuko, & Carbone, Marco. (2008). Multiparty asynchronous session types. *Pages 273–284 of: POPL '08*.
- Jeurig, Johan. (1995). Polytypic pattern matching. *Pages 238–248 of: FPCA '95*.
- Leivant, Daniel. (2001). Termination proofs and complexity certification. *Pages 183–200 of: Theoretical Aspects of Computer Software*.
- Miller, Mark S., Tribble, E. Dean, & Shapiro, Jonathan. (2005). Concurrency among strangers. *Pages 195–229 of: Int. Symp. on Trustworthy Global Computing*.
- Milner, Robin. (1999). *Communicating and Mobile Systems: the pi Calculus*. Cambridge University Press.
- Mostinckx, Stijn, Scholliers, Christophe, Philips, Eline, Herzeel, Charlotte, & De Meuter, Wolfgang. (2007). Fact spaces: Coordination in the face of disconnection. *Pages 268–285 of: Proc. COORDINATION 2007*.
- Mostrous, Dimitris, & Vasconcelos, Vasco T. (2011). Session typing for a featherweight erlang. *Pages 95–109 of: COORDINATION*.
- Murphy, Amy L., Picco, Gian Pietro, & Roman, Gruia-Catalin. (2006). Lime: A coordination model and middleware supporting mobility of hosts and agents. *Acm Trans. on Software Engineering and Methodology*, **15**(3), 279–328.
- Newell, Allen, & Simon, Herbert A. (1972). *Human problem solving*. Prentice Hall.
- Picco, Gian Pietro, Balzarotti, Davide, & Costa, Paolo. (2005). LighTS: A lightweight, customizable tuple space supporting context-aware applications. *Pages 413–419 of: Proc. SAC '05*.
- Pierce, Benjamin C. (1991). *Programming with intersection types, union types, and polymorphism*. Tech. rept. Technical Report CMU-CS-91-106, Carnegie Mellon University.
- Pierce, Benjamin C., & Turner, David N. (1998). Local type inference. *Pages 252–265 of: POPL*.
- Scholliers, Christophe, Harnie, Dries, Tanter, Eric, De Meuter, Wolfgang, & D'Hondt, Theo. (2011). Ambient contracts: verifying and enforcing ambient object compositions à la carte. *Personal and ubiquitous computing*, **15**(4), 341–351.
- Srinivasan, Sriram, & Mycroft, Alan. (2008). Kilim: Isolation-typed actors for Java. *Pages 104–128 of: ECOOP*.
- Strom, Robert E, & Yemini, Shaula. (1986). Tpestate: A programming language concept for enhancing software reliability. *IEEE Transactions on Software Engineering*, 157–171.
- Summers, Alexander J., & Müller, Peter. (2016). Actor services: Modular verification of message passing programs. *ESOP*.
- Tasharofi, Samira, Dinges, Peter, & Johnson, Ralph E. (2013). Why do scala developers mix the actor model with other concurrency models? *ECOOP*.
- Tobin-Hochstadt, Sam, & Felleisen, Matthias. (2008). The design and implementation of Typed Scheme. *Pages 395–406 of: POPL*.
- Toninho, Bernardo, Caires, Luís, & Pfenning, Frank. (2011). Dependent session types via intuitionistic linear type theory. *PPDP*.
- Van Cutsem, Tom, Gonzalez Boix, Elisa, Scholliers, Christophe, Lombide Carreton, Andoni, Harnie, Dries, Pinte, Kevin, & De Meuter, Wolfgang. (2014). AmbientTalk: programming responsive mobile peer-to-peer applications with actors. *Computer Languages, Systems & Structures*, **40**(3–4), 112–136.

- van der Goot, Roel. (2000). *High performance linda using a class library*. PhD dissertation, Erasmus University Rotterdam.
- Vieira, Hugo T., Caires, Luís, & Seco, João C. (2008). The conversation calculus: A model of service-oriented computation. *Pages 269–283 of: ESOP*.
- Wright, Andrew K., & Felleisen, Matthias. (1994). A syntactic approach to type soundness. *Information and Computation*, **115**(1), 38–94.
- Wyckoff, P., McLaughry, S. W., Lehman, T. J., & Ford, D. A. (1998). T Spaces. *IBM Syst. J.*, **37**(3), 454–474.

A Auxiliary Metafunction Definitions

A.1 λ_{ds} Reduction Metafunctions

The reduction semantics of λ_{ds} (figure 7) mentions several metafunctions that create and analyze assertion sets. This section collects their formal definitions.

Definition 15. The *make-set* metafunction creates an assertion set from a vector of values:

$$\begin{aligned} \text{make-set} & : \mathbf{Val} \longrightarrow_{\text{partial}} \mathbf{ASet} \\ \text{make-set}(\vec{v}) & = \bigcup \vec{\pi} \quad \text{where } \vec{\pi} = \overrightarrow{\text{interp}(v)} \end{aligned}$$

Definition 16. The *interp* function maps λ_{ds} values to sets of assertions:

$$\begin{aligned} \text{interp} & : \mathbf{Val} \longrightarrow_{\text{partial}} \mathbf{ASet} \\ \text{interp}(\star) & = \mathbf{Assertion} \\ \text{interp}(b) & = \{b\} \\ \text{interp}(m(\vec{v})) & = \{m(\vec{v}') \mid (\vec{v}') \in \text{interp}(\vec{v})\} \\ \text{interp}(\cdot) & = \{\cdot\} \\ \text{interp}((v, \vec{v}')) & = \{(x, \vec{y}') \mid x \in \text{interp}(v), (\vec{y}') \in \text{interp}(\vec{v}')\} \\ \text{interp}(\text{cons } v_1 \ v_2) & = \{\text{cons } x \ y \mid x \in \text{interp}(v_1), y \in \text{interp}(v_2)\} \\ \text{interp}(\text{observe } v) & = \{\text{observe } x \mid x \in \text{interp}(v)\} \end{aligned}$$

Definition 17. The *project* function analyzes assertion sets with a pattern:

$$\begin{aligned} \text{project} & : \mathbf{ASet} \times \mathbf{Val} \times \mathbf{Expr} \longrightarrow_{\text{partial}} \mathbf{Expr} \\ \text{project}(\pi, v_p, M) & = \text{unroll}(m) \quad \text{if } m \text{ is finite} \\ & \quad \text{where } m = \{\gamma(M) \mid v \in \pi, \text{match}(v, v_p) = \gamma\} \end{aligned}$$

Successful pattern matches yield substitutions:

$$\text{Substitutions } \gamma \in \mathbf{Sub} = \mathbf{Var} \longrightarrow_{\text{partial}} \mathbf{Val}$$

where composition $\gamma_1 \circ \gamma_2$ is defined in the usual manner.

Pattern matching is defined in straightforward fashion:

$$\begin{aligned} \text{match} & : \mathbf{Val} \times \mathbf{Val} \longrightarrow_{\text{partial}} \mathbf{Sub} \\ \text{match}(v, \$x : \tau) & = \{(x, v)\} \\ \text{match}(v, _) & = \{\} \\ \text{match}(b, b) & = \{\} \\ \text{match}(\text{observe } v, \text{observe } v_p) & = \text{match}(v, v_p) \\ \text{match}(m(\vec{v}_n), m(\vec{v}_{pn})) & = \text{match}((\vec{v}_n), (\vec{v}_{pn})) \\ \text{match}((v, \vec{v}_n), (v_p, \vec{v}_{pn})) & = \text{match}(v, v_p) \circ \gamma \quad \text{where } \text{match}((\vec{v}_n), (\vec{v}_{pn})) = \gamma \\ \text{match}(\text{cons } v_1 \ v_2, \text{cons } v_{p1} \ v_{p2}) & = \text{match}(v_2, v_{p2}) \circ \gamma \quad \text{where } \text{match}(v_1, v_{p1}) = \gamma \end{aligned}$$

The *unroll* function translates the set of results to a list expression:

$$\begin{aligned} \text{unroll} & : \mathcal{P}(\mathbf{Expr}) \longrightarrow_{\text{partial}} \mathbf{Expr} \\ \text{unroll}(\emptyset) & = \text{nil} \\ \text{unroll}(\{M\} \uplus S) & = \text{cons } M \ \text{unroll}(S) \end{aligned}$$

While the unrolling operation does not specify the order of elements, we assume a fixed ordering for the selection of elements to form the given set to make the definition deterministic.

A.2 λ_{ds}^U Complete Type Judgment

The following definitions complete the type judgment for λ_{ds}^U terms (section 4.3), providing the rules elided from figures 9 and 11.

Definition 18. The following inference rules specify the complete type judgment of λ_{ds}^{\cup} terms, $\Gamma \vdash M : \tau$, as well as the judgment $\vdash I$ for complete dataspace programs.

$$\frac{\text{flat}(\tau_c) \quad \vdash M_{boot} : \text{List}(\text{Actor } \tau_c)}{\vdash \text{dataspace } \tau_c M_{boot}} \boxed{\text{T-DATASPACE}}$$

$$\frac{\Gamma \vdash x : \tau \quad \tau <: \sigma}{\Gamma \vdash x : \tau} \boxed{\text{T-SUB}}$$

$$\frac{\Gamma \vdash M_{behavior} : (\text{Patch } \tau_{in} \tau_{in}, \tau_{state}) \rightarrow (\text{List}(\text{Action } \tau_{out} \tau_c), \tau_{state}) \quad \Gamma \vdash M_{state} : \tau_{state} \quad \Gamma \vdash M_{assertions} : \text{AssertionSet } \tau_{out} \quad \tau_{out} <: \tau_c \quad \text{predict-routing}(\tau_{out}, \tau_c) <: \tau_{in} \quad \text{flat}(\tau_c)}{\Gamma \vdash \text{actor } \tau_c M_{behavior} M_{state} M_{assertions} : \text{Actor } \tau_c} \boxed{\text{T-ACTOR}}$$

$$\frac{\Gamma \vdash M_s : \text{AssertionSet } \tau_s \quad \Gamma \vdash \text{PAT} : \tau_p \quad \text{safe}(\tau_s, \tau_p) \quad \text{bindings}(\text{PAT}) = \Gamma' \quad \Gamma, \Gamma' \vdash M_b : \tau_b}{\Gamma \vdash \text{project } M_s \text{ with PAT in } M_b : \text{List } \tau_b} \boxed{\text{T-PROJECT}}$$

$$\frac{\overline{\Gamma \vdash \text{SK} : \tau}}{\Gamma \vdash \{\overline{\text{SK}}\} : \text{AssertionSet } \bigcup \overline{\tau}} \boxed{\text{T-SET}} \quad \frac{\pi \models \tau}{\Gamma \vdash \pi : \text{AssertionSet } \tau} \boxed{\text{T-}\pi}$$

$$\frac{}{\Gamma \vdash \text{error}_\eta : \tau} \boxed{\text{T-ERROR}} \quad \frac{\Gamma, x : \tau \vdash M : \sigma}{\Gamma \vdash (\lambda x : \tau. M) : \tau \rightarrow \sigma} \boxed{\text{T-FUN}}$$

$$\frac{\Gamma \vdash M_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \vdash M_2 : \tau_1}{\Gamma \vdash M_1 M_2 : \tau_2} \boxed{\text{T-APP}} \quad \frac{\Gamma(x) = \tau}{\Gamma \vdash x : \tau} \boxed{\text{T-VAR}}$$

$$\frac{\overline{\Gamma \vdash M : \tau}}{\Gamma \vdash m(\overline{M}) : m(\overline{\tau})} \boxed{\text{T-MSG}} \quad \frac{\overline{\Gamma \vdash M : \tau}}{\Gamma \vdash (\overline{M}) : (\overline{\tau})} \boxed{\text{T-TUPLE}}$$

$$\frac{\Gamma \vdash M_1 : \text{List } \tau \quad \Gamma \vdash M_2 : \sigma}{\Gamma \vdash \text{cons } M_1 M_2 : \text{List}(\tau \cup \sigma)} \boxed{\text{T-CONS}}$$

$$\frac{\Gamma \vdash M_c : \tau_a \rightarrow \tau_l \rightarrow \tau_a \quad \Gamma \vdash M_n : \tau_a \quad \Gamma \vdash M_l : \text{List } \tau_l}{\Gamma \vdash \text{fold } M_c M_n M_l : \tau_a} \boxed{\text{T-FOLD}}$$

$$\frac{\Gamma \vdash M : \tau}{\Gamma \vdash \text{observe } M : \text{Observe } \tau} \boxed{\text{T-OBSERVE}} \quad \frac{}{\Gamma \vdash b : \mathbf{B}(b)} \boxed{\text{T-BASE}}$$

$$\frac{\overline{\Gamma \vdash M : \tau} \quad \Delta(p, \overline{\tau}) = \sigma}{\Gamma \vdash p \overline{M} : \sigma} \boxed{\text{T-PRIM}}$$

The judgment employs several additional metafunctions.

Definition 19. The \mathbf{B} and Δ metafunctions assign sound types to primitive values and operations:

$$\begin{aligned} \mathbf{B} &: \mathbf{BasicVal} \longrightarrow \mathbf{BaseTy} \\ \Delta &: \mathbf{Prim} \times \vec{\tau} \longrightarrow_{\text{partial}} \tau \end{aligned}$$

Definition 20. The *bindings* function extracts the type annotations from a pattern; repeated occurrences of the same identifier result in shadowing:

$$\begin{aligned} \text{bindings} &: \mathbf{Pat} \longrightarrow \mathbf{Env} \\ \text{bindings}(\$x : \tau) &= x : \tau \\ \text{bindings}(m(\overrightarrow{\text{PAT}})) &= \text{bindings}(\overrightarrow{\text{PAT}}) \\ \text{bindings}((\text{PAT}, \overrightarrow{\text{PAT}}_n)) &= \text{bindings}(\text{PAT}_1), \text{bindings}(\overrightarrow{\text{PAT}}_n) \\ \text{bindings}(\text{observe PAT}) &= \text{bindings}(\text{PAT}) \\ \text{bindings}(\text{inbound PAT}) &= \text{bindings}(\text{PAT}) \\ \text{bindings}(\text{outbound PAT}) &= \text{bindings}(\text{PAT}) \\ \text{bindings}(_) &= \cdot \quad \text{otherwise} \end{aligned}$$

The type rules also employ several auxiliary judgments.

Definition 21. The judgment $\Gamma \vdash_P \text{PAT} : \tau$ checks patterns separately, which allows us to limit the expressions that may be used in patterns:

$$\begin{array}{c} \frac{\text{flat}(\tau)}{\Gamma \vdash_P (\$x : \tau) : (\$: \tau)} \boxed{\text{P-CAPTURE}} \qquad \frac{}{\Gamma \vdash_P _ : \text{Discard}} \boxed{\text{P-DISCARD}} \\ \\ \frac{\Gamma \vdash M : \tau \quad \text{flat}(\tau)}{\Gamma \vdash_P M : \tau} \boxed{\text{P-EXP}} \qquad \frac{\overrightarrow{\Gamma \vdash_P \text{PAT} : \tau}}{\Gamma \vdash_P (\overrightarrow{\text{PAT}}) : (\overrightarrow{\tau})} \boxed{\text{P-TUPLE}} \\ \\ \frac{\overrightarrow{\Gamma \vdash_P \text{PAT} : \tau}}{\Gamma \vdash_P m(\overrightarrow{\text{PAT}}) : m(\overrightarrow{\tau})} \boxed{\text{P-MSG}} \qquad \frac{\Gamma \vdash_P \text{PAT} : \tau}{\Gamma \vdash_P \text{observe PAT} : \text{Observe } \tau} \boxed{\text{P-SUB}} \end{array}$$

Definition 22. Similarly, the judgment $\Gamma \vdash_{\text{SK}} \text{SK} : \tau$ checks assertion-set creation:

$$\begin{array}{c} \frac{}{\Gamma \vdash_{\text{SK}} * : *} \boxed{\text{SK-STAR}} \qquad \frac{\Gamma \vdash M : \tau \quad \text{flat}(\tau)}{\Gamma \vdash_{\text{SK}} M : \tau} \boxed{\text{SK-EXP}} \\ \\ \frac{\overrightarrow{\Gamma \vdash_{\text{SK}} \text{SK} : \tau}}{\Gamma \vdash_{\text{SK}} (\overrightarrow{\text{SK}}) : (\overrightarrow{\tau})} \boxed{\text{SK-PROD}} \qquad \frac{\overrightarrow{\Gamma \vdash_{\text{SK}} \text{SK} : \tau}}{\Gamma \vdash_{\text{SK}} m(\overrightarrow{\text{SK}}) : m(\overrightarrow{\tau})} \boxed{\text{SK-MSG}} \\ \\ \frac{\Gamma \vdash_{\text{SK}} \text{SK} : \tau}{\Gamma \vdash_{\text{SK}} \text{observe SK} : \text{Observe } \tau} \boxed{\text{SK-SUB}} \end{array}$$

Definition 23. The judgment $\text{flat}(\tau)$ holds for types that correspond to legal assertions:

$$\begin{array}{c}
\frac{}{\text{flat}(B)} \boxed{\text{F-BASE}} \qquad \frac{\overrightarrow{\text{flat}(\tau)}}{\text{flat}(\bigcup \vec{\tau})} \boxed{\text{F-UNION}} \qquad \frac{\overrightarrow{\text{flat}(\tau)}}{\text{flat}((\vec{\tau}))} \boxed{\text{F-PROD}} \\
\frac{\overrightarrow{\text{flat}(\tau)}}{\text{flat}(m(\vec{\tau}))} \boxed{\text{F-MSG}} \qquad \frac{\text{flat}(\tau)}{\text{flat}(\text{List } \tau)} \boxed{\text{F-LIST}} \\
\frac{\text{flat}(\tau)}{\text{flat}(\text{Observe } \tau)} \boxed{\text{F-OBSERVE}} \qquad \frac{}{\text{flat}(\star)} \boxed{\text{F-}\star}
\end{array}$$

Definition 24. The judgment $\text{finite}(\tau)$ describes assertion types that do not contain any uses of \star :

$$\begin{array}{c}
\frac{}{\text{finite}(B)} \boxed{\text{FIN-BASE}} \qquad \frac{\overrightarrow{\text{finite}(\tau)}}{\text{finite}(\bigcup \vec{\tau})} \boxed{\text{FIN-UNION}} \\
\frac{\overrightarrow{\text{finite}(\tau)}}{\text{finite}((\vec{\tau}))} \boxed{\text{FIN-TUPLE}} \qquad \frac{\overrightarrow{\text{finite}(\tau)}}{\text{finite}(m(\vec{\tau}))} \boxed{\text{FIN-MSG}} \\
\frac{\text{finite}(\tau)}{\text{finite}(\text{List } \tau)} \boxed{\text{FIN-LIST}} \qquad \frac{\text{finite}(\tau)}{\text{finite}(\text{Observe } \tau)} \boxed{\text{FIN-SUB}}
\end{array}$$

Definition 25. Rule T-PROJECT employs the $\text{safe}(\tau, \sigma)$ judgment to determine if projecting a pattern of type σ against a set of τ -typed assertions could yield an infinite or ill-typed result:

$$\begin{array}{c}
\frac{\text{finite}(\tau) \quad \tau <: \sigma}{\text{safe}(\tau, \$: \sigma)} \boxed{\text{PS-CAPTURE}} \qquad \frac{}{\text{safe}(\tau, \text{Discard})} \boxed{\text{PS-DISCARD}} \\
\\
\frac{}{\text{safe}(\tau, B)} \boxed{\text{PS-BASE}} \qquad \frac{\overrightarrow{\text{safe}(\tau, \sigma)}}{\text{safe}(\tau, \bigcup \vec{\sigma})} \boxed{\text{PS-UNIONR}} \\
\\
\frac{\overrightarrow{\text{safe}(\tau, \sigma)}}{\text{safe}(\bigcup \vec{\tau}, \sigma)} \boxed{\text{PS-UNIONL}} \qquad \frac{\text{disjoint}(\tau, \sigma)}{\text{safe}(\tau, \sigma)} \boxed{\text{PS-DISJOINT}} \\
\\
\frac{\overrightarrow{\text{safe}(\tau_i, \sigma_i)}}{\text{safe}((\vec{\tau}), (\vec{\sigma}))} \boxed{\text{PS-TUPLE}} \qquad \frac{\overrightarrow{\text{safe}(\star, \tau_i)}}{\text{safe}(\star, (\vec{\tau}))} \boxed{\text{PS-TUPLE}\star} \\
\\
\frac{\overrightarrow{\text{safe}(\tau_i, \sigma_i)}}{\text{safe}(m(\vec{\tau}), m(\vec{\sigma}))} \boxed{\text{PS-MSG}} \qquad \frac{\overrightarrow{\text{safe}(\star, \tau_i)}}{\text{safe}(\star, m(\vec{\tau}))} \boxed{\text{PS-MSG}\star} \\
\\
\frac{\text{safe}(\tau, \sigma)}{\text{safe}(\text{List } \tau, \text{List } \sigma)} \boxed{\text{PS-LIST}} \qquad \frac{\text{safe}(\star, \tau)}{\text{safe}(\star, \text{List } \tau)} \boxed{\text{PS-LIST}\star} \\
\\
\frac{\text{safe}(\tau, \sigma)}{\text{safe}(\text{Observe } \tau, \text{Observe } \sigma)} \boxed{\text{PS-SUB}} \qquad \frac{\text{safe}(\star, \tau)}{\text{safe}(\star, \text{Observe } \tau)} \boxed{\text{PS-SUB}\star}
\end{array}$$

where

$$\text{disjoint}(\tau, \sigma) = \tau \tilde{\cap} \sigma <: \perp$$

Definition 26. Assertion sets π are given types through the judgment $\pi \models \tau$:

$$\begin{array}{c}
\frac{\pi \subseteq \mathbb{Z} \quad \exists n. |\pi| = n}{\pi \models \text{Int}} \boxed{\text{M-NUMBER}} \qquad \frac{\pi \subseteq \text{Assertion}}{\pi \models \star} \boxed{\text{M-}\star} \\
\\
\frac{\pi \subseteq \{\overrightarrow{x_i} \mid x_i \in \overrightarrow{\pi_i}\} \quad \overrightarrow{\pi_i} \models \overrightarrow{\tau_i}}{\pi \models (\overrightarrow{\tau})} \boxed{\text{M-PROD}} \\
\\
\frac{\pi \subseteq \{m(\overrightarrow{x_i}) \mid x_i \in \overrightarrow{\pi_i}\} \quad \overrightarrow{\pi_i} \models \overrightarrow{\tau_i}}{\pi \models m(\overrightarrow{\tau})} \boxed{\text{M-MSG}} \\
\\
\frac{\pi \subseteq \{\text{observe } v \mid v \in \pi'\} \quad \pi' \models \tau}{\pi \models \text{Observe } \tau} \boxed{\text{M-OBSERVE}} \\
\\
\frac{\pi \subseteq \{\text{nil}\} \cup \{\text{cons } x y \mid x \in \pi_1, y \in \pi_2\} \quad \pi_1 \models \tau \quad \pi_2 \models \text{List } \tau}{\pi \models \text{List } \tau} \boxed{\text{M-LIST}} \\
\\
\frac{\pi = \bigcup \pi_i \quad \overrightarrow{\pi_i} \models \overrightarrow{\tau_i}}{\pi \models \bigcup \overrightarrow{\tau}} \boxed{\text{M-UNION}}
\end{array}$$

In rule M-UNION, the premise $\pi = \bigcup \pi_i$ refers to (semantic) set union, while the conclusion $\bigcup \overrightarrow{\tau}$ uses (syntactic) type union. Furthermore, the rule does not require that the π_i s are disjoint or non-empty.

Definition 27. The subtyping judgment $\tau <: \sigma$ relates compatible types:

$$\begin{array}{c}
\frac{}{\tau <: \tau} \boxed{\text{S-REFL}} \qquad \frac{\tau_1 <: \tau_2 \quad \tau_2 <: \tau_3}{\tau_1 <: \tau_3} \boxed{\text{S-TRANS}} \qquad \frac{\text{flat}(\tau)}{\tau <: \star} \boxed{\text{S-}\star} \\
\\
\frac{\tau_2 <: \sigma_2 \quad \sigma_1 <: \tau_1}{\tau_1 \rightarrow \tau_2 <: \sigma_1 \rightarrow \sigma_2} \boxed{\text{S-FUN}} \qquad \frac{\overrightarrow{\tau} <: \overrightarrow{\sigma}}{\bigcup \overrightarrow{\tau} <: \sigma} \boxed{\text{S-UNIONSUB}} \\
\\
\frac{\exists i. \tau <: \sigma_i}{\tau <: \bigcup \overrightarrow{\sigma}} \boxed{\text{S-UNIONSUPER}} \qquad \frac{\tau <: \sigma \quad \text{predict-routing}(\tau, \sigma) <: \tau}{\text{Actor } \tau <: \text{Actor } \sigma} \boxed{\text{S-ACTOR}} \\
\\
\frac{\tau <: \sigma}{\text{List } \tau <: \text{List } \sigma} \boxed{\text{S-LIST}} \qquad \frac{\tau <: \sigma}{\text{AssertionSet } \tau <: \text{AssertionSet } \sigma} \boxed{\text{S-SET}} \\
\\
\frac{\tau <: \sigma}{\text{Observe } \tau <: \text{Observe } \sigma} \boxed{\text{S-SUB}} \qquad \frac{\overrightarrow{\tau_i} <: \overrightarrow{\sigma_i}}{(\overrightarrow{\tau}) <: (\overrightarrow{\sigma})} \boxed{\text{S-PROD}} \\
\\
\frac{\overrightarrow{\tau_i} <: \overrightarrow{\sigma_i}}{m(\overrightarrow{\tau}) <: m(\overrightarrow{\sigma})} \boxed{\text{S-MSG}}
\end{array}$$

A.3 Dataspace Coordination Metafunctions

The reduction semantics of dataspace (section. 5.1) employs several metafunctions.

Definition 28. The \oplus operation incorporates a patch into a dataspace's stored knowledge:

$$\begin{aligned} \oplus & : \mathbf{DS} \times (\mathit{Lift}(\mathbf{Loc}) \times \mathbf{Patch}) \longrightarrow \mathbf{DS} \\ R \oplus (k, \pi_{add}/\pi_{del}) & = R \cup \{(k, c) \mid c \in \pi_{add}\} - \{(k, c) \mid c \in \pi_{del}\} \end{aligned}$$

Definition 29. The *round-robin* metafunction implements a scheduling policy, rotating the actors in a configuration:

$$\begin{aligned} \mathit{round-robin} & : E^\Sigma \times \mathbf{AState} \longrightarrow \mathbf{Config} \\ \mathit{round-robin}([\cdot; R; \vec{A}_I(\ell \mapsto \square)\vec{A}_Q], \Sigma) & = [\cdot; R; \vec{A}_I\vec{A}_Q(\ell \mapsto \Sigma)] \end{aligned}$$

B Extensions for Nested Dataspaces

As mentioned in section 6, dataspace can be nested, forming a tree structure of actors. To keep the model simple, our presentation removes this ability. However, because nested dataspace are a useful feature, and they are present in both our typed and untyped prototypes, they warrant additional discussion.

Here, we give an outline of how to extend the type system to accommodate hierarchical dataspace. The operational semantics of nested dataspace systems requires non-trivial machinery, but the semantics of an internal actor language like λ_{ds}^U is largely the same. The extensions to the type system are likewise straightforward. We refer the interested reader to Garnock-Jones et. al.'s work (2016; 2017) for the operational semantics of nested dataspace systems.

Actors communicate with one another across dataspace using special assertion constructors. The assertion *outbound* c is pertinent to the parent of the current dataspace; routing recognizes such assertions, in much the same way as it pays special note to interests *observe*, and replicates them in the parent context. Similarly, *inbound* c is an assertion from the parent dataspace pertinent to the current one. Again, routing pays special attention to interest in such assertions, *observe inbound* c , replicating the interest *observe* c in the parent context and wrapping received events with *inbound*.

To accommodate this hierarchical runtime, we must change the type syntax in three ways. The first change to λ_{ds}^U is the addition of expression, value, and type forms for the new assertions. The second alteration changes the status of dataspace; rather than being reserved for descriptions of complete programs, *dataspace* is now an expression form, akin to an actor action:

$$M = \dots \mid \text{outbound } M \mid \text{inbound } M \mid \text{dataspace } \tau_c M$$

Rule T-DATASPACE changes to reflect the status of dataspace as actor actions:

$$\frac{\text{flat}(\tau_c) \quad \Gamma \vdash M : \text{List Actor } \tau_c \quad \tau_{ctx} = \widehat{ds\text{-route}}(\tau_c)}{\Gamma \vdash \text{dataspace } \tau_c M : \text{Actor } \tau_{ctx}} \quad \boxed{\text{T-DATASPACE}}$$

The third change concerns the metafunction $\widehat{ds\text{-route}}$, which encapsulates the outbound- and inbound-sensitive routing described above. It synthesizes a type, τ_{ctx} in which the described dataspace operates. When viewed from the parent context,

- $\text{strip-out}(\tau)$ is the type of assertions produced by the nested dataspace;
- $\text{relay-interests}(\tau)$ is the type of interests it produces; and
- $\text{strip-in}(\tau)$ is the type of assertions it expects to match stated interests.

Definition 30.

$$\begin{aligned} \widehat{ds\text{-route}} & : \text{Type} \longrightarrow \text{Type} \\ \widehat{ds\text{-route}}(\tau) & = \tau_{out} \cup \tau_{in} \cup \tau_{relay} \\ & \text{where } \tau_{out} = \text{strip-out}(\tau) \\ & \quad \tau_{in} = \text{strip-in}(\tau) \\ & \quad \tau_{relay} = \text{relay-interests}(\tau) \end{aligned}$$

The *strip-out*, *strip-in*, and *relay-interests* metafunctions follow the same structure as *strip-obs* (defined in figure 10).

Definition 31.

$$\begin{aligned}
& \textit{strip-out} & : & \mathbf{Type} \longrightarrow \mathbf{Type} \\
\textit{strip-out}(\text{Outbound } \tau) & = & \tau \\
\textit{strip-out}(\star) & = & \star \\
\textit{strip-out}(\bigcup \vec{\tau}) & = & \bigcup \overrightarrow{\textit{strip-out}(\tau)} \\
\textit{strip-out}(\tau) & = & \perp & \text{otherwise}
\end{aligned}$$

Definition 32.

$$\begin{aligned}
& \textit{strip-in} & : & \mathbf{Type} \longrightarrow \mathbf{Type} \\
\textit{strip-in}(\text{Inbound } \tau) & = & \tau \\
\textit{strip-in}(\star) & = & \star \\
\textit{strip-in}(\bigcup \vec{\tau}) & = & \bigcup \overrightarrow{\textit{strip-in}(\tau)} \\
\textit{strip-in}(\tau) & = & \perp & \text{otherwise}
\end{aligned}$$

Definition 33.

$$\begin{aligned}
& \textit{relay-interests} & : & \mathbf{Type} \longrightarrow \mathbf{Type} \\
\textit{relay-interests}(\text{Observe } (\text{Inbound } \tau)) & = & \text{Observe } \tau \\
\textit{relay-interests}(\star) & = & \star \\
\textit{relay-interests}(\bigcup \vec{\tau}) & = & \bigcup \overrightarrow{\textit{relay-interests}(\tau)} \\
\textit{relay-interests}(\tau) & = & \perp & \text{otherwise}
\end{aligned}$$

All of the properties shown in sections 4.4 and 5.3 can be adapted to the system including nested dataspaces as well.

