

Programming and Reasoning about Actors that Share State

SAM CALDWELL

e-mail: samc@ccs.neu.edu

Northeastern University

TONY GARNOCK-JONES

e-mail: tony.garnock-jones@maastrichtuniversity.nl

Maastricht University

MATTHIAS FELLEISEN

e-mail: matthias@ccs.neu.edu

Northeastern University

Abstract

Actor languages realize concurrency via message passing, which most of the time is easy to use. Empirical code inspection provides evidence, however, that on occasion, programmers wish to have an actor share some of its state with others. The dataspace model adds a tightly controlled state-exchange mechanism, dubbed dataspace, to the actor model for just this purpose. Experience with dataspace suggests that this form of sharing calls for linguistic constructs that allow programmers to state temporal aspects of actor conversations. In response, this paper presents the facet notation: its theory, its type system, its behavioral type system, and some first experiences with an implementation.

1 Introduction

Actor languages and libraries support an easy-to-use mechanism for expressing concurrent computations. The program spawns as many actors as needed, and actors strictly communicate via messages. Pure message passing eliminates the possibility of race conditions arising from a lack of mutual exclusion, a major problem in shared-state concurrency. An investigation (Tasharofi et al., 2013) of actively used actor programming libraries suggests, however, that programmers often use actors with threads and shared state, if available. When interviewed, programmers tend to point to the ease of sharing some state between computational components running in parallel. In short, programmers wish for some amount of sharing among actors.

The dataspace actor model supplements the message-exchanges of the actor model with a space for sharing state in a controlled manner (Garnock-Jones et al., 2014; Garnock-Jones and Felleisen, 2016; Caldwell et al., 2020). It accepts the actor model's premise that computation is conversation—exchange of messages—but it also embraces the demand of

47 conversations for context. Roughly speaking, a context consists of the public information—
 48 corresponding to fragments of an actor’s current state—that the participants remain aware
 49 of for a period of time. Such public information, dubbed an assertion, is placed into a
 50 shared dataspace. From a programmer’s perspective, an assertion is a part of an actor’s
 51 state that it wishes to share publicly.

52 From the dataspace’s point of view, an actor is a function that reacts to changes in the
 53 published states of other actors. Even a message is a change to this shared state: it shows up
 54 and disappears immediately. Indeed, basic message-passing may be expressed with datas-
 55 spaces in just this manner. The functional, black-box interface for actors lends itself to
 56 system-level reasoning but, when it comes to expressing the behavior of an individual actor,
 57 a purely functional notation is *not* a good fit. The core programming tasks—situationally
 58 engaging and disengaging in behavior and managing and reacting to changes in shared
 59 state—are inherently temporal. Hence, the model calls for a novel way of expressing a
 60 concurrent actor in a context-sensitive conversation.

61 In response, this paper presents the *design of a language of facets* and a system for
 62 *reasoning about their temporal behavior*. From the perspective of this facet language,
 63 an actor consists of several pieces—the facets—each of which represents the actor as a
 64 participant in distinct (parts of) conversations. The concept addresses the central concerns
 65 of actors:

- 66 • engaging and disengaging in conversations;
- 67 • maintaining local state;
- 68 • permitting access to shared, public state; and
- 69 • reacting to external events.

71 The benefits of facets are evident in improving the organization of dataspace actor pro-
 72 grams and supporting a conversational style of concurrency. Furthermore, facets directly
 73 enable type-based static reasoning about communication.
 74

75 **Dining Concurrency Researchers.** To illustrate our meaning of conversational concu-
 76 rency, imagine a table of dining and conversing attendees at a conference on concurrency
 77 research. Experience suggests a number of common characteristics for such a scenario:
 78

- 79 • The conversation is likely to branch into several subconversations as new topics are
 80 broached. Previous topics may be returned to—or not.
- 81 • The participants may engage in one subconversation or several simultaneously; they
 82 may leave one conversation to join another, only to return later to the first one.
- 83 • The set of participants changes over time. Late-arrivals may join the table after
 84 the conversation commences while early-departers exit the conversation before its
 85 conclusion.
- 86 • The conversation is contextual. A newcomer to the conversation is likely to listen for
 87 a while to pick up the context. Alternatively, this newcomer may explicitly state an
 88 interest in the contextual background.
- 89 • Finally, these participants identify one another through a number of characteris-
 90 tics, including their presentations at the conference, their interests, their association
 91

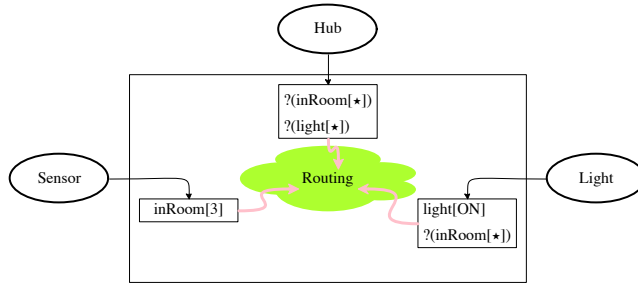


Fig. 1: Dataspace With Actors Engaged In a Smart Home Conversation

with particular institutions, and so on. They may or may not know each others' names. Moreover, a particular participant's identifying characteristic may vary from conversation to conversation.

While Dijkstra's dining philosophers illustrate how coordination avoids certain pitfalls of concurrency, namely deadlocked program states, the scenario of dining conference attendees paints a portrait of how developers desire well-designed concurrent programs to behave.

The notions of dataspace and facet equip developers with mechanisms to program in this conversational style: means for representing and sharing context, evolving actor behavior in the face of shifting demands, and a flexible notion of identity and participation in a conversation. Sections 2 and 3 return to these ideas.

Roadmap. The presentation proceeds in three steps: sections 2 through 4 introduce the ideas via examples; the formal models of facet actors and dataspace programs are the subject of sections 5, 6, and 7, respectively; and section 8 states the key theorem concerning the temporal type checking of facet actors. Section 9 briefly explains the implementation and section 10 reports the results of evaluating the ability to verify behavioral properties of realistic programs. The remaining sections are about related and future work.

2 Background: Dataspaces, Actors

The dataspace actor model generalizes the publish/subscribe protocol of Linda's Tuplespaces (Carriero et al., 1994). When actors wish to share pieces of state with others, they deposit those with a dedicated actor, the *dataspace*. All actors connected to a dataspace can query this shared space. The dataspace thus controls how published state is shared within a group of conversing actors.

A simulation of a smart home allows a comprehensive illustration. Let us say the smart home supports three kinds of devices: smart lights; a presence sensor, which detects how many people are in a room; and a control hub for accessing and controlling the other devices. The simulation expresses each of these devices as an actor (*Light*, *Sensor*, *Hub*) embedded in a common dataspace.

Figure 1 depicts a moment in such a simulation. The connecting dataspace is the large rectangle. Each actor has deposited aspects of its state pertinent to the others; these are called *assertions*.¹ For example, the Sensor actor has published the assertion `inRoom[3]`, denoting the number of people in the room. Published pieces of state—assertions—are made up from values (e.g. numbers, strings) as well as immutable data structures (lists, structs).

An actor subscribes to publications of state of others via *assertions of interest*. In the simulation, the Light actor turns a bulb on or off based on the occupancy of the room, meaning it needs to know about the `inRoom` assertion. To get this information and keep it up-to-date, it registers `?(inRoom[*])` with the dataspace: `inRoom[*]` is an *assertion pattern* where `*` is a match-anything wildcard and `?` turns it into an assertion of interest, or just interest for short.

Communication in dataspace programs arises from matches between interests and assertions. The dataspace provides a routing mechanism, depicted in figure 1 as a cloud, for detecting such matches. Routing tracks changes to the set of assertions matching each actor's interests. When this set changes, the dataspace generates an *event* describing the change and informs the interested actor(s). In the running example of the smart home, the presence Sensor actor updates its assertion from `inRoom[3]` to `inRoom[2]` if someone leaves. In turn, the dataspace informs the Light actor of this change, because of its interest.

An actor's response to an event is one of the following actions:

- it may add and/or subtract assertions from its published state;
- it may spawn additional actors; and
- it may exit, by choice or by an uncaught run-time exception.

When an actor exits, the dataspace removes all of its assertions. Since this removal is a change to the assertion set, it triggers an event for all interested actors. In response, other actors may react to component failures in a graceful manner. For example, when a light burns out in the smart home simulation, its corresponding Light actor exits. This triggers the removal of its `light[ON]` assertion. In response, the dataspace informs the Hub actor because it is interested in such assertions. This actor may then take appropriate action, such as display a notification.

Dataspace communication is anonymous. That is, the model does not provide actors with an address or process-ID (PID) to serve as a communication handle. Rather, correlation information—that is, data dictating to which actor(s) some data is relevant—must be included in the exchanged assertions and interests. This allows identifying information to be encoded in a domain-specific fashion. For example, the Light actor could include its serial number in its assertion `light[serial, ON]`, allowing future communications specific to that device to refer to it by its serial number.

Dining Concurrency Researchers. The dataspace model addresses several needs of the conversational style introduced in the dining concurrency setting of the introduction:

¹ The term assertion is inspired by the resemblance to a fact in the database of a Prolog program (Clocksin and Mellish, 1981).

- Assertions directly represent conversational context.
- Anonymous actors give a flexible notion of identity and conversational participant.
- The notion of a change-in-state as an event merges the means through which a new participant receives prior context in a conversation with normal communication.
- Similarly, notification of withdrawn assertions is the same—and automatic—whether via normal behavior or exceptional crashes.

In sum, the dataspace model abstracts over the processing of routing events by actors. The model requires only that an actor has a functional event-transducer interface:

$$Event \times State \rightarrow Actions \times State$$

How to express an implementation of this interface is an entirely different question.

3 Facets

With the facet notation, a programmer expresses an individual actor as a tree of facets. Roughly speaking, a facet groups some of the actor's state with the behavior related to a particular conversation. The behavior must specify the contributions (assertions and interests) to a conversation together with reactions to the utterances of other participants. The tree structure reflects the tendency of conversations to branch from one another while accruing context. This section introduces the facet language with the code for some of the actors from the smart home simulation of the preceding section.

```

1 (define (light-actor)
2   (start-facet light-facet
3     (field [light-state : Bool ON])
4     (assert (light (! light-state)))
5     (on (asserted (in-room 0)) (:= light-state OFF))
6     (on (retracted (in-room 0)) (:= light-state ON))))
7
8 (spawn (light-actor))

```

Fig. 2: First Faceted Actor

Figure 2 presents the code for the `Light` actor written in our Racket-based implementation. The listing includes the `spawn` expression that launches the actor (line 8). The actor engages in a single conversation concerning the presence of people in the room and the state of the light. Hence, it starts a single facet (line 2), named `light-facet`. The `start-facet` form imperatively extends the actor's facet tree with additional behavior. The new facet joins the tree as a child of the enclosing facet. The first facet started by the actor, as is the case for `light-facet`, becomes the root of the tree.

A facet keeps the actor-relevant state in *fields*. The `light-facet` facet creates a single field, which keeps track of the current state of the light (line 3). This `light-state` field holds values of type `Bool` and is initialized to `ON`, a constant defined elsewhere in the program. The expression `(! light-state)` accesses the current value of the field, while `(:= light-state OFF)` updates it. The type of the `light-state` field handle itself is

(`Field Bool`), analogous to the type of a mutable reference cell. A type annotation on a field is optional; when elided, it defaults to the type of the initialization expression.

A facet interacts with the connected dataspace via *endpoints*. Endpoints come in two varieties—assertions and event handlers—corresponding to incoming and outgoing information. The `light-facet` facet comes with three endpoints:

- The first one is an assertion endpoint (line 4). It shares information about the state of the light with the dataspace. The assertion is a struct of type `light`. The `light` assertion carries one value, the contents of the `light-state` field accessed via `!`. The endpoint establishes an assertion in the dataspace while the facet it belongs to is active. In this case, since the actor boots with the `light-facet` facet, it initially makes such an assertion. Moreover, the facet keeps its assertion up to date. When an assertion endpoint refers to a field, the endpoint and field establish a dataflow link. The dataflow mechanism propagates field updates in order to keep the actor’s assertions up-to-date with respect to its fields. Thus, every time `light-state` changes, the actor automatically withdraws the previous `light` assertion and replaces it with a version reflecting the new `light-state`.
- The other endpoints are on (event) handlers (lines 5-6). Each consists of two pieces:
 1. The event specification describes the nature and structure of the event. The nature of an event is expressed in terms of a keyword—`asserted` or `retracted`—corresponding to the appearance or disappearance of assertions in the dataspace. The structure of the event is expressed with a pattern, articulating a query for specific structures in the dataspace. The actor automatically issues an assertion of interest corresponding to the event handler’s pattern. In the running example, the event handler wishes to know about `in-room` assertions in the dataspace; it does so with placing `?(in-room 0)` in the dataspace. An event handler’s assertion of interest is initialized in the same manner as an `assert` endpoint. Thus, the actor boots with assertions of interests corresponding to its initial event handlers. Moreover, if the patterns of the event handlers reference the values of any fields, the assertions of interest are kept up to date through the same dataflow mechanism.
 2. The body part of the handler is code that executes in response to every matching event. In the running example, each event handler’s body updates the facet’s state via the `light-state` field. Recall that updates to this field automatically propagate to the facet’s assertion via dataflow.

The `light-actor` function is effectful. When called from inside an actor `spawn` statement (line 8), or the body of a facet’s event handler endpoint, it has the effect of starting a new facet with the described behavior. In our implementation programmers may mix facet operations, field creation, and endpoint creation together with definitions and expressions. With this flexibility, any aspect of an actor’s behavior may be defined inside a procedure. This can improve the readability of the code as well as facilitate re-use and sharing of code among different actor definitions.

Figure 3 visualizes the `light-facet` facet of the `Light` actor. The facet encapsulates the `light-state` field and the two endpoints. The solid-blue arrow indicates the dataflow

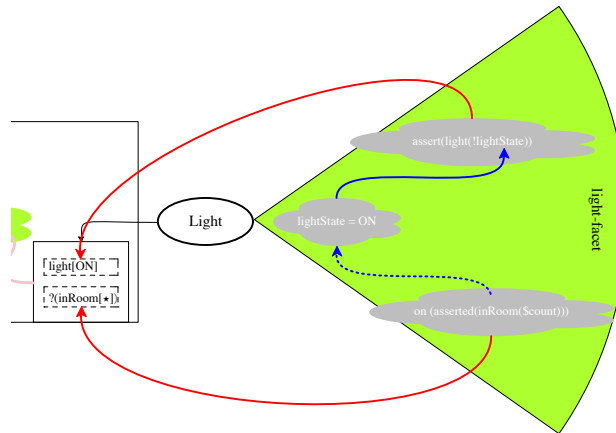


Fig. 3: Visualization of an Actor's Facets and Facet's Endpoints

link between the field and the assertion endpoint; the dotted-blue arrow indicates that a value flows from the body of the `on` handler into the field. Finally, the red arrows show the manifestation of the endpoints in the dataspace as an assertion and an interest for the two endpoints, respectively.

Let's equip the smart home simulation with additional features and observe how the language of facets facilitates the corresponding adaptation of the `Light` actor:

Wall Switches control the power to lights in a room. When a switch is off, the light must be off and is unable to communicate with other devices. Flipping the switch on turns the light on and allows for communication with, and control by, other devices.

Multiple Rooms complicate the simulation, too. Certain interactions are possible only among devices located within the same room.

Configuration is essential for software systems. Initially, the only devices are the control hub and the wall switches. The user may install lights and presence sensors. After installing a device, the user must assign it to a specific room via the hub.

Figure 4 shows how to adapt the implementation of the `Light` actor to this revised scenario. An unrelated difference to figure 2 is the use of `spawn` inside of the `spawn-light-actor` function—a simplification of the presentation. This function is now called by a UI component when the user installs a light. It takes one `String`-typed argument. The argument designates to which wall switch the light is connected; that is, the argument represents the inherent connection between elements on a power circuit. Once the function is called, it generates a unique ID for the light (line 2), akin to a manufacturer serial number. After that, it spawns the light actor with one initial facet (lines 3-4).

```

1 (define (spawn-light-actor [wall-switch-id : String])
2   (define my-id (generate-unique-id "light"))
3   (spawn
4     (start-facet light-facet
5       (during (wall-switch wall-switch-id ON)
6         (field [light-state ON])
7         (assert (light my-id (! light-state))))
8       (during (room-assignment my-id $room)
9         (on (asserted (in-room room 0)) (:= light-state OFF))
10        (on (retracted (in-room room 0)) (:= light-state ON))))))

```

Fig. 4: Implementation of the Light actor, draft

Equipped with a basic understanding of this function, let us look at the revised `light-facet` facet. Recall that the actor should *not* engage in any behavior unless the connected switch is flipped on. Once it is on, the actor participates in conversation(s) until the switch is turned off. A basic facet-oriented actor would start a facet in response to the first kind of event and shut it down when the second kind occurs. This pattern is so common that the facet language comes with a `during` expression (line 5), which combines an assertion pattern that can be used with either `asserted` or `retracted`.²

Here the facet operates while the assertion of `(wall-switch wall-switch-id ON)` exists in the dataspace—representing the interval of time during which the room’s switch is in the on position. As the switch is flipped into the on position, the facet announces its existence and (ON) state using a field and assertion endpoint as before (lines 6-7). This `light` assertion allows the Hub actor to detect the installation of a new device. If this is the first time the Hub actor encounters this device ID, it prompts the user to assign it to a room. Once the assignment is complete, the Hub actor deposits a `room-assignment` assertion (line 8) to inform the Light actor. The sub-pattern `$room` binds the corresponding portion of the `room-assignment` assertion to the name `room`. Knowing what room it is in allows the Light actor to converse with any presence sensor in the same room, turning on and off accordingly (lines 9-10).

Multi-faceted Actors. The revised implementation of the Light actor starts several facets. Eventually it may run three concurrent facets: `light-facet` and one per `during` expression. Each of these facets corresponds to a particular context within an ongoing conversation. Furthermore, the facets are nested, just like actor conversations. When one facet starts another, the first is the parent to the second. Stopping a facet (via `stop`) means terminating it and all of its children. In the context of the Light actor, the `wall-switch` facet may shut down when the switch is flipped off; in this case, shutting down this facet also shuts down the nested facets associated with `during` expressions. A `stop` form may optionally specify some continuation behavior, as in `(stop orig (start-facet cont . . .))`. In that case, the `cont` facet takes the place of `orig` in the tree. Continuation facets address the need to transition from one task to another.

Wider Facet Trees. While the facets in the Light actor are nested in a linear fashion, bushy trees are equally common. Consider the Hub actor. A user interacts with the Hub to

² In Racket, `during` is just a notational definition (macro); see section 3.1

configure each light. Once configured, the Hub both monitors and controls these lights. To accomplish this task, the actor starts a facet for each separate conversation with the various Light actors.

369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393

```
(define-type LightDict (Dictionary String Light))

(define (launch-hub)
  (spawn
    (start-facet hub
      (field [lights : LightDict (create-dictionary)])
      (during (light $id _)
        (on start
          (match (dict-ref (! lights) id #false)
            [(light room _) (control-light lights id room)]
            [#false (get-room-from-user id)]))))))

;; auxiliary functions:
(define (get-room-from-user [id : String])
  (start-facet get-room
    (define room ... interact with user to assign a room ...)
    (stop get-room (control-light lights id room))))

(define (control-light [lights : (Field LightDict)] [id : String]
  [room : String])
  (start-facet control
    (assert (room-assignment id room))
    (on (asserted (light id $o?))
      (set-light-state lights id room o?))
    (on stop (set-light-state id room OFF))))

(define (set-light-state [lights : (Field LightDict)] [id : String]
  [room : String] [on? : Bool])
  (:= lights (dict-set (! lights) id (light room on?))))
```

Fig. 5: Implementation of the Hub actor

394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414

Figure 5 shows a portion of the actor’s implementation. The actor starts a facet named hub (line 5), which creates one field (lights, line 6) for storing information about the lights in the system. The field is a dictionary. For each registered light, this dictionary associates the light’s ID with a light struct that records its room and its known state. The auxiliary function set-light-state (lines 27–29) takes a reference to the lights field as an argument, of type (Field LightDict), and updates the dictionary for a particular light.

Using during, the actor converses with each individual light. The during reacts in response to new light assertions.³ Such assertions can mean one of two things: the installation of a light or the restoration of power to an already installed light. In the first case, lights does not contain an entry with that ID; in the second case, it does. Each of these situations calls for different behavior, but no matter what, the reaction should happen only when the facet starts up.

³ The “don’t care” pattern _ means the during reacts to the initial appearance of a light assertion with a particular ID, but not to subsequent updates to the on/off state of that light.

To express this idea directly, the language enables facets to react to the start and stop events, respectively. That is, instead of a reaction to a dataspace assertion, a facet's end-points may specify a start or stop specific reaction. Here the `during` facet uses `start` to decide how to act (line 8). In the handler (lines 9–11), it uses `lights` to look up what it knows about the light. In the first case and `match` branch, the facet retrieves the light's room assignment and starts another facet to control it via `control-light` (lines 19–25). In the second case, the actor calls `get-room-from-user`, which engages in yet another dialog with the user to assign the new light to a room. This nested facet stops itself, with some explicit continuation behavior: a call to `control-light`.

The `control-light` function starts a facet to assert the light's room-assignment and to react to changes in its state. When the state changes, the facet updates `lights`. Finally, when the facet stops—due to termination of the outer `during`—it records the light as off. Stop handlers, like exception handlers, enable separation between the decision to disengage in behavior and the definition of how to clean up in that event. In this case, the decision is implicit in the `during` on line 7's retracted event handler. Meanwhile, the code for cleaning up is in the `control-light` function (line 24). Crucially, the need to clean up only arises after an initial provisioning step, so it would not make sense to have the functionality at the same level as the `during`. In the case of a crash, an actor does not run any stop event handlers. Since the state of the actor could be inconsistent, with fields partially updated, orderly shutdown may not be possible.

Figure 6 visualizes the facets of the `Light` and `Hub` actors. In the case of the `Light` actor, the tree is deeply nested, as the actor engages in new behaviors as it accumulates conversational context. By contrast, the `Hub` actor's facet tree exhibits breadth, as it starts a new facet in reaction to each installed light. In turn, each of those facets have sub-trees either to interact with the user or to control the light. Each sub-tree of the hub facet is independent of the others, evolving based on its interactions with its controlled light and the user without needing to know the current state (or existence) of the others.



Fig. 6: Visualization of the `Light` and `Hub` actors' facet trees

3.1 Derived Forms

Figure 7 displays derived forms for common patterns. One exceedingly common pattern is for a facet to be constructed in response to the appearance of an assertion and torn down in response to its disappearance. The `Light` and `Hub` actors both fall into this category. We abstract this pattern into the `during` form. Each time the asserted event fires, any pattern

variables bound in *pattern* are instantiated to yield a concrete assertion, *pattern'*. The `(on (retracted ...))` endpoint monitors this assertion in the new facet.

Figure 7 also shows a related derived form, `during/spawn`. In contrast to `during`, this form does not create a facet within the actor but spawns an actor in the dataspace. The critical difference concerns failure. While a failure in `during` tears down the current actor, a failure in `during/spawn` terminates the separate actor but leaves the spawning one alone.

<pre> 468 (during pattern during-body ...) 469 <u>def</u> 470 (on (asserted pattern) 471 (start-facet theBody 472 (on (retracted pattern') 473 (stop theBody)) 474 during-body ...)) </pre>	<pre> 468 (during/spawn pattern during-body ...) 469 <u>def</u> 470 (on (asserted pattern) 471 (spawn (start-facet theBody 472 (on (retracted pattern') 473 (stop theBody)) 474 during-body ...))) </pre>
<pre> 474 (define/query-value id expr₀ pattern expr) 475 <u>def</u> 476 (field [id expr₀]) 477 (on (asserted pattern) 478 (:= id expr)) 479 (on (retracted pattern) 480 (:= id expr₀)) </pre>	<pre> 474 (define/query-set id pattern expr) 475 <u>def</u> 476 (field [id (Set τ) empty-set]) 477 (on (asserted pattern) 478 (:= id (set-add (! id) expr))) 479 (on (retracted pattern) 480 (:= id (set-remove (! id) expr))) </pre>

Fig. 7: Derived Forms

Both `during` and `during/spawn` allow the programmer to directly express *matching of supply to demand*. Assertions matched by the *pattern* are interpreted as demand; the *during-body* constitutes the supply. As demand increases via new matching assertions in the dataspace, the supply expressions are executed to match. As demand *decreases* again due to the withdrawal of assertions, `(on (retracted ...))` endpoints automatically terminate the facet or actor in response. Resources are allocated via `(on start ...)` clauses and released in `(on stop ...)` clauses in response to changing needs.

Figure 7 defines a second family of derived forms concerning the integration of newly arrived assertions into local fields. While actors often just deal with singleton assertions, other scenarios call for a set, a hash table, or even an aggregate summary of several assertions. To support this idiom, the facet language provides derived constructs called *queries*. Queries define and subsequently update fields, making their results available for use in neighboring facet endpoints.

One such form, `define/query-value`, is useful for local tracking of an assertion kind in which there may be zero or one instance in the dataspace. For example, we could extend the smart home program with persistence, so that the user's assignments of devices to rooms are saved between runs of the program. A separate actor would handle persisting such configurations to the file system or other location and loading them when the program starts. Then, when the Hub actor boots, it needs to operate based on a reloaded configuration or start with a default fresh one. It may use

```

504 (define/query-value config DC (configuration $lights ...)
505   (configuration lights ...))

```

to create a field that defaults to DC (short for default configuration) value unless there is a configuration assertion, in which case it uses the value of that assertion.

The second such form, `define/query-set`, collects information from relevant assertions in a set.⁴ A user-interface component for the smart home may use it to show the user all of the lights currently installed in the system:

```
(define/query-set lights (light $id _) id)
```

The query scheme directly generalises to structures such as hash tables, etc., and also allows *aggregations* analogous to SQL's `COUNT(*)` and `GROUP BY`.

3.2 The Case for Facets

Returning to the dining concurrency researchers example, facets support a conversational style in several ways:

- Engaging in a new conversation is as simple as starting a facet. Likewise, stopping a facet disengages from a conversation.
- The `during` form directly connects context to behavior.
- The tree of facets mirrors the way (sub)conversations branch from existing conversations.

Defining actors using facets provides additional advantages over basic functions and objects, as well. Appendix A provides an implementation of the light actor using a procedural notation as a concrete basis of comparison, demonstrating and explaining these benefits.

Generally put, the facet notation combines and extends a number of concepts from other languages and empowers programmers to use them in a synergistic manner. For example, end-point programming may remind the reader of the event-handling and reactive programming style of Esterel (Berry and Gonthier, 1992) and CRIME (Mostinckx et al., 2008). Dataflow exists in many forms (Demetrescu et al., 2011). Updates to behavior have been around as long as the Actor model (Hewitt et al., 1973). The `during` behavior is the clearest example of the synergistic use of these features. It precisely describes a temporal engagement with the dataspace; sets up and tears down local behavior as needed; and simultaneously takes into account the accumulated conversational context.

4 Static Checks: Types & Behavior

Programming actors requires the development of protocols. A protocol coordinates the flow of data among actors in a conversation. Designing the program requires reasoning about protocols while writing code. An actor language ought to assist programmers with this reasoning process and ideally with checking it statically.

The language of facets comes with a structural type system and a behavioral one. The first is somewhat conventional (section 4.1); also see Caldwell et al. (2020)'s work on

⁴ The type of the set's elements is inferred from the provided *expr*.

structural types for dataspace. The second enables programmers to specify an actor’s dynamic behavior (section 4.2) and perform some static checking (section 4.4). The bridge between the two is an effect-type system (section 4.3). Section 4.5 applies this checking to an example specification from the smart home.

4.1 Basic Types

The starting point is a structural type system. Base types describe data such as `Int`, `Bool`, etc. A struct-type definition such as `(struct light ([on? : Bool]))` introduces the type constructor `LightT`.⁵ Thus, `(light #true)` has type `(LightT Bool)`.

Our previous work on structural types (Caldwell et al., 2020) for purely functional data-space actors validates that it prevents mistakes involving the shape of assertions. That is, type checking incorporates knowledge of how a dataspace routes assertions among actors and uses the types to confirm that assertions cannot upset routing. It is straightforward to adapt this system and its proof of soundness to facet-oriented dataspace actors.

4.2 Specification

Like for structural properties of routing assertions, programmers need a language for behavioral—that is, temporal—properties, too. Like structural types give a syntax and semantics to claims about properties of assertions, behavioral types should allow the programmer to make claims concerning the behavior of actors in a dataspace. We employ a form of linear temporal logic (LTL) for this purpose, though of course, its basic propositions must somehow involve the assertions that actors deposit into and withdraw from dataspace. In addition to introducing this language, this section also informally explains what it means for a facet-based actor program to live up to such a specification.

LTL has a long history as a specification language for concurrent programs and the basis for mechanized checks (Pnueli, 1977; Manna and Pnueli, 1991; Holzmann, 1997). When compared to other temporal logics, such as the branching-time computation tree logic (CTL), LTL is typically considered more intuitive and at least as amenable to efficient model checking (Vardi, 2001; Rozier, 2011). Moreover, compared to typical protocol description languages such as choreographies (Carbone and Montesi, 2013) or session types (Honda et al., 2008), LTL allows the focus to be on *what* is communicated rather than *who* does the communicating, suiting the anonymous style of dataspace programming.

The base LTL predicates are types that describe assertions in the dataspace. For example, the formula `(LightT Bool)` holds at any moment iff some actor has deposited an assertion `(light #true)` or `(light #false)` in the dataspace. Specifications may also use negation (`Not`), and `(Not (LightT Bool))` means no assertion of this structure type is in the dataspace.

Other connectives are conjunction (`And`), disjunction (`Or`), and logical implication (`Implies`). With temporal connectives, a programmer can specify the evolution of data-space programs. The proposition `(Always (LightT Bool))` holds if the dataspace

⁵ The final “T” is a convention for such type constructors. The `struct` form defines the plain `Light` type name to be an alias for the type constructor fully instantiated with the default type parameters, such as `(LightT Bool)`.

contains a `light` assertion at every moment during execution, regardless of which actor has deposited it there. Meanwhile,

```
(Until (LightT Bool) (Not WallSwitchOn))
```

says that eventually there will not be a `(wall-switch-on)` assertion in the dataspace, but until then there is a `light` assertion. The formula

```
(Eventually (LightT Bool))
```

describes an execution state with an assertion matching `(LightT Bool)` in the dataspace or the execution arrives at such a state within a finite number of steps.

A dataspace program satisfies a specification if its dataspace evolves according to the interpretation of the LTL formula. Thus, a program lives up to the just-mentioned formula if, at some point during execution, an actor makes a matching `light` assertion. In particular, any dataspace program that includes the `light-facet` actor from the preceding section satisfies this formula, because it makes a `light` assertion as it boots.

Consider another proposition:

(Always		
(Implies		
(LightT Bool)		It is always true that, if a <code>light</code> assertion is
(Until		in the dataspace, eventually there will not be
(LightT Bool)		a <code>wall-switch-on</code> assertion but until that
(Not WallSwitchOn))))		time a <code>light</code> assertion will remain in the
		dataspace.

It is the strong version of `Until`, requiring that `(Not WallSwitchOn)` eventually be true.

Take the dataspace program that runs the two actors on the left:

actor 1	actor 2	dataspace evolution	
		t =	assertions
		0	\emptyset
(start light-facet	(start wall-switch	1	{(light ON),
(field [state ON])	(assert (wall-switch-on))	2	(wall-switch-on)}
(assert (light (! state)))	(on (asserted (light _)))	3	{(light ON)}
(on (retracted (wall-switch-on)))	(stop wall-switch)))		{(light OFF)}
(:= light OFF)))			

The evolution of the dataspace is on the right. It shows which assertions are in the dataspace for the first four time steps, starting with the boot state.

After booting, actor 1 makes a `light` assertion and actor 2 adds a `(wall-switch-on)` to the dataspace. Hence the dataspace configuration at that point matches the antecedent of the implication. Consequently, the `Until` property must also describe that state so that the latter models the complete formula. Due to actor 1's assertion, actor 2 stops its only facet, leading to the withdrawal of the `(wall-switch-on)` assertion. Since the `light` assertion remains active until that point, the implication still holds. Furthermore, actor 1 notices the disappearance of the `(wall-switch-on)`, and in response, replaces its own assertion the dataspace. As a result, the next configurations also matches the antecedent of the implication as does the until property.

While the notion of satisfaction is intuitive, the eventual goal is to mechanize this type-checking of specifications. In this regard, the sample explanations are also suggestive.

What is needed is an abstract description of the behavior of the actors and their interaction with dataspace via assertions. And again, just like for structural type checking, it is the *shape* of exchanges that can serve in this role, not the precise assertions.

4.3 Effect Types for Facets

The end of the preceding section almost dictates the third step. Historically, type systems research uses effect-type systems for descriptions of program behavior (Lucassen and Gifford, 1988). In the setting of facet-oriented actors, the to-be-observed effects are caused by actors overall, facets, and endpoints. Hence, each corresponding term-level construct comes with a type-level construct for describing the type of its effects. See figure 8 for the actual notation.

Actor and Facet Descriptors		Event Descriptors	
Term	Type	Term	Type
(spawn t ...)	(Spawn T ...)	(asserted p)	(Asserted P)
(start-facet x t ...)	(StartFacet x T ...)	(retracted p)	(Retracted P)
(stop x t ...)	(Stop x T ...)	start	Start
(assert t)	(Assert T)	stop	Stop
(on evt t ...)	(On Evt T ...)		

Also, $\$x:T$ and $_$ have types $\$T$ and $_$, respectively.

Fig. 8: Effect Types

Let us illustrate the synthesis of effect types with the code for spawning the simplistic `Light` actor from section 3 (without rooms and power switches). Ignoring fields and expressions, the code describes a single actor and a single facet nested within this actor. The facet itself comes with three end points: a plain assertion endpoint, an on-asserted one and an on-retracted one. The behavioral type mirrors this description, including the ordering of the endpoints within the facet:

```

674 (spawn                                     (Spawn
675   (start-facet light-facet                 (StartFacet light-facet
676     (field [light-state ON])                (Assert (LightT Bool))
677     (assert (light (! light-state)))        :: (On (Asserted (InRoomT Int)))
678     (on (asserted (in-room 0))              (On (Retracted (InRoomT Int))))))
679     (:= light-state OFF))
678     (on (retracted (in-room 0))
679       (:= light-state ON)))

```

Here is how to read the type in terms of the actor's communication behavior:

- it makes an assertion of type `(LightT Bool)`;
- it expresses interest to assertions of type `(InRoomT Int)`;
- it reacts to their appearance and disappearance; and
- its reactions do not change its communication behavior.

For this last point, note that the bodies of the `On` types are empty.

From this perspective, a behavioral type is a *simplified* actor. This simplified actor communicates via *types* of assertions rather than value assertions. By taking the behavior type

of each actor in a program—say, the smart home program—we get a communication-only program, that is, a collection of simplified actors and their types of assertion exchanges with dataspace. Accordingly, the number of possible behaviors is much lower—and that is precisely what enables a mechanical checking of (some) specifications.

4.4 Checking Specifications Mechanically

The transition of possible interactions from concrete value assertions to types of assertions enable specification checking. At a rather high level, the checker infers effect types from a program and compiles them to the language of a model checker. The compilation encodes effect types as simplified actors. This model-checking program is combined with the programmer’s LTL formula. If the model-checking attempt fails, the error is reported in terms of a trace of these simplified actor programs.

The theory behind this idea is presented in sections 5 and 7. Section 8 verifies why this approach works. The implementation is sketched in section 9. Section 10 evaluates its ability to check behavioral properties of realistic programs. The following subsection illustrates the working of the model checker with an example.

4.5 Revisiting the Smart-Home Example

To demonstrate the power of the specification language, let us revisit the smart home example from section 3. The code presented in figure 4 is buggy, and attempting to check it against a natural specification exposes the bug. Here is the specification of the expected behavior of the lights:

```
(Always (And (Implies (InRoomT Room NonZero)
                  (Implies WallSwitchOn (Eventually LightOn)))
             (Implies (InRoomT Room Zero)
                  (Eventually (Not LightOn)))))
```

This specification describes the expected outcomes of interactions between the presence sensor, light, and light-switch actors. It states that the light turns on and off in response to the sensor actor’s `InRoomT` assertions, as long as it is powered (`WallSwitchOn`).

When analyzed with respect to this specification, our implementation of LTL checking reports the error and provides a counterexample in terms of the involved actors and their assertions:

```
...
Process Sensor RETRACTS RoomEmpty
...
Process Switch ASSERTS WallSwitchOn
...
```

If a presence sensor is installed and configured *after* the light in the same room is turned on even though it is empty, the light stays on due to the sensor’s initial reading.

Based on this trace, we can identify the active assertions and thus the active facets of the `Light` actor. This assessment provides a starting point for determining where to locate the bug. A close look reveals that the `light-facet` facet in figure 4 reacts only to the

retraction of `(in-room room 0)` to determine that the room is occupied. If the room is occupied when the light regains power, no such retraction will take place. To fix this problem, we can amend the facet to query the state of `in-room` assertions on startup. The boxed code in figure 9 is all that is needed to make the actor satisfy its LTL specification.

```

737
738
739
740
741
742 (define (spawn-light wall-switch-id)
743   (define my-id (generate-unique-id "light"))
744   (spawn
745     (start-facet light-facet
746       (during (wall-switch wall-switch-id ON)
747         (field [light-state ON])
748         (assert (light my-id (! light-state))))
749       (during (room-assignment my-id $room)
750         (on (asserted (in-room room 0)) (:= light-state OFF))
751         (on (retracted (in-room room 0)) (:= light-state ON))))
752     (start-facet init
753       (on (asserted (in-room room $n))
754         (:= light-state (not (zero? n)))
755         (stop init))))
756   ))))

```

Fig. 9: Implementation of the Light actor, fixed (see figure 4)

5 Facets: The Semantics

A proper language design should come with a rigorous blueprint, especially if its rationale includes claims about the reasoning power of its type system. The presented language consists of two pieces: the facet notation and its connection to dataspaces. The former deserves a formal semantics; the latter has been presented thoroughly elsewhere (Caldwell et al., 2020; Garnock-Jones and Felleisen, 2016) but section 7.1 provides an abbreviated treatment so that the meta-theorems can be stated in a self-contained manner.

This section presents the formal model of a *single* facet-based actor (sections 5.1 and 5.2). It describes how a facet-based actor updates its state in response to an event: the evolution of the facet tree; the invocation of event handlers; and field maintenance.

5.1 Formal Syntax

Figure 10 defines the abstract syntax of facet-oriented actors. Statements `Pr` imperatively update the actor's state in one of several ways:

- `start fn \vec{e} (\overrightarrow{D} Pr)`—starting a facet with name `fn`, assertions \vec{e} , and event handlers (\overrightarrow{D} Pr);
- `stop fn Pr`—terminating a facet, along with all of its children, and engaging in some continuation behavior Pr;
- `field x = e in Pr`—creating a mutable field for Pr;
- `x := e`—assigning a new value to a field; and
- `spawn Pr`—spawning an additional actor.

<pre> 783 Pr ∈ Prog = start fn \vec{e} (\overrightarrow{D} Pr) 784 stop fn Pr 785 spawn Pr 786 field x = e in Pr 787 x := e 788 let x = e in Pr 789 Pr; Pr 790 skip 791 792 D ∈ EvtDsc = asserted e 793 retracted e 794 start 795 stop </pre>	<pre> 783 e ∈ Expr = b 784 x 785 !x 786 p(\vec{e}) 787 ? e 788 m(\vec{e}) 789 * 790 x : τ 791 792 fn ∈ FName = facet names 793 x ∈ Var = variables 794 p ∈ PrimOp = primitive operations 795 m ∈ MsgCtor = message labels 796 b ∈ BasicVal = basic values: strings, etc. </pre>
---	--

Fig. 10: Program Syntax

Statements may introduce local, lexically scoped variables with `let x = e in Pr`. Otherwise, they are just like statements in other programming languages that compose sequentially (`Pr; Pr`). For technical reasons, the grammar includes no-op (`skip`).

The set of expressions comprises

- b —basic values;
- $p(\vec{e})$ —invocations of potentially partial primitive operations;
- x —references to variable names;
- $!x$ —field access;
- $? e$ —assertion of interest in e ;
- $m(\vec{e})$ —tuple with tag m ;
- $*$ —wildcard assertion patterns/templates; and
- $x : \tau$ —binding patterns.

A vector of assertion templates (\vec{e}) defines the assertions made by a facet. Templates describe families of assertions built with labels (m), interests ($?$), wildcards ($*$), and expressions (e).

Event handlers consist of an event description (D) and a body (Pr). Events describe occurrences that are either internal to the actor (`start`, `stop`) or particular assertions in the dataspace (`asserted e`, `retracted e`). In the latter case, the description includes a pattern expression e defining the relevant assertions. In order to construct patterns, expressions additionally include match-anything wildcards ($*$) and binding variables ($x : \tau$).

In comparison to the implemented syntax, the model uses some obvious short-hands and some less obvious simplifications:

- The assertion endpoints of a facet are all declared together.
- The event handlers come as a sequence of description/handler body pairs.
- The fields shared by a facet’s endpoints are declared outside the facet.
- The assertion labels replace `struct` names.
- Binding variables in patterns require a type annotation.

Here is the basic `Light` actor from figure 2 in the model’s syntax:

```

829     field state = ON in
830       start light-facet
831         light(!state)
832         (asserted in-room(count: Int)
833           state := if(equal?(count, 0), OFF, ON))

```

5.2 Semantics

Since individual actors consume events and produce assertions, we choose to specify the semantics of facet-oriented actors via a labeled transition system. Each transition acts on a machine state and constructs the next one. The transition labels l describe either inputs in the form of event patches (Δ) or outputs in the form of spawned actors ($\overline{\text{Pr}}$). An unlabeled transition relation (\bullet), represents actor-internal work. Patches Δ comprise two disjoint assertions sets π , representing added and removed assertions. Assertions u range over basic values, labeled tuples, and assertions of interest.

The formulation of machines requires syntax for describing the states in addition to the syntax of programs. Figure 11 defines this syntax and the machine states. The `error` state represents an actor that has crashed due to an unhandled error.

An actor machine state consists of five components: $\langle \text{FT}; \overrightarrow{\text{I}}; \overrightarrow{\text{PS}}; \pi; \sigma \rangle$. The components track different kinds of state and knowledge:

- *the activity state of facets*

The machine must express which facets are currently active as well as the parent-child relationship among them. This aspect takes the form of a tree: FT . A node $\text{fn}[\overrightarrow{e}(\overrightarrow{\text{D Pr}})].\overrightarrow{\text{FT}}$ in this tree describes an active facet in terms of its name, its assertion and event handler endpoints, and its children, while the empty tree takes the form ε .

- *the instruction stream within the active facets*

A machine state identifies a sequence of ready-to-perform instructions I , and it also maintains a queue of pending scripts PS . A pending script (fid, Pr) is a sequence of statements to be performed and pertinent context information. The context is a facet ID $\langle \overrightarrow{\text{fn}} \rangle$ identifying the facet from which the script originated. A facet ID is a sequence of facet names forming the path from the root of the tree to that particular facet. Instructions I can `start` a new facet at a designated location in the tree; `stop` a running one; and `spawn` an actor.

- *the information known to the actor*

Specifically, the store σ holds the contents of the fields, while the set of assertions π keeps track of the actor's current knowledge of assertions in the dataspace. Newly booted facets receive these assertions as an event.⁶

Finally, we need some miscellaneous pieces of syntax. An `Evt` is either a patch or facet `start/stop`. Evaluation reduces expressions to values v .

⁶ Due to the particulars of dataspace routing (section 7), the actor will not receive another event from the dataspace describing these assertions until they disappear and subsequently reappear. That is, the dataspace tracks the set of assertions known to an *actor*; not facet. A new facet of the actor may include an event handler for which an already-received assertion is relevant.

875 $M \in \mathbf{Machine} = \langle \overrightarrow{FT}; \overrightarrow{I}; \overrightarrow{PS}; \pi; \sigma \rangle$ error 876 877 $FT \in \mathbf{FctTree} = \varepsilon$ $\text{fn}[\overrightarrow{c}(\overrightarrow{D Pr})].\overrightarrow{FT}$ 878 879 $C \in \mathbf{Context} = \square$ $\text{fn}[\overrightarrow{c}(\overrightarrow{D Pr})].\overrightarrow{FT} \cdot C \cdot \overrightarrow{FT}$ 880 881 $I \in \mathbf{Instr} = \text{start fn } \overrightarrow{c}(\overrightarrow{D Pr}) @ \text{fid}$ stop fn spawn Pr 882 883 884 $PS \in \mathbf{PScript} = (\text{fid}, Pr)$ 885 886 $\text{fid} \in \mathbf{FacetID} = \langle \overrightarrow{\text{fn}} \rangle$ 887 888 $\text{fn} \in \mathbf{FName} = \text{facet names}$ 889 890 $\sigma \in \mathbf{Store} = x \xrightarrow{\text{fn}} v$	$u \in \mathbf{Assertion} = b$ $m(\overrightarrow{u})$ $? u$ 876 $\pi \in \mathbf{ASet} = \mathcal{P}(\mathbf{Assertion})$ 878 $\Delta \in \mathbf{Patch} = \pi^+ / \pi^-$ where $\pi^+ \cap \pi^- = \emptyset$ 879 $\text{Evt} \in \mathbf{Event} = \Delta$ start stop 882 $v \in \mathbf{Val} = b$ $? v$ $m(\overrightarrow{v})$ $*$ $x : \tau$ 884 $l \in \mathbf{Label} = \bullet$ Δ \overrightarrow{Pr} 886 887 888 889 890
---	---

Fig. 11: Evaluation Syntax

892 $\langle l-c[I_0]; :: ; \emptyset; \{ \text{state} \mapsto \text{ON} \} \rangle$ 893 where $l-c = \text{light-facet}$ 894 $I_0 = [\text{light}(!\text{state})$ 895 $(\text{asserted in-room}(\text{count} : \text{Int})$ 896 $\text{state} := \text{if}(\text{equal}?(\text{count}, 0), \text{OFF}, \text{ON}))].\varepsilon$ <hr style="width: 50%; margin: 10px auto;"/> 897 $\frac{\{ \text{in-room}(0) \} / \emptyset}{\bullet} \rightarrow$ 898 $\langle l-c[I_1]; :: ; \text{PS}; \{ \text{in-room}(0) \}; \{ \text{state} \mapsto \text{ON} \} \rangle$ 899 where $\text{PS} = (\langle l-c \rangle, \text{state} := \text{if}(\text{equal}?(0, 0), \text{OFF}, \text{ON}))$ 900 $I_1 = [\text{light}(!\text{state})$ 901 $(\text{asserted in-room}(\text{count} : \text{Int})$ 902 $\text{state} := \text{if}(\text{equal}?(\text{count}, 0), \text{OFF}, \text{ON}))].\varepsilon$ <hr style="width: 50%; margin: 10px auto;"/> 903 $\frac{\bullet}{\bullet} \rightarrow$ 904 $\langle l-c[I_2]; :: ; \{ \text{in-room}(0) \}; \{ \text{state} \mapsto \text{OFF} \} \rangle$ 905 where $I_2 = [\text{light}(!\text{state})$ 906 $(\text{asserted in-room}(\text{count} : \text{Int})$ 907 $\text{state} := \text{if}(\text{equal}?(\text{count}, 0), \text{OFF}, \text{ON}))].\varepsilon$	(by inject) (by transfer)
--	------------------------------

Fig. 12: The light-facet facet as an initial machine state and its transitions

Figure 12—above the horizontal line—shows the initial machine state M for the light-facet facet from above. The rest of the figure displays the first two transitions of the machine.

Transition Relations. Figure 13 defines the transition relation on machine states that describes the behavior of a single facet-based actor in response to incoming events. For the moment we ignore how such events arrive; section 7 covers that aspect of the semantics. We use $M \longrightarrow M'$ as shorthand for $M \xrightarrow{\bullet} M'$.

Generally speaking, the machine works in the following fashion. Initially, the machine state describes an inert actor, meaning there are no instructions to perform nor pending scripts to execute. When the machine receives an event patch, it is matched against the

event handlers of the actor's facets. For each successful match, the body of the event handler is enqueued as a pending script.

Once the matching process is completed, the machine executes the scripts. It dequeues the first one and interprets the instructions, one at a time. Some instructions generate internal events (facet start and stop). For those the machine matches them again against the facet tree, and this process may enqueue additional scripts. The machine continues in this manner until it reaches inertness again. At that time, another external event may be injected.

Here are descriptions of the machine's five transition rules:

inject —The rule's purpose is to dispatch an event to an inert actor. Due to dataspace routing, the event is guaranteed to conform to the actor's interests. Incorporating the incoming patch via \oplus yields a new set π' of assertions known by the actor. The *dispatch* metafunction matches the event against each handler in the tree of facets, yielding a pending script for each match.

transfer —The machine transfers the next pending script to the current register and partially evaluates (*p-e-s*) its internal statements. In doing so, it updates and creates fields in the store. The process eliminates expressions and yields a sequence of instructions that alter the active facet tree (*start/stop*) or correspond to an output by the actor (*spawn*). Partially evaluating the script means *eval*-ing expressions, which may include applications of partial primitives. In the event that such a primitive fails, the actor crashes. The *error* state represents a crashed actor without any behavior.

start —The next instruction demands the start of a facet. The new facet is given a fresh name; the bodies of each of its event handlers refer to the new name. The *dispatch* function is used to create two scripts: one to signal a *start* event and one to boot the event handlers that correspond to the known assertions of the actor. These scripts are enqueued and the new facet is inserted in the tree. The instruction specifies the location for the new facet in the tree as a path (*fid*) to the new facet's parent. The *locate* metafunction constructs the appropriate context. It is possible that no such context exists, which happens when the desired parent, or one of its ancestors, is terminated via a *stop* instruction in the time between the script containing this *start* statement being enqueued and the instruction executing. In that case, the new facet is immediately terminated. The *dispatch* function informs it of the *stop* event. The resulting scripts are enqueued, and the facet is never inserted into the tree.

stop —The machine terminates an active facet and its children, *dispatch*-ing a *stop* event to allow these facets to shut down in an orderly manner. The instruction may designate a facet that is already eliminated. In this case there is no more.

spawn —The facet wishes to create an actor, which the machine treats as an output instruction. Formally, the transition comes with a $\overline{\text{Pr}}$ label.

The top-most machine state in figure 12 transitions to an inert state in two steps. The first one injects the script, consuming the patch $\{\text{in-room}(0)\}/\emptyset$ as the label indicates. The second enqueues the script's instruction and performs it.

967	$\langle \overrightarrow{\text{FT}}; \cdot; \cdot; \pi; \sigma \rangle \xrightarrow{\Delta} \langle \overrightarrow{\text{FT}}; \cdot; \overrightarrow{\text{PS}}; \pi'; \sigma \rangle$	(inject)
968	where	
969	$\pi' = \pi \oplus \Delta$	
970	$\overrightarrow{\text{PS}} = \text{dispatch}(\overrightarrow{\text{FT}}, \Delta, \pi, \pi', \sigma, \langle \cdot \rangle)$	
971	$\langle \overrightarrow{\text{FT}}; \cdot; (\text{fid}, \text{Pr}) \cdot \overrightarrow{\text{PS}}'; \pi; \sigma \rangle \xrightarrow{\bullet} M'$	(transfer)
972	where	
973	$M' = \begin{cases} \langle \overrightarrow{\text{FT}}; (\text{fid}, \overrightarrow{\text{I}}); \overrightarrow{\text{PS}}; \pi; \sigma' \rangle & \text{if } \overrightarrow{\text{I}}, \sigma' = p\text{-e}(\text{Pr}, \sigma, \text{fid}) \\ \text{error} & \text{otherwise} \end{cases}$	
974	$\langle \overrightarrow{\text{FT}}; \text{start fn } \overrightarrow{e} (\overrightarrow{\text{D Pr}}) @ \text{fid} \cdot \overrightarrow{\text{I}}; \overrightarrow{\text{PS}}; \pi; \sigma \rangle \xrightarrow{\bullet} \langle \overrightarrow{\text{FT}}'; \overrightarrow{\text{I}}; \overrightarrow{\text{PS}} \cdot \overrightarrow{\text{PS}}_{\text{start}} \cdot \overrightarrow{\text{PS}}_{\text{boot}} \cdot \overrightarrow{\text{PS}}_{\text{stop}}; \pi; \sigma \rangle$	(start)
975	where	
976	fn_{new} fresh in $\overrightarrow{\text{FT}}$	
977	$\overrightarrow{\text{FT}}_{\text{new}} = \text{fn}_{\text{new}}[\overrightarrow{e} (\overrightarrow{\text{D Pr}}[\text{fn} \mapsto \text{fn}_{\text{new}}])].\varepsilon$	
978	$\overrightarrow{\text{PS}}_{\text{start}} = \text{dispatch}(\overrightarrow{\text{FT}}_{\text{new}}, \text{start}, \emptyset, \emptyset, \sigma, \text{fid})$	
979	$\overrightarrow{\text{PS}}_{\text{boot}} = \text{dispatch}(\overrightarrow{\text{FT}}_{\text{new}}, \pi/\emptyset, \emptyset, \pi, \sigma, \text{fid})$	
980	$\overrightarrow{\text{FT}}' = \begin{cases} \text{C}[\overrightarrow{\text{FT}}_{\text{new}}] & \text{if } \text{locate}(\overrightarrow{\text{FT}}, \text{fid}) = \text{C} \\ \overrightarrow{\text{FT}} & \text{otherwise} \end{cases}$	
981	$\overrightarrow{\text{PS}}_{\text{stop}} = \begin{cases} \cdot & \text{if } \text{locate}(\overrightarrow{\text{FT}}, \text{fid}) \text{ defined} \\ \text{dispatch}(\overrightarrow{\text{FT}}_{\text{new}}, \text{stop}, \pi, \pi, \sigma, \text{fid}) & \text{otherwise} \end{cases}$	
982	$\langle \overrightarrow{\text{FT}}; \text{stop fn } \cdot \overrightarrow{\text{I}}; \overrightarrow{\text{PS}}; \pi; \sigma \rangle \xrightarrow{\bullet} \langle \overrightarrow{\text{FT}}'; \overrightarrow{\text{I}}; \overrightarrow{\text{PS}} \cdot \overrightarrow{\text{PS}}_{\text{stop}}; \pi; \sigma \rangle$	(stop)
983	where	
984	$\overrightarrow{\text{FT}}' = \begin{cases} \text{C}[\varepsilon] & \text{if } \overrightarrow{\text{FT}} = \text{C}[\text{fn}[\dots].\overrightarrow{\text{FT}}] \\ \overrightarrow{\text{FT}} & \text{otherwise} \end{cases}$	
985	$\text{fid} = \text{facet-context}(\text{C})$	
986	$\overrightarrow{\text{PS}}_{\text{stop}} = \begin{cases} \text{dispatch}(\text{fn}[\dots].\overrightarrow{\text{FT}}, \text{stop}, \pi, \pi, \sigma, \text{fid}) & \text{if } \overrightarrow{\text{FT}} = \text{C}[\text{fn}[\dots].\overrightarrow{\text{FT}}] \\ \cdot & \text{otherwise} \end{cases}$	
987		
988	$\langle \overrightarrow{\text{FT}}; \text{spawn Pr} \cdot \overrightarrow{\text{I}}; \overrightarrow{\text{PS}}; \pi; \sigma \rangle \xrightarrow{\overrightarrow{\text{Pr}}} \langle \overrightarrow{\text{FT}}; \overrightarrow{\text{I}}; \overrightarrow{\text{PS}}; \pi; \sigma \rangle$	(spawn)

Fig. 13: Transition Relation

5.2.1 Metafunctions

The transition system of figure 13 refers to several metafunctions. This section provides their formal definitions. The definitions make use of secondary metafunctions, which are defined in appendix B.

The boot_{Pr} function initializes the machine from an actor description. It recognizes only scripts that create some number of fields and then `start` a single facet. It creates each field and then boots the facet as a child of a synthetic *root* facet with no behavior. The synthetic root facet allows the initial facet to be replaced by any number of facets via `stop`:

$$\begin{aligned}
 \text{boot}_{\text{Pr}} & : \text{Pr} \times \sigma \xrightarrow{\text{partial}} \mathbf{M} \\
 \text{boot}_{\text{Pr}}(\text{field } x = e \text{ in Pr}, \sigma) & = \text{boot}_{\text{Pr}}(\text{Pr}[x \mapsto x'], \sigma[x' \mapsto v]) \\
 & \quad \text{if } v = \text{eval}(e, \sigma) \\
 & \quad \text{where} \\
 & \quad x' \text{ fresh in } \sigma, \text{Pr} \\
 \text{boot}_{\text{Pr}}(\text{let } x = e \text{ in Pr}, \sigma) & = \text{boot}_{\text{Pr}}(\text{Pr}[x \mapsto v], \sigma) \\
 & \quad \text{if } v = \text{eval}(e, \sigma) \\
 \text{boot}_{\text{Pr}}(\text{start fn } \overrightarrow{e} (\overrightarrow{\text{D Pr}}), \sigma) & = \langle \text{fn}_{\text{root}}[\cdot].\overrightarrow{\text{FT}}; \cdot; \overrightarrow{\text{PS}}; \emptyset; \sigma \rangle \\
 & \quad \text{where} \\
 & \quad \overrightarrow{\text{FT}} = \text{fn}[\overrightarrow{e} (\overrightarrow{\text{D Pr}})].\varepsilon \\
 & \quad \overrightarrow{\text{PS}} = \text{dispatch}(\overrightarrow{\text{FT}}, \text{start}, \emptyset, \emptyset, \sigma, \text{fn}_{\text{root}}) \\
 & \quad \text{fn}_{\text{root}} \text{ free in } \overrightarrow{\text{FT}}
 \end{aligned}$$

The *dispatch* function matches an event against each handler in a tree of facets, accumulating the successes as a list of pending scripts:

$$\begin{aligned}
1013 \quad & \text{dispatch} & : & \text{FT} \times \text{Evt} \times \pi \times \pi \times \sigma \times \text{fid} \longrightarrow \vec{\text{PS}} \\
1014 \quad & \text{dispatch}(\varepsilon, \text{Evt}, \pi, \pi', \sigma, \text{fid}) & = & \cdot \\
1015 \quad & \text{dispatch}(\text{fn}[\vec{e}(\text{D Pr})].\vec{\text{FT}}, \text{Evt}, \pi, \pi', \sigma, \langle \vec{\text{fn}}_{\text{ctx}} \rangle) & = & \frac{\text{dispatch1}(\langle \vec{\text{fn}}_{\text{ctx}} \cdot \text{fn} \rangle, \text{D}, \text{Pr}, \text{Evt}, \pi, \pi', \sigma)}{\text{dispatch}(\text{FT}, \text{Evt}, \pi, \pi', \sigma, \langle \vec{\text{fn}}_{\text{ctx}} \cdot \text{fn} \rangle)}
\end{aligned}$$

It utilizes a helper function, *dispatch1*, to instantiate the body of a single event handler with the substitution(s) arising from matching against a single event. Substitutions γ map variables to values and matching a pattern against a set of assertions yields a set of substitutions \mathcal{S} :

$$\begin{aligned}
1024 \quad & \text{dispatch1} & : & \text{fid} \times \text{D} \times \text{Pr} \times \text{Evt} \times \pi \times \pi \times \sigma \longrightarrow \vec{\text{PS}} \\
1025 \quad & \text{dispatch1}(\text{fid}, \text{D}, \text{Pr}, \text{Evt}, \pi, \pi', \sigma) & = & \begin{cases} (\text{fid}, \text{unroll}(m)) & \text{if } \mathcal{S} = \text{match}_{\text{D}}(\text{D}, \text{Evt}, \pi, \pi', \sigma), \mathcal{S} \neq \emptyset \\ & \text{where } m = \{\gamma(\text{Pr}) \mid \gamma \in \mathcal{S}\} \\ \cdot & \text{otherwise} \end{cases}
\end{aligned}$$

Finally, *unroll* sequentializes a set of statements using an arbitrary yet fixed ordering:

$$\begin{aligned}
1029 \quad & \text{unroll} & : & \mathcal{P}(\text{Pr}) \longrightarrow \text{Pr} \\
1030 \quad & \text{unroll}(\emptyset) & = & \text{skip} \\
1031 \quad & \text{unroll}(\{\text{Pr}\} \uplus m) & = & \text{Pr}; \text{unroll}(m)
\end{aligned}$$

The *p-e* function partially-evaluates a script, yielding a sequence of instructions and an updated store:

$$\begin{aligned}
1034 \quad & p-e & : & \text{Pr} \times \sigma \times \text{fid} \xrightarrow{\text{partial}} \vec{\text{I}} \times \sigma \\
1035 \quad & p-e(\text{skip}, \sigma, \text{fid}) & = & \cdot \\
1036 \quad & p-e(\text{Pr}_1; \text{Pr}_2, \sigma, \text{fid}) & = & \vec{\text{I}}_1 \cdot \vec{\text{I}}_2, \sigma'' \\
1037 \quad & & & \text{if} \\
1038 \quad & & & \vec{\text{I}}_1, \sigma' = p-e(\text{Pr}_1, \sigma, \text{fid}) \\
1039 \quad & & & \vec{\text{I}}_2, \sigma'' = p-e(\text{Pr}_2, \sigma', \text{fid}) \\
1040 \quad & p-e(\text{start fn } \vec{e}(\text{D Pr}), \sigma, \text{fid}) & = & \text{start fn } \vec{e}(\text{D Pr}) @ \text{fid}, \sigma \\
1041 \quad & p-e(\text{stop fn } \text{Pr}, \sigma, \langle \vec{\text{fn}} \cdot \vec{\text{fn}}' \rangle) & = & \text{stop fn } \vec{\text{I}}, \sigma' \\
1042 \quad & & & \text{if } p-e(\text{Pr}, \sigma, \langle \vec{\text{fn}}' \rangle) = \vec{\text{I}}, \sigma' \\
1043 \quad & p-e(\text{spawn Pr}, \sigma, \text{fid}) & = & \text{spawn Pr}, \sigma \\
1044 \quad & p-e(\text{let } x = e \text{ in Pr}, \sigma, \text{fid}) & = & p-e(\text{Pr}[x \mapsto v], \sigma, \text{fid}) \text{ if } v = \text{eval}(e, \sigma) \\
1045 \quad & p-e(\text{field } x = e \text{ in Pr}, \sigma, \text{fid}) & = & p-e(\text{Pr}', \sigma', \text{fid}) \text{ if } v = \text{eval}(e, \sigma) \\
1046 \quad & & & \text{where} \\
1047 \quad & & & x' \text{ fresh in } \sigma, \text{Pr} \\
1048 \quad & & & \sigma' = \sigma[x' \mapsto v] \\
1049 \quad & & & \text{Pr}' = \text{Pr}[x \mapsto x'] \\
1050 \quad & p-e(x := e, \sigma \uplus \{x \mapsto v\}, \text{fid}) & = & \cdot, \sigma[x \mapsto v] \text{ if } v = \text{eval}(e, \sigma)
\end{aligned}$$

The *facet-context* function produces the path of facet names (a fid) leading to the hole in a context:

$$\begin{aligned}
1049 \quad & \text{facet-context} & : & \text{C} \longrightarrow \text{fid} \\
1050 \quad & \text{facet-context}(\square) & = & \langle \cdot \rangle \\
1051 \quad & \text{facet-context}(\text{fn}[\vec{e}(\text{D Pr})].\vec{\text{FT}} \cdot \text{C} \cdot \vec{\text{FT}}') & = & \langle \text{fn} \cdot \vec{\text{fn}}' \rangle \\
1052 \quad & & & \text{where } \langle \vec{\text{fn}}' \rangle = \text{facet-context}(\text{C})
\end{aligned}$$

Meanwhile, *locate* performs the reverse direction, producing the context for a particular fid in a tree of facets:

$$\begin{aligned}
1055 \quad & \text{locate} & : & \text{FT} \times \text{fid} \xrightarrow{\text{partial}} \text{C} \\
1056 \quad & \text{locate}(\text{FT}, \langle \cdot \rangle) & = & \square \\
1057 \quad & \text{locate}(\text{fn}[\vec{e}(\text{D Pr})].\vec{\text{FT}}, \langle \text{fn} \cdot \vec{\text{fn}}' \rangle) & = & \text{fn}[\vec{e}(\text{D Pr})].\vec{\text{FT}}_1 \cdot \text{C} \cdot \vec{\text{FT}}_2 \\
1058 \quad & & & \text{if} \\
1059 \quad & & & \text{locate}(\text{FT}', \langle \vec{\text{fn}}' \rangle) = \text{C} \\
1060 \quad & & & \vec{\text{FT}} = \vec{\text{FT}}_1 \cdot \text{FT}' \cdot \vec{\text{FT}}_2
\end{aligned}$$

1059	$T \in \mathbf{FacetTy} =$	$\tau \in \mathbf{Type} =$	
1060	\mid	\mid	
1061	\mid	\mid	
1062	\mid	\mid	
1063	$D_T \in \mathbf{EvtDscTy} =$	$B \in \mathbf{BaseTy} =$	\mid
1064	\mid	\mid	\mid
1065	\mid	\mid	\mid
1066	\mid	\mid	\mid
1067	\mid	\mid	\mid
1068	\mid	\mid	\mid
1069	\mid	\mid	\mid
1070	\mid	\mid	\mid
1071	\mid	\mid	\mid
1072	\mid	\mid	\mid
1073	\mid	\mid	\mid
1074	\mid	\mid	\mid
1075	\mid	\mid	\mid
1076	\mid	\mid	\mid
1077	\mid	\mid	\mid
1078	\mid	\mid	\mid
1079	\mid	\mid	\mid
1080	\mid	\mid	\mid
1081	\mid	\mid	\mid
1082	\mid	\mid	\mid
1083	\mid	\mid	\mid
1084	\mid	\mid	\mid
1085	\mid	\mid	\mid
1086	\mid	\mid	\mid
1087	\mid	\mid	\mid
1088	\mid	\mid	\mid
1089	\mid	\mid	\mid
1090	\mid	\mid	\mid
1091	\mid	\mid	\mid
1092	\mid	\mid	\mid
1093	\mid	\mid	\mid
1094	\mid	\mid	\mid
1095	\mid	\mid	\mid
1096	\mid	\mid	\mid
1097	\mid	\mid	\mid
1098	\mid	\mid	\mid
1099	\mid	\mid	\mid
1100	\mid	\mid	\mid
1101	\mid	\mid	\mid
1102	\mid	\mid	\mid
1103	\mid	\mid	\mid
1104	\mid	\mid	\mid

Fig. 14: Type Syntax

6 Types for Facets

Facet types (section 6.1) play two roles. First, they ensure basic type safety in the form of a standard soundness theorem. Second, they reflect facet operations to the type level with a type-and-effect style. By lifting the individual actor semantics to the type level (section 6.2), and then dataspace programs, type-level programs can be assigned an operational semantics that enables model checking temporal logic claims.

6.1 Type Syntax and Judgment

Figure 14 defines the syntax of types. Effect types T directly reflect the core facet syntax (figure 10), including facet *start/stop* and *spawn* operations of Pr. Event descriptor types D_T correspond directly to the term level descriptors D . Types τ summarize the result of expressions, with base types B describing primitive values. Type environments Γ associate variable names from patterns and fields with their types and facet names with the marker *FacetName*.

Figure 15 specifies the core type judgment on facets: $\Gamma \vdash_{Pr} Pr : T$. It ensures that facet and field names are used appropriately while collecting certain facet operations as an effect type T . As mentioned, this effect type is used for behavioral analysis. Auxiliary judgments apply to the different categories of syntax; appendix D provides their definitions.

Here are explanations of the key typing rules:

T-START assigns types to a facet-creation instruction. It appeals to an auxiliary judgment for checking the assertion expressions e , which also produces a description of the assertion types τ . The condition $\text{assertable}(\tau)$, defined in appendix D, ensures that assertion types do not contain binding patterns or field names. The event descriptors D for each event handler are checked in the same manner, yielding a type D_T for each. The final aspect to check is the body Pr of each event handler. For each such body, the environment is extended with the name of the facet being started as well as the binding variables from the pattern in the event descriptor (via the bindings_D metafunction). Checking each body produces a type level description of its behavior, T . The product of the rule is a *Start* effect, combining the types of the assertions plus each event handler description and its body.

$$\begin{array}{c}
\boxed{\Gamma \vdash_{Pr} Pr : T} \\
\text{T-SKIP} \quad \frac{}{\Gamma \vdash_{Pr} \text{skip} : \text{Skip}} \quad \text{T-SEQ} \quad \frac{\Gamma \vdash_{Pr} Pr_1 : T_1 \quad \Gamma \vdash_{Pr} Pr_2 : T_2}{\Gamma \vdash_{Pr} Pr_1; Pr_2 : T_1; T_2} \\
\text{T-START} \quad \frac{\overrightarrow{\text{assertable}(\tau)} \quad \overrightarrow{\Gamma \vdash_D D : D_T} \quad \overrightarrow{\Gamma \vdash_e e : \tau} \quad \overrightarrow{\Gamma, \text{fn} : \text{FacetName}, \text{bindings}_D(D) \vdash_{Pr} Pr : T}}{\Gamma \vdash_{Pr} \text{start fn } \vec{c} \overrightarrow{(D Pr)} : \text{Start fn } \vec{c} \overrightarrow{(D_T T)}} \\
\text{T-STOP} \quad \frac{\Gamma(\text{fn}) = \text{FacetName} \quad \text{prune-up-to}(\text{fn}, \Gamma) \vdash_{Pr} Pr : T}{\Gamma \vdash_{Pr} \text{stop fn Pr} : \text{Stop fn T}} \\
\text{T-SPAWN} \quad \frac{\text{prune}(\Gamma) \vdash_{Pr} Pr : \text{Start}[\dots]}{\Gamma \vdash_{Pr} \text{spawn Pr} : \text{Spawn Start}[\dots]} \\
\text{T-FIELD} \quad \frac{\Gamma \vdash_e e : \tau \quad \Gamma, x : \text{Field } \tau \vdash_{Pr} Pr : T}{\Gamma \vdash_{Pr} \text{field } x = e \text{ in Pr} : T} \quad \text{T-ASSIGN} \quad \frac{\Gamma \vdash_e e : \tau \quad \Gamma(x) = \text{Field } \tau}{\Gamma \vdash_{Pr} x := e : \text{Skip}} \\
\text{T-LET} \quad \frac{\Gamma \vdash_e e : \tau \quad \Gamma, x : \tau \vdash_{Pr} Pr : T}{\Gamma \vdash_{Pr} \text{let } x = e \text{ in Pr} : T}
\end{array}$$

Fig. 15: Facet Typing

T-STOP is the analogue of T-START for facet termination. It consults the environment to ensure that the named facet is in context. The *prune-up-to* function removes the target facet's name from the environment, as well as the names of any of its descendants, when checking the continuation Pr. The resulting effect type is to Stop the same facet with a continuation behavior.

T-FIELD, T-ASSIGN check field creation and assignment, respectively. The judgment for field creation checks the type of the initial expression and associates the name with an appropriate type in the environment for the remaining statements. The judgment for field assignment looks up that type in the environment, making sure that the type of the assigned expression matches. Field updates and references are not included in effect types. Thus, field creation yields the effect type from its enclosed statements while field assignment yields a Skip effect.

T-SPAWN concerns the creation of actors. The rule utilizes the *prune* function (appendix D) to remove identifiers associated with fields and facet names from the running actor, because these names are local. The spawned actor's behavior Pr is checked within this restricted context, and it has the (sole) effect of Start-ing a facet. The effect type of the spawn statement is a Spawn with a type description of the actor's initial facet.

6.2 Machine Types and Type-level Machines

The judgment $\overrightarrow{\Gamma}_M \vdash_M M : M_T$ relates a machine M to a machine type M_T . A machine type $\langle \overrightarrow{\text{FT}}_T; \overrightarrow{\text{IT}}_T; \overrightarrow{\text{PS}}_T; \pi_\tau; \sigma_\tau \rangle$ includes a facet tree type ($\overrightarrow{\text{FT}}_T$), sequence of instruction types ($\overrightarrow{\text{IT}}_T$), sequence of pending script types ($\overrightarrow{\text{PS}}_T$), set of assertion types (π_τ), and store type (σ_τ). Appendix F defines the full judgment and syntax. In brief, the syntax of machine

types reflects the syntax of machines (figure 11) in the same manner that effect types T (figure 14) reflect the syntax of statements Pr (figure 10).

The statement of soundness for facet-based actors relies on two auxiliary notions. First, we say M is an inert machine state, $\text{inert}(M)$, iff $M = \langle \text{FT}; \cdot; \cdot; \pi; \sigma \rangle$. That is, it has neither instructions to perform nor pending scripts to execute. Second, we designate a set of *internal* transition labels, l_\bullet . Internal labels correspond to the machine processing instructions and scripts. Thus, they consist of each \bullet and $\overline{\text{Pr}}$ label, but not Δ .

Theorem 1 (Soundness). *If $\vdash_M M : M_T$ and $M \xrightarrow{\Delta} M'$ then either:*

- $M' \xrightarrow{l_\bullet}^* M''$ and $\text{inert}(M'')$; or
- $M' \xrightarrow{l_\bullet}^* \text{error}$

Proof By the standard progress and preservation properties (Wright and Felleisen, 1994). Appendix G contains the full statement and proof of each property. To show termination of internal reduction sequences, note the lack of recursive facilities in the language, and that each event handler in a machine's facet tree may activate a maximum of one time between external events. \square

If a well typed single-actor machine configuration can take an *inject* step, it eventually transitions to an inert machine state or a state that represents an exceptional state (due to a misapplication of a partial primitive).

The transition system for machine states operates on machine types with only minor syntactic adjustments. Appendix F defines the full type-level transition system. Section 8 makes use of this notion for behavioral analysis of actors. Here, we note that transitions on type machines are well-defined.

Theorem 2 (Well-definedness). *Either $M_T \xrightarrow{l_\tau} M'_T$ or $\text{inert}(M_T)$.*

Proof Follows that of theorem 1. \square

7 Dataspaces: The Semantics of Programs

While the preceding section describes the semantics of a *single* actor, a full understanding necessitates a semantics of entire dataspace programs, that is, programs with many, dynamically created and terminated, actors. Figure 16 defines the abstract syntax (left) and the semantics (right) of dataspace programs.

A dataspace configuration consists of a queue of pending patches (ℓ, Δ) , labeled by the originating actor; a table of assertions (R); and a series of actors (A). The table R is a set of pairs associating each assertion with the identity ℓ of the originating actor. An active actor in the dataspace $\ell \mapsto \Sigma$ associates an address with an actor state, Σ . The state consists of a pending event Δ and behavior B . The pending event accumulates a patch for the actor between its turns. An actor behavior (f, v) pairs a behavior function with a private state. The behavior function f is a total mapping between a pending event and the private state of the actor to a new state and possibly some actions (*act*). The model

1197	$DS \in \mathbf{Dataspace}$	=	$\overrightarrow{[(\ell, \Delta); R; \vec{A}]}$	$\overrightarrow{[(\ell_{act}, \Delta); R; \vec{A}]} \longrightarrow_{DS} \overrightarrow{[(\ell, \Delta') \cdot (\ell', \pi/\emptyset); R'; (\ell_A \mapsto \Sigma') \cdot (\ell' \mapsto \Sigma'')]}$ where
1198	$R \in \mathbf{Table}$	=	$\mathcal{P}(u \times \ell)$	$\overrightarrow{R', \ell_A \mapsto \Sigma} = \mathbf{foldL} \mathit{update} \overrightarrow{(\ell_{act}, \Delta)} (R, \vec{A})$
1199	$\ell \in \mathbf{Address}$	=	\mathcal{L}	
1200	$A \in \mathbf{Actor}$	=	$\ell \mapsto \Sigma$	$\overrightarrow{\Sigma', \Delta', \vec{P}} = \mathit{dispatch}_{DS}(\vec{\Sigma})$
1201	$\Sigma \in \mathbf{ActorState}$	=	$\Delta \triangleright B$	
1202	$act \in \mathbf{Action}$	=	$\Delta \mid P$	$\overrightarrow{\Sigma'', \pi} = \mathit{boot}_{\Sigma}(\vec{P})$ with ℓ' fresh in $\vec{\ell}_A$
1203	$B \in \mathbf{Behavior}$	=	(f, v)	
1204	$P \in \mathbf{ProcDesc}$	=	$\mathbf{actor} f v \pi$	
1205	$f \in \mathbf{BehFun}$	=	$\exists State. \Delta \times State \longrightarrow State \times \vec{act}$	

Fig. 16: Dataspace Syntax and Semantics

abstracts over the language for specifying behavior functions so both facet-based code as well as alternatives work. Actions include actor specifications P and patches Δ . An actor specification $\mathbf{actor} f v \pi$ consists of three pieces: a behavior function; its initial state, which is any value from the language for expressing behavior functions; and its initial assertions.

The semantics is defined in terms of a transition relation on dataspace configurations. The relation interprets actor actions and dispatches events in a single rule (**step-par**). It starts by interpreting the queue of pending patches. The *update* metafunction handles a single queue item, updating the dataspace's assertion table and dispatching events. The pending event for each actor aggregates a patch describing relevant updates. Next, the *dispatch*_{DS} metafunction applies the behavior of each actor to its pending patch, producing an updated actor state and some outputs. Any spawn actions are booted, yielding a new address, actor state, and assertions. Finally, all of the patches and boot assertions are collected into a new queue for the next configuration. It includes the updated assertion table and freshly booted actors. Appendix E contains the full definition of these metafunctions.

Definition 3 ($DS ;$). A *run* of a dataspace configuration DS is a potentially infinite sequence of dataspace configurations, where each neighboring pair is related by the transition relation:

$$DS \longrightarrow_{DS} DS_1 \longrightarrow_{DS} DS_2 \longrightarrow_{DS} \dots$$

It is a run that represents the temporal behavior of a dataspace program and its individual actors, and it is thus the foundation for verifying the value of our behavioral types.

Definition 4 (inert_{DS}). A dataspace configuration $\overrightarrow{[(\ell, \Delta); R; \vec{A}]}$ is *inert* iff each patch and pending event is empty.

Theorem 5 (Dataspace Soundness). *Either $\mathit{inert}_{DS}(DS)$ or there exists DS' such that $DS \longrightarrow_{DS} DS'$.*

Proof Garnock-Jones's dissertation proves the soundness of dataspace systems, as well as a number of desirable properties (Garnock-Jones, 2017, theorem 4.17).

7.1 Facet Dataspace Programs

Caldwell et al. (2020) show how to link a semantics for individual actors with the dataspace semantics and prove the soundness of the composite system. Their semantics coincides with ours, so that the soundness of the system follows from theorem 5, because each actor implements the functional interface of a **BehFun**. The $interp_M$ function is the key component connecting the semantics of dataspace with the semantics of facet actors. It implements the dataspace behavior interface **BehFun** for facet machine states, and it applies a pending patch to a machine state via an `inject`-step, then allows the machine to continue via internal transitions until reaching either inertness or an error state. The result of the function is the new machine state together with a patch describing the difference(s) in assertions between the initial and final states. Furthermore, any \overline{Pr} transitions are translated to dataspace specifications P:

$$\begin{aligned}
 &interp_M : \Delta \times M \longrightarrow M \times \overrightarrow{\text{act}} \\
 &interp_M(\Delta, \text{error}) = (\text{error}, \cdot) \\
 &interp_M(\Delta, M) = \begin{cases} (\text{error}, \text{patch}(M, \text{error})) & \text{if } M \xrightarrow{\Delta} M' \xrightarrow{l \bullet}^* \text{error} \\ (M'', \text{patch}(M, M'') \cdot \text{label-to-action}(l \bullet)) & \text{if } M \xrightarrow{\Delta} M' \xrightarrow{l \bullet}^* M'', \text{inert}(M'') \end{cases}
 \end{aligned}$$

The $boot_P$ metafunction translates a facet-level `spawn` to a dataspace process description:

$$\begin{aligned}
 &boot_P : Pr \longrightarrow P \\
 &boot_P(Pr) = \text{actor } interp_M M_{boot} \text{ assertions-of}_M(M_{boot}) \\
 &\text{where } M_{boot} = \begin{cases} M & \text{if } boot_{Pr}(Pr, \emptyset) = M \\ \text{error} & \text{otherwise} \end{cases}
 \end{aligned}$$

7.2 Type-Level Dataspace Programs

Section 6.2 sketches how to lift the transition system on a facet machine to operate on a facet machine type. Similarly, the dataspace transition system above naturally lifts to the type level with only minor adjustments. A type-level dataspace consists of actors communicating in terms of assertion types τ . Facet machine types provide behavior to actors in a type-level dataspace the same way that facet machines give behavior to term-level dataspace actors. We write DS_T to refer to a dataspace configuration with type assertions.

7.3 Example

Figure 17 shows a dataspace configuration based on the smart home example and its evolution. For readability, the example employs a notation different from that of figure 16, giving each actor a name rather than an address and aggregating various other parts of the syntax. Each actor is displayed as a triple, as in $\Delta \triangleright \boxed{nm} \triangleright \pi$ for an actor with name nm . The leftmost element is the current pending event for the actor, while on the right are its current set of assertions. The example treats the behavior and private state of each actor as a black box, displaying them as a box containing the actor's name. For brevity, we use two additional notational shorthands. As a pending event, \cdot represents the empty patch \emptyset/\emptyset , while a single assertion u stands for the patch $\{u\}/\emptyset$.

The configuration contains three actors, representing the control hub, a light with the ID "l1", and a presence sensor with the ID "ps1". Each transition of the relation from

figure 16 is broken into two steps. The first routes assertions among the actors and updates pending events when there is relevant information. The second dispatches each actor on its pending event and updates its current assertions.

The system proceeds in the following manner:

1. Initially, the light actor announces its presence and on/off state, as well as a subscription to a room assignment.
2. The hub actor's interest in `light` assertions yields a corresponding event.
3. In response, the hub actor introduces a room assignment to the "bed" room.
4. Routing results in an event informing the light actor of the assignment.
5. The light actor's reaction is to introduce a subscription for assertions that the "bed" room is `empty`.
6. The subscription matches the existing `empty` assertion, immediately yielding an event for the light actor.
7. The light actor updates its `light` assertion to the `off` state in response to the room being `empty`.
8. The hub actor receives a patch describing the change to the `light` assertion.
9. Finally, the hub actor processes the latest patch without changing its assertions, yielding an inert system.

Figure 18 shows the corresponding type-level system. The primary difference is that the assertions in the dataspace are types rather than values; the names of the actors have also been capitalized to emphasize the difference between the two examples. The transitions follow the same logic as those of figure 17, with one exception. The example encodes the on/off state of the light as a `Bool`. Consequently, the types of assertions in the dataspace do not change when the light switches from on to off in response to an `empty` assertion (type). Thus, the type-level system reaches inertness in fewer transitions than the term-level one.

8 The Predictive Nature of Behavioral Types

Equipped with a semantics of facets and dataspaces, we are now in a position to verify the predictive nature of our behavioral types. Specifically, we wish to show that if a program's behavioral types satisfy a specification, the program itself does, too. We start with a model-appropriate definition of our LTL specifications.

Definition 6 (LTL). $\psi \in \text{LTL} = \tau \mid \bullet\psi \mid \psi \text{ U } \psi \mid \neg\psi \mid \psi \vee \psi$

The atomic propositions are types τ of dataspace assertions. Our basic syntax comes with two temporal operators: $\bullet\psi$ (next) and $\psi_1 \text{ U } \psi_2$ (strong-until). The set of formulas includes negation ($\neg\psi$) and disjunction ($\psi \vee \psi$). **Note** Standard constructs such as conjunction ($\psi \wedge \psi$), eventually ($\diamond\psi$), and always ($\square\psi$) are derived forms.

Formulating the meaning of the LTL specifications takes two steps. The first introduces the set of atomic propositions of a dataspace.

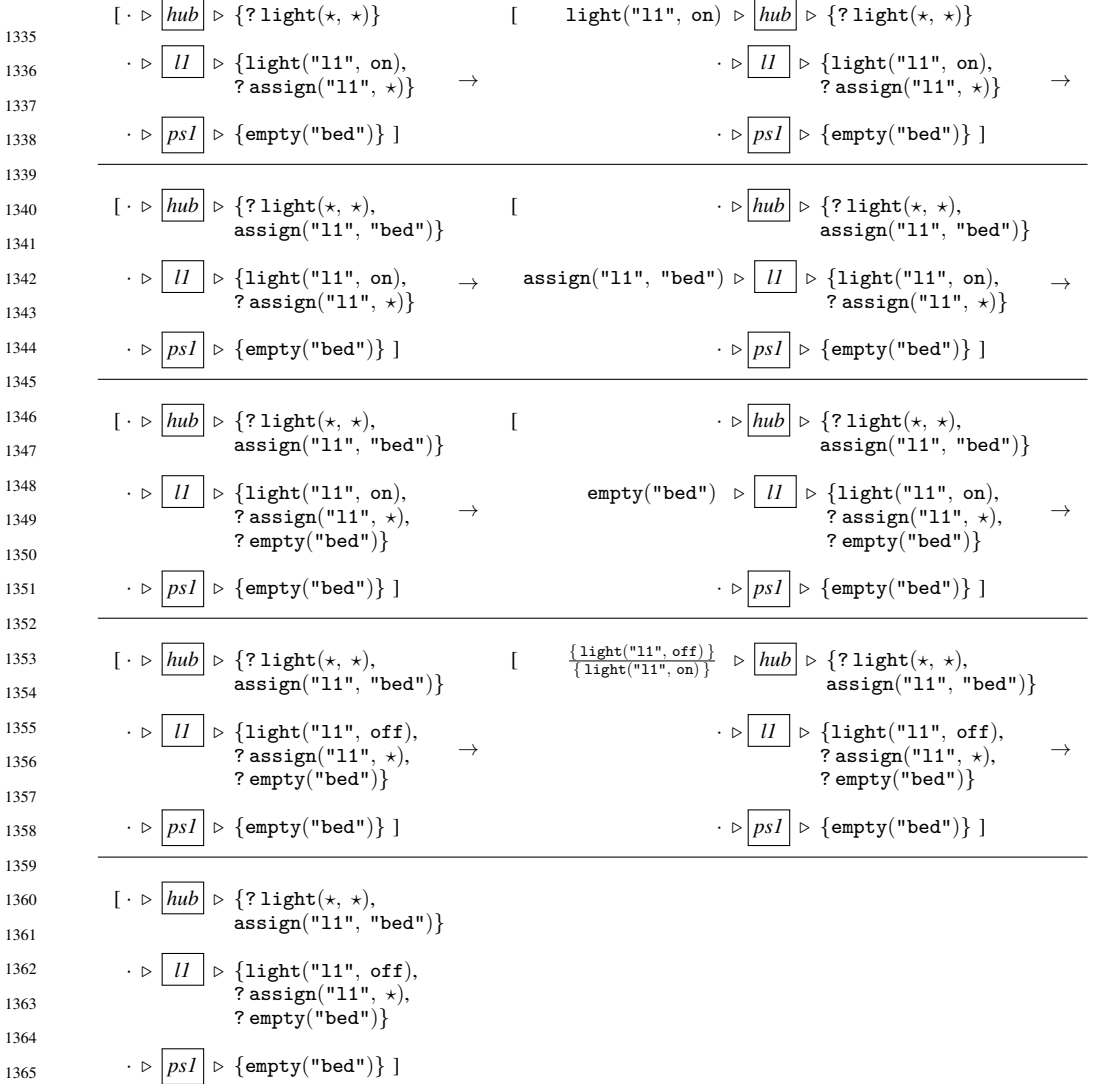


Fig. 17: Example Dataspace Transitions

Definition 7 (AP(DS)).

$$\begin{aligned}
 \tau \in \mathbf{AP}(\langle \overline{(\ell, \Delta)}; \mathbf{R}; \overline{\mathbf{A}} \rangle) & \quad \text{iff there exists } (u, \ell) \text{ s.t. } (u, \ell) \in \mathbf{R}, \vdash_e u : \tau \\
 \tau \in \mathbf{AP}(\langle \overline{(\ell, \Delta_\tau)}; \mathbf{R}_\tau; \overline{\mathbf{A}} \rangle) & \quad \text{iff there exists } \ell \text{ s.t. } (\tau, \ell) \in \mathbf{R}_\tau
 \end{aligned}$$

That is, an atomic proposition belongs to a dataspace configuration, if its table of assertions u contains an assertion (by any actor l) that type checks as τ . We overload the notation for type-level dataspace configurations, $\mathbf{AP}(\mathbf{DS}_\tau)$. At the type level, an atomic proposition τ belongs to the configuration if such a type is present in the configuration's assertions (\mathbf{R}_τ).

The second step concerns the notion of satisfaction, that is, when a dataspace configuration—and its run—satisfies an LTL specification.

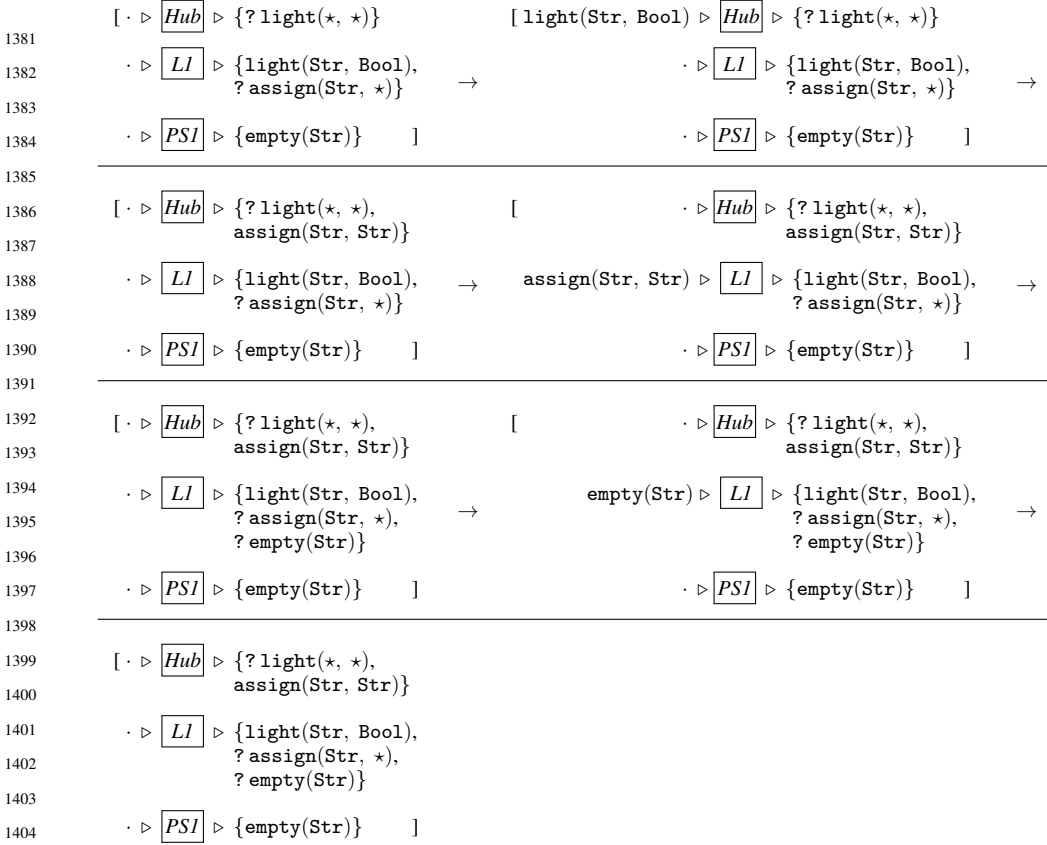


Fig. 18: Example Type Dataspace Transitions

Definition 8 ($DS \models \tau$). For both term and type-level configurations,

$$\begin{aligned}
 DS \models \tau & \quad \text{iff} \quad \tau \in \mathbf{AP}(DS) \\
 DS \models \neg\psi & \quad \text{iff} \quad DS \not\models \psi \\
 DS \models \psi_1 \vee \psi_2 & \quad \text{iff} \quad DS \models \psi_1 \text{ or } DS \models \psi_2 \\
 DS \models \bullet\psi & \quad \text{iff} \quad DS_1 \models \psi \\
 DS \models \psi_1 \mathbf{U} \psi_2 & \quad \text{iff} \quad \text{there exist an } i \text{ s. t. } DS_i \models \psi_2 \text{ and } \forall j, 0 \leq j < i, DS_j \models \psi_1
 \end{aligned}$$

At the term level, atomic propositions hold in dataspace configurations where an actor is making an assertion of that type. At the type level, they hold for configurations where an actor is asserting that type. The temporal constructor $\bullet\psi$ holds in a dataspace configuration when ψ holds starting from the next step in the run. Similarly, the other temporal operator, strong until ($\psi_1 \mathbf{U} \psi_2$), holds when ψ_1 is true for some finite number of steps in a dataspace run, at which point ψ_2 is true. All other syntactic forms have their usual meanings.

At this point, we can define what it means for a program—a description of the initial facet-based actors in the dataspace, or their types—to satisfy a specification.

Definition 9 ($\cdot \models \psi$).

1427

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

1458

1459

1460

1461

1462

1463

1464

1465

1466

1467

1468

1469

1470

1471

1472

$$\begin{aligned} \vec{\text{Pr}} \models \psi & \text{ iff } \text{boot}_{\text{DS}}(\vec{\text{Pr}}) = \text{DS} \wedge \text{DS}_1 \models \psi \\ \vec{\text{T}} \models \psi & \text{ iff } \text{boot}_{\text{DS}_T}(\vec{\text{T}}) = \text{DS}_T \wedge \text{DS}_{T_1} \models \psi \end{aligned}$$

Note how the definitions use the first successor of the initial dataspace configuration to check the desired formula. They ignore the initial one because all of the actors' initial assertions are in the pending action queue, waiting to be interpreted. The first transition step moves these assertions into the configuration's assertion table.

Formula	Term Meaning	Type Meaning
$\text{light}(\text{Bool})$	Either $\text{light}(\text{true})$ or $\text{light}(\text{false})$ is currently asserted in the dataspace.	$\text{light}(\text{Bool})$ asserted.
$\neg \text{light}(\text{Bool})$	Neither $\text{light}(\text{true})$ nor $\text{light}(\text{false})$ is currently asserted in the dataspace.	$\text{light}(\text{Bool})$ not asserted.
$\diamond \text{light}(\text{Bool})$	The execution reaches a $\text{light}(\text{Bool})$ configuration in a finite number of steps.	Same as term meaning.
$\square \text{light}(\text{Bool}) \Rightarrow \bigcup (\text{light}(\text{Bool}) \neg \text{wall-switch-on}())$	Whenever $\text{light}(\text{Bool})$ is true, there is eventually no $\text{wall-switch-on}()$ assertion in the dataspace, and $\text{light}(\text{Bool})$ stays true until then.	Same as term meaning.

Fig. 19: Formal LTL Examples

Figure 19 illustrates how to translate the example specification from section 4.2 into the formal syntax. It also indicates what it means at the term and the type level.

The type system and the behavioral type system properly predict the behavior of actors and programs. Naturally, the first is the basis for the second. That is, all correspondence between terms and types relies on standard soundness (theorem 1). The point of soundness is to confirm that there are distinct classes of actor programs: those that live up to (type) expectations and those that raise exceptions. We collect the first kind in a set $\text{safe}(M)$ for a configuration M that does not error due to the application of a partial primitive, such as division by zero. We write $\text{safe}(\text{Pr})$ if $\text{boot}_{\text{Pr-s}}$ evaluates to a safe machine state.

Next we establish a correspondence between the behavior of a single actor and its type. Specifically, a term-level actor machine is related to a type-level actor machine if they make the same assertion types; the type-level machine types the term-level one; they react to the same (types of) events; and their reactions to those events yield related states.

Definition 10. $M \approx M_T$ if and only if:

- $\vdash_{\pi} \text{assertions-of}_M(M) : \text{assertions-of}_{M_T}(M_T)$
- $\forall l$ such that $M \xrightarrow{l} M'$, $\exists l_{\tau}, \vdash_l l : l_{\tau}$ such that $M_T \xrightarrow{l_{\tau}} M'_T$ and $M' \approx M'_T$
- $\forall l_{\tau}$ such that $M_T \xrightarrow{l_{\tau}} M'_T$, $\exists l, \vdash_l l : l_{\tau}$ such that $M \xrightarrow{l} M'$ and $M' \approx M'_T$

Appendix F provides the straightforward definitions of the \vdash_π and \vdash_l judgments.

Assuming the machine belongs to the set of safe ones, typing implies correspondence.

Theorem 11. *If $\vdash_M M : M_T$ and $\text{safe}(M)$ then $M \approx M_T$.*

Proof The initial machine states for a term and its type are related, as are their assertions. The key meta functions preserve typing, and a suitably related event can always be constructed. See appendix G for the full statement of these properties and their proofs. \square

Theorem 11 establishes a similarity between term and type level machines based on their inputs (injected events) and output labels (spawned actors).

The point of this theorem is that an actor machine and its type implement the same interface, up to typing. A dataspace program consisting of actor machines interacts with its context in the same way as a dataspace program consisting of those actor machines' types.

Finally, we can state the key theorem, namely, that type-level behavior carries over to term-level dataspace programs. That is, if a collection of actor terms \vec{Pr} have types \vec{T} and boots to exception-free dataspace programs, then LTL properties satisfied by the type-level dataspace program also hold for the term-level one.

Theorem 12 (LTL Transference). *If*

- $\overline{\vdash_{Pr} Pr : \vec{T}}$
- $\overline{\text{safe}(\vec{Pr})}$
- $\vec{T} \models \psi$

then $\vec{Pr} \models \psi$

Proof Due to the simulation property (theorem 11), the table of assertions in each step of the term and type level dataspace runs is related by typing. Therefore, the same atomic predicates (**AP**) hold in each, so the same LTL properties apply. \square

Note The assumption of safety for the single-actor machines is standard for the statement of partial correctness. In essence, it is equivalent to the conventional assumption “if the statement terminates properly” in Floyd-Hoare logic (Floyd, 1967; Hoare, 1969).

In detail, our type analysis does not seek to capture the behavior of actors in the case of uncaught exceptions. While dealing with failures is a core principle of the actor and dataspace model, the current type system does not capture this notion.

If we were to account for the possibility of unexpected exceptions, every non-trivial actor would include the behavior of immediately crashing. Every program would then include the situation where all its actors crash and other such degenerate cases, vastly increasing the difficulty of stating LTL properties of the system. Even the addition of supervision actors would not avoid this collapse in reasoning capability.

We conjecture that by generalizing work on exception analysis via type systems (Yi and Ryu, 2002) it might be possible to extend our type system to reason about such cases precisely, at least for some limited class of programs. We leave this problem as future work.

9 Implementation

An implementation of the theoretical design requires three pieces: the language itself (its syntax and semantics); a type checker that derives effect types; and a compilation of these effect types plus the specification to a model-checking system. The following three subsections sketch these three pieces of our implementation.

9.1 Implementing the Behavior of Facet-Oriented Actors

The basic language implementation of facet-oriented actors consists of three layers:

- a syntax layer, which provides the facet notation of the first half of this paper;
- a runtime system, which provides data structures and functions that implement the behavior of facets and endpoints; and
- an imperative dataflow network for tracking changes to fields and scheduling the re-evaluation of dependent computations.

These pieces are built atop the existing implementation of the dataspace model. The dataspace implementation does not require any modifications in order to support these elements, just as the dataspace model predicts (Garnock-Jones and Felleisen, 2016).

The syntax layer essentially turns the surface forms into calls to functions provided by the runtime system. It makes use of `syntax/parse`, Racket’s high-level syntax extension system (Culpepper and Felleisen, 2010, 2012). Implementing the surface language with `syntax/parse` greatly simplifies the addition of actor syntax. Moreover, the resulting interface can be extended and grown as the language develops, and the implementation can benefit from any improvements to the underlying dataspace library. It also facilitates the addition of a type checker.

The run-time system is the most sophisticated of the three layers. The key elements strongly mirror the data structures of the machine and the meta-functions of the formal model of section 5. Technically speaking, the run-time system maintains data structures corresponding to the different parts of machine states M and defines functions for dealing with facets, enqueueing scripts, etc. The implementation differs from the model primarily in its support for efficient re-evaluation of the assertions of an actor as fields are updated.

Finally, the runtime maintains a bipartite, directed dataflow graph for each facet-oriented actor. A source node represents a field, a target node an endpoint, and edges the dependencies of the endpoints on the fields. Each endpoint representation contains a procedure that is used to compute the set of assertions to be associated with the endpoint. By recording field dependencies during the execution of such procedures, the implementation learns which endpoints must have their assertion sets recomputed in response to a field change.

9.2 Implementing the Type Checker

The type checker is implemented with the Turnstile meta-DSL (Chang et al., 2017). Roughly speaking, Turnstile macros supplement `syntax/parse` with a mechanism to check types first and elaborate the given source syntax into target terms later. That is,

each syntactic form is defined as a macro that performs some type checking before elaborating to syntax that implements the run-time behavior. Type information is propagated via metadata on syntax objects.

For example, the `define-typed-syntax` macro for `start-facet` looks like this:

```

1569 (define-typed-syntax (start-facet name:id ep:expr ...+) >>
1570   [[name : FacetName] ⊢ ep >> ep- (⇒ v effs)] ...
1571   #:fail-unless (all-endpoint-effects? #'(effs ...))
1572     "only endpoint installation effects allowed"
1573   -----
1574   [⊢ (install-facet! name (lambda () ep- ...)) (⇒ v (StartFacet name effs ...))])

```

Line by line, this definition reads the following way:

- The macro applies to syntax that matches the `(start-facet id expr ...+)` pattern.⁷ The pattern expects an identifier and a non-empty sequence of endpoint expressions.
- The code between `>>`—pronounced “elaborates to”—and the dashed line analyzes and checks these endpoint forms:
 1. The first clause elaborates each endpoint expression, `ep`, individually in an environment extension that records `name` as the name of a facet. This allows Turnstile to resolve each identifier reference within each `ep` to the proper type.
 2. The result of a successful⁸ elaboration is bound to the name `ep-`, while the types of effects performed by `ep` are bound to the name `effs`.
 3. The trailing ellipse dictates that each `ep` form is analyzed in this manner.
- The `#:fail-unless` clause enforces the side-condition on the corresponding type-checking rule. It makes sure that the body has only endpoint installation effects, such as from the use of `on`, `assert`, and `field`; it disallows the following effect types: `start-facet`, `stop`, and `spawn`. The actual work is left to a helper function: `all-endpoint-effects?`. If the check fails, the checker signals a type error.
- The clause below the dashed line consists of the macro’s two outputs:
 1. The first output is the target expression that implements the behavior of `start-facet`. The produced expression calls the `install-facet!` procedure, which is provided by the run-time support system for facets. The procedure consumes the `name` and a thunk that runs the results of elaborating the endpoint forms—in the specified order. These endpoint forms then elaborate to calls to other procedures from the run-time system, and so on.
 2. The second output attaches a type to the elaboration. Here it describes the effect type of the form: `StartFacet`, with endpoint types described by `effs ...`.

Generally speaking, the `define-typed-syntax` macros are almost verbatim transliterations of the typing rules presented in figure 15.

⁷ For expressions that use the keyword `start-facet` and don’t match the pattern `syntax/parse` raises an error with an informative error message.

⁸ Any type error discovered while checking an `ep` is reported as soon as it is discovered, aborting the rest of the elaboration of `start-facet`.

9.3 Implementing the Model Checker

A programmer initiates a system verification with respect to some specification by adding the following formula to the program:

```
(verify-actors spec actor-type ...)
```

In terms of theorem 12, this formula requests a check of $\vec{T} \models \psi$ where the LTL formula ψ corresponds to `spec`, and T is the series of `actor-types`, one per actor in the dataspace program. If the check succeeds, then the actors of these types jointly realize the specification.

The implementation of `verify-actors` utilizes the model checker SPIN (Holzmann, 1997) to check such uses. That is, it translates each use of `verify-actors` into a Promela program, that is, SPIN's input forms.

More precisely, the translation turns each `actor-type` into a Promela process with a state-machine driven behavior. Each state corresponds to a set of active facets within the actor. An encoding of the dataspace routing algorithm allows processes to react to the appearance and disappearance of assertions, while an additionally generated process performs message dispatching.

SPIN is invoked on this Promela program, plus a translation of the LTL ψ . If verification fails, the trace provided by SPIN is back-translated into the actions of the type-level actors.

Figure 20 displays a sample Promela program. It is the result of translating the `Light` actor's type. The figure annotates the Promela code with in-line explanations. We italicize the annotations to distinguish them from the actual code and comments, which are enclosed within `/* . . . */`. The rest of the Promela process follows the same form.

Blocks of code within the actor (lines 15, 23) execute *atomically* to avoid any unnecessary interleaving with other actor's initialization in the SPIN search space. Additionally, comments within the code, such as those on lines 16 and 30, embed information from the Racket source program; this information is used to back-translate counterexamples of the SPIN model to the source level.

10 Evaluation

The usefulness of this model-checking approach depends on two primary factors:

- Whether the specification language can express important properties of dataspace communication as they arise in actual programs; and
- Whether the implementation can successfully check those properties that are expressible.

This section presents the results of an analysis of these factors in the context of a corpus of facet-based dataspace programs implemented in Racket. Inspecting each program yields a number of behavioral properties important to program correctness. These properties provide the basis of the evaluation.

```

1  Declare an enumeration type for the possible combinations of active facets:
1657 2  mtype = {dur_im_lgt_dur, dur_lgt, dur_lgt_dur, lgt, dur_im_lgt}
3  Create a global variable that controls scheduling:
1658 4  bool light_proc_clock = true;
1659 5  Name a Promela process that is active when the program starts:
6  active proctype light_proc() {
1660 7      Use a local variable to track the actor's currently active facets:
1661 8      mtype current = lgt;
9      Local variables for tracking the assertions known to the actor:
1662 10     bool know_WallSwitchOnT_Symbol = false;
1663 11     bool know_RoomOccupiedT_Symbol = false;
1664 12     bool know_RoomAssignmentT_Symbol_Symbol = false;
13     Issue the actor's initial assertions:
1665 14     atomic {
15         ASSERT(Obs_WallSwitchOnT_Symbol); /*#s(assert #s(Observe ...))*/
1666 15     }
1667 16     Launch an infinite loop for the actor's ongoing behavior:
1668 17     do
19         Wait until the actor is enabled:
1669 20         :: true ->
1670 21             light_proc_clock == GLOBAL_CLOCK;
22             atomic {
1671 23                 light_proc_clock = !light_proc_clock;
1672 24                 Dispatch based on which facets are active:
25                 do
26                     :: current == dur_lgt_dur ->
1674 27                     Dispatch to the event handlers within the active facets:
28                     if
29                         The condition for one of the active event handlers.
1676 30                         A retracted event fires if there are no matching assertions:
1677 31                         :: RETRACTED(RoomAssignmentT_Symbol_Symbol) && /*#s(Retracted ...)*/
32                         and the actor must have seen such a prior assertion:
1678 33                         know_RoomAssignmentT_Symbol_Symbol ->
1679 34                         Perform effects dictated by the type of the event handler's body:
35                         stop a facet by updating current:
36                         current = dur_lgt;
1681 37                         Update the actor's knowledge of the retracted assertion,
38                         and no-longer-relevant ones:
1682 38                         know_RoomAssignmentT_Symbol_Symbol = false;
1683 39                         know_RoomOccupiedT_Symbol = false;
1684 40                         Retract assertions from the stopped facet:
41                         RETRACT(Obs_LightOnCmdT_Symbol); /*#s(retract #s(Observe ...))*/
1685 42                         RETRACT(Obs_RoomOccupiedT_Symbol); /*#s(retract ...)*/
1686 43                         RETRACT(Obs_LightOffCmdT_Symbol); /*#s(retract ...)*/
44                         The beginning of the next event handler:
1687 45                         :: ASSERTED(RoomOccupiedT_Symbol) && /*#s(Asserted ...)*/
1688 46                         !know_RoomOccupiedT_Symbol ->
47                         ...
48                     }
1689 48             }
49     }
1690 49 }

```

Fig. 20: Annotated excerpt of a Promela program for the SPIN model checker

10.1 Corpus

This section describes the programs in the corpus, as well as several behavioral properties for each program that comprise the subject of the evaluation. One property—deadlock-freedom—is relevant to several examples and thus factored into its own section (10.2), which discusses the topic in the context of dataspace programs and our behavioral checks. Each program consists of about 500 lines of code, comprising the concurrent core of a larger system.

Smarthome. This program is an extended version of the earlier examples, including a text-based user interface and temperature sensors and controls.

- *Light Presence.* This property describes the desired interaction between presence sensors and lights in the smarthome program. The specification essentially states that once a presence sensor and light have been installed in a room, the light turns on or off as the user moves in or out of the room, respectively.⁹ Note that while on the surface this property only talks about light and presence sensor actors, the behavior of those actors depends on additional interactions from the user interface, the actor representation of the user, and the hub actor.
- *Steady State Temperature.* This property specifies the desired behavior of the thermostat in the smarthome: when the user sets a desired temperature, the system gradually heats or cools the home until the target is reached, before turning off.

Data Processing. This program implements the core of a streaming data processing framework, inspired by Flink.¹⁰ Clients submit jobs consisting of some number of interdependent tasks. Several actors collaborate to manage the underlying computational resources for executing tasks, tracking intermediate progress, and delivering the end result.

- *Task Delegation.* This property checks the working of three actors that collaborate to perform tasks by treating them as a black box: each assigned task is eventually performed. Performing a task involves delegating it from the job manager to a task manager, and then from a task manager to a task runner, and finally propagating the results back in the opposite direction.
- *Job Completion.* This is another black-box correctness property: when a client submits a job, the job's results eventually become available. When a client submits a job, a job manager actor analyzes the request and computes a directed, acyclic graph of tasks. It then processes the task graph by assigning tasks as they become ready to other actors, before finally announcing the results of the job. A weak variant of this property states that each job either completes or processes tasks in an infinite loop.
- *Load Balancing.* Task-runner actors are capable of performing one task at a time. A task manager actor monitors the status of task performers and assigns them tasks when they are idle. Likewise, each task manager has a capacity based on the number of task runners it manages. The job manager assigns tasks to task managers based on their current free capacity. This property states that the task manager never assigns a task to a busy runner and the job runner never assigns a task to a task manager that is at capacity. A weak variant of the property pertains to the scenario where each task manager oversees exactly one task runner actor.

Caucus. This program represents a geographically distributed, iterative election in the style of a caucus, inspired by the “two-buyer problem” from the behavioral types literature (Honda et al., 2008). A distinct actor represents each voter and candidate. After a

⁹ This property is a response to a bug uncovered through manual testing, where a light installed in a room after a presence sensor sometimes did not properly react to the initial presence information.

¹⁰ <https://flink.apache.org>

1749 registration period, voting commences across a collection of regions. Within each region,
 1750 voting proceeds in rounds until a single candidate receives a majority share of the vote.
 1751 Once every region has reached a decision, the winner of the most regions is pronounced the
 1752 winner of the election. In each region an actor guards against various forms of misbehavior
 1753 from voters and candidates, such as voting twice.

- 1754
1755
1756
- *Resolution*. The election completes. A weak variant of the property states that either the election completes or holds infinitely many rounds of voting.
 - *Candidate (Mis)Behavior*. This property is the specification for a well-behaved candidate actor in the election. A well-behaved candidate actor announces its candidacy as an assertion, and maintains it until either the election is over or an election agent informs it that it is no longer in the running.
 - *Voter (Mis)Behavior*. This property is the specification for a well-behaved voter actor. A well-behaved voter actor registers in exactly one region, and then participates in each round of voting in that region by casting exactly one vote for one of the eligible candidates. A weak variant of the property states that the voter registers and always votes at least once in each round.
- 1767

1768 **Windowing System.** This program implements a basic graphical windowing system. A
 1769 collection of driver actors collaborate to provide the primary interface. A layout-solving
 1770 actor powers graphical output by combining requested sizes and layout styles (such as
 1771 vertical or tabular) with the actual dimensions of the system window to compute the actual
 1772 size and location of each item. Another interface actor provides descriptions of mouse
 1773 events, including when the mouse is touching a window, mouse presses, and releases. The
 1774 mouse interface allows for an actor *mix-in* (a function that starts facets or endpoints) to
 1775 implement drag-and-drop behavior, which can then be instantiated and reused freely. The
 1776 rest of the program consists of actors implementing rudimentary windowed applications
 1777 and menus.

- 1778
1779
1780
1781
- *Layout*. This property describes the behavior of the layout engine, which spawns solvers for horizontal, vertical, and tabular layouts on demand. The specification is based on a comment in the original, untyped program, describing the desired behavior; it essentially demands that requests for layouts are answered with solutions:

```
1786     ;;      (Observe LayoutSolution)+ ==>
1787     ;;      RequestedLayoutSize ==>
1788     ;;      ComputedLayoutSize ^ LayoutSolution+
```

1789 This comment closely corresponds to the specification language of assertion types
 1790 and LTL connectives provided by the implementation.

- *Menu Duration*. This property checks the lifetime of menu items: that they appear in response to selecting a menu, and that they remain until either a selection is made or a mouse click occurs, either selecting a menu item or closing the menu.
- 1793
1794

Web Chat. This program implements the server for a chat service loosely based on Slack.¹¹ It allows a user to connect, sign up for an account, and then create and join conversations. Connected users may request to follow one another. Accepted follow requests lead to each user being added to the contact list of the other. Users may join the conversations of their contacts. The system implements an option for permission delegation, so that if user A chooses to invite user B to a conversation, user B may then invite user C, and so on.

- *Conversation Release.* This property is based on a comment in the original, untyped program, essentially calling for the release of resources in the event that a request is canceled before the response has materialized:

```
;; TODO: CHECK THE FOLLOWING: When the 'invitation' vanishes (due
;; to satisfaction or rejection), this should remove the question
;; from all eligible answerers at once
(during (invitation $cid $inviter $invitee)
...)
```

A weak version of the property relaxes the need for the resources to all be removed “at once.”

- *Contact Release.* This property is also based on a similar comment in the code to the conversation release property, but in the module for managing the contacts list for users:

```
;; TODO: CHECK THE FOLLOWING: When the 'permission-request'
;; vanishes (due to satisfaction or rejection), this
;; should remove the question from all eligible answerers
;; at once
(during (permission-request $who $grantee ($ p (p:follow _)))
...)
```

A weak version of the property relaxes the need for the resources to all be removed “at once.”

10.2 Deadlock Freedom

Dataspaces implement a form of asynchronous message-passing, so in a technical sense every dataspaces program is free from deadlocks. However, dataspaces programs may reach a stuck state where each actor waits for further communication before continuing. Such a situation is sometimes referred to as a soft deadlock. The corpus includes several programs that feature numerous kinds of actors participating in interwoven conversations, making freedom from such soft deadlocks a desirable property to check.

Figure 21 demonstrates a simple program simulating two friends attempting to make plans. One friend actor, once it knows the time of the meeting, is ready to suggest a location. Meanwhile, the other friend actor, is ready to suggest a time depending on the location. The result is no communication at all. Such situations can arise from poorly designed protocols and buggy implementations. Figure 21 is a case of the former.

One way of describing soft deadlocks is that an actor states an interest, that interest is never withdrawn (that is, the interest remains relevant to the actor), and no matching assertion ever arises. The following LTL formula states this property with respect to the

¹¹ <https://slack.com>


```

1 (define (friend1)
1841 2 (spawn
3 (start-facet f1
1842 4 (on (asserted (time $t))
1843 5 (start-facet (assert (place "yours"))))))))
6
1844 7 (define (friend2)
1845 8 (spawn
9 (start-facet f2
1846 10 (on (asserted (place $p))
1847 11 (start-facet (assert (time 0))))))
12
1848 13 (run-dataspace (friend1) (friend2))
1849

```

Fig. 21: Soft Deadlocked Actors

friend1 actor's interest in Time assertions:

```

1854 (define-ltl friend1-deadlock
1855 (Always (Implies ?(TimeT *)
1856 (Eventually (Or Time
1857 (Not ?(TimeT *)))))))

```

Checking this LTL property against the implementations of the `friend1` and `friend2` actors fails, as should be expected for figure 21. The resulting trace illustrates the two actors waiting for one another with active subscriptions but no matching assertions.

A similar property describes the deadlock from the perspective of the `friend2` actor's interest in Place assertions, which fails a similar check:

```

1863 (define-ltl friend2-deadlock
1864 (Always (Implies ?(PlaceT *)
1865 (Eventually (Or Place
1866 (Not ?(PlaceT *)))))))

```

Generalizing, a program may be checked for deadlocks by taking each assertion of interest $?\tau$ and checking the following LTL formula:

```

1868 (Always (Implies ?\tau
1869 (Eventually (Or \tau
1870 (Not ?\tau))))))

```

Such a check succeeds if the model checker can show that the program *never* deadlocks. This property may be too strong for the level of type precision. An alternative approach is to state in LTL that the program *definitely* deadlocks, and see if model checking produces a counterexample, implying that the program has some non-deadlocking executions:

```

1876 (Eventually (And (Always ?\tau)
1877 (Not (Eventually \tau))))

```

This section refers to the former property as *strong* (soft) deadlock freedom and the latter as *weak* (soft) deadlock freedom.

10.3 Results

Table 1 summarizes the results. Each row of the table describes a property from the corresponding program. A ✓ in the “Expressible” column indicates that the implementation can express the property and a corresponding ✓ in the “Checkable” column indicates that the

implementation successfully checked the property. Otherwise, the process of stating and checking could not be completed. Section 10.4 discusses these cases.

Program	Property	Expressible	Checkable
Smarthome	Light Presence	✓	✓
	Steady State Temperature	✓	✓
	Deadlock Freedom	✓	✓
Data Processing	Task Delegation	✓	✓
	Job Completion	✓	
	Job Completion (Weak)	✓	✓
	Load Balancing		
	Load Balancing (weak)	✓	✓
	Deadlock Freedom	✓	✓
Caucus	Resolution	✓	
	Resolution (Weak)	✓	✓
	Cand. Misbehavior	✓	✓
	Voter Misbehavior		
	Voter Misbehavior (weak)	✓	✓
	Deadlock Freedom	✓	✓
Windowing System	Layout	✓	✓
	Menu Duration	✓	✓
Web Chat	Conv. Release		
	Conv. Release (weak)	✓	✓
	Cont. Release		
	Cont. Release (weak)	✓	✓

Table 1: Evaluation Results

10.3.1 Performance

One possibility is that, even if the implementation *can* express and check a desired property, the result may not return within a reasonable amount of time. The potential for a program’s state space to grow exponentially is unavoidable. Eventually, even the most optimized checkers will require enormous amounts of memory and face slowdowns. Determining a threshold for a “reasonable” amount of time is tricky, as different developers have different conceptions. This evaluation sets the upper bound for a check that could still be useful at eight hours. That is a long time, but short enough to be a part of continuous integration (CI) tests that run on a nightly basis, with the results ready for review the next morning. Table 2 details the running time and peak memory usage for each successfully checked property.

Program	Property	Time (s)	Mem. (MB)
Smarthome	Light Presence	4.6	285
	Steady State Temperature	14.3	338
	Deadlock Freedom	131.2	333
Data Processing	Task Delegation	2.6	382
	Job Completion (weak)	2.9	467
	Load Balancing (weak)	2.7	467
	Deadlock Freedom	11.2	382
Caucus	Resolution (Weak)	1.6	620
	Cand. Misbehavior	1.7	626
	Voter Misbehavior (weak)	1.8	626
	Deadlock Freedom	6.9	624
Windowing System	Layout	11.0	894
	Menu Duration	11.1	898
Web Chat	Conv. Release (weak)	8.5	1252
	Cont. Release (weak)	1.7	251

Table 2: Running time (seconds) and memory usage (megabytes) of each completed check. The checks ran on a laptop with an Apple M1 Pro CPU, 32GB of memory, and MacOS Ventura 13.5 installed.

The performance clears the “integration test” threshold by a wide margin. All of the properties take seconds or a few minutes rather than hours to check. The property that takes the longest to check, deadlock freedom for the smarthome program, consists of 22 individual properties, one for each type of subscription in the program. Each sub-property leads to a separate invocation of the model checker, one for each type of interest in the program. These results put the performance in the realm of “unit test” acceptability, where the checks can be run along with a unit test suite before committing each change to a project.

10.4 Interpretation and Discussion

This section discusses each of the properties that the implementation failed to support and the primary reason(s). In a few cases, interesting details for successfully checked properties are provided as well.

Data Processing.

- *Job Completion.* The implementation is not able to check this property. The checker is not able to reason about the termination of the task-processing loop in the protocol. Even though the loop follows an inductive structure—processing a DAG—this information is not propagated to the type level. Moreover, since an LTL formula always talks about all possible executions of a system, the specification language cannot express that it is at least possible to reach the desired end state.
- *Load Balancing.* The implementation is unable to state and check the property. For a task manager with an arbitrary number of associated task runners, stating the specification requires stating relations between numbers. Checking the property would similarly require reasoning about arithmetic operations. Moreover, the property inherently deals with the multiplicity of assertions of a certain types, namely task-assignment assertions. Violating the property means having more than the expected number of such assertions.
Checking the weak variant of the specification makes use of a coincidence in the program for working around the inability to reason about the multiplicity of assertions. Because there happen to exist two distinct types of tasks, a limited form of multiplicity can be checked—namely, if there are some of one type of task assigned to a manager, there are none of the other, and vice versa.

Caucus.

- *Resolution.* Checking this property also encounters difficulty due to the potential non-termination of a loop in the protocol. The implementation is unable to reason about the termination of the loop, even though the loop matches a simple inductive structure (one candidate is removed each round).
- *Voter (Mis)Behavior.* The implementation lacks the precision to state the full specification. The primary impediments are the inability to correlate information from one assertion to another, such as the name of the candidate the voter chooses with the names of the candidates on the ballot, as well as reason about the multiplicity of assertions. However, the weak specification is still able to distinguish a correct implementation of a voter from several malicious ones.
- *Deadlock Freedom.* While section 10.2 discusses deadlock freedom, one interesting point to note is that this check uncovered a bug in the program: at the end of each round of voting, a region manager actor announces intermediary results that inform the candidates of their status in the race, except for the final round of voting, the manager actor simply announces the winner. Due to an oversight, the original implementation of the candidate actors only listened for the intermediate results, not the winner announcement. As a result, they were soft deadlocked waiting for an intermediate update that never materializes. The fix to the bug is to have the candidate actor listen and react to the final election result.

Web Chat.

- *Conversation Release*. This property is checkable with the implementation, with one caveat: the phrase “at once” entails a level of precision that is beyond the implementation. That is, it can check that the required assertions disappear, but it does not have a way of expressing that the concerned parties, “all eligible answerers,” all react within a set timeframe. However, that part of the specification is of dubious importance when scrutinized: dataspace routing will always dispatch an event to all interested actors when the question is withdrawn. Each such actor will have the opportunity to react to the event in due time. Since dataspace scheduling is fair, there is little reason to worry about such timing details. Thus, the full specification is beyond the implementation, but the weak version, without the timing constraint, is checkable.
- *Contact Release*. Just as with the conversation release property, the “at once” phrasing creates difficulty but is of questionable importance. So again, the full specification is not checkable but a weak variant without the timing constraint is.

Threats to Validity. There are several threats to validity of this evaluation. The corpus comprises a small number of programs, authored by only three individuals. The programs for the corpus were not selected randomly. Rather, we sought programs whose implementations involved a high degree of concurrency and communication, as opposed to business logic or other concerns. Moreover, the identification of properties and subsequent verification efforts were conducted by the authors of the tools. We are highly familiar with both the capabilities and limitations, potentially biasing the search for properties that are a good match for the current capabilities.

11 Related Work

The design of the facet language has two distinct goals. First, the notation greatly facilitates the expression of the temporal relationships with respect to —their conversations with— other actors. See appendix A for an extended comparison of facets to an implementation of the same behavior in a functional style. Second, in contrast to higher-order functional programming of actors, the first-order character of the facet language allows the construction of reasoning tools that assist programmers with the validation of essential protocol properties. That is, the facet language forms capture the input/output behavior of an actor without the need for additional data- or control-flow analysis. See section 10 for a basic illustration. Our results are thus related to two distinct bodies of research: design of concurrent languages (section 11.1) and reasoning with automated behavioral types (section 11.2). The dissertations of Garnock-Jones (Garnock-Jones, 2017) and Caldwell (Caldwell, 2023) also feature extended discussions of related languages and analysis tools.

11.1 Facets and the Design of Concurrent Languages

Programming languages provide limited support for organizing conversations among groups of actors and their internals. Garnock-Jones et al. (2014, 2016) compare dataspace communication with related coordination mechanisms such as the Conversation

2071 Calculus (Vieira et al., 2008), the Mobile Ambient Calculus (Cardelli and Gordon, 2000),
 2072 the join calculus (Fournet and Gonthier, 2000), and tuplespaces (Carriero et al., 1994;
 2073 Murphy et al., 2006), as well as various actor systems and process calculi. Here, we focus
 2074 on linguistic features for organizing the code of individual actors

2075 *State-Machine Actors.* Actor systems such as Akka (Akka Project, 2022) and
 2076 Erlang/OTP (Armstrong, 2003) provide means of organizing actors as state machines.
 2077 Erlang’s `gen_statem` interface exemplifies these mechanisms. The programmer describes
 2078 each state in the machine as a procedure mapping an incoming message to the next state
 2079 and some actions to carry out. These state transition functions also operate on the actor’s
 2080 private store, handled separately from the state of the machine.

2081 Actors organized as event-driven state machines lose access to contextual informa-
 2082 tion. Instead, such notations force programmers to encode context into each state and
 2083 save related information explicitly in the private store. Such encodings, however, make
 2084 it cumbersome to create actors that simultaneously participate in multiple conversations or
 2085 alternatively engage and disengage in multiple intertwined behaviors, such as the Hub and
 2086 Light actors from section 3.

2087 The `gen_statem` interface and its sibling `gen_server` address only one of the concerns
 2088 of facets in a limited fashion: abstracting control state. Language support ends at instan-
 2089 tiating the interface with callbacks. An actor that deals with multiple types of messages
 2090 in a single state must often manually demultiplex the message to match it with the proper
 2091 response. Callbacks use the familiar state-passing style, operating on a monolithic store.
 2092 Though callbacks may coordinate startup and shutdown of the actor, initiating and closing
 2093 a conversation comes without linguistic support.

2094 *Fact Spaces and CRIME.* The inspirational *fact spaces model* (Mostinckx et al.,
 2095 2007) shares many similarities with the dataspace model. Fact spaces build on
 2096 TOTAM (Scholliers et al., 2009, 2010; González Boix et al., 2014), a form of tuplespace,
 2097 to equip actors with a means of moving and sharing state. In fact spaces, programs react
 2098 to both the appearance and disappearance of facts from a shared repository. Reactions are
 2099 described in a logic coordination language, allowing the computation of new facts based
 2100 on current facts via forward-chaining. The language implementation (implicitly) records
 2101 the dependencies between facts and invokes application actions specified by logic rules.

2102 The implementation of the fact space model, dubbed CRIME (Mostinckx et al., 2008),
 2103 integrates the fact spaces model with the AmbientTalk language (Van Cutsem et al., 2007,
 2104 2014). AmbientTalk is an instance of the Actor model (Agha, 1986; De Koster et al., 2016)
 2105 in the mold of E (Miller et al., 2005). In E and AmbientTalk, objects are organized into *vats*;
 2106 each vat runs its own event loop, dispatching incoming messages to the contained objects.
 2107 Combining AmbientTalk and CRIME requires bridging the gap between the events corre-
 2108 sponding to assertion and retraction of facts and the message exchange of AmbientTalk.
 2109 The solution incorporates reactive programming, in the mold of FrTime (Cooper and
 2110 Krishnamurthi, 2006). The result is that actors may define time-varying *behaviors*, with
 2111 values shifting based on the available facts.

2112 The differences between CRIME and facet-oriented dataspace actors stem from the
 2113 absence of a unifying notion of a conversational frame. CRIME comes without any rep-
 2114 resentation of a conversation, hence programs lack the conversational structure of (nested)
 2115 facets. Time-varying collections do not offer a way of propagating changes to the tuples
 2116

in the shared space, unlike the fields and query forms that connect to an actor’s endpoints. Because CRIME does not group the components of conversational behavior, no automatic support exists for the release of associated resources. Programmers must carefully reason about the relationships between components to ensure that the state of the actor and associated tuples remain consistent. The underlying E language offers object references to denote specific conversations within the heap of a given actor and employs method dispatch as a limited pattern matcher over received messages.

Concurrency Within Active Objects. Some Active Objects languages allow for concurrency *within* an Actor or Actor-like entity (Schäfer and Poetzsch-Heffter, 2008, 2010; de Boer et al., 2007). Typically, this concurrency is achieved by allowing multiple message dispatches to be active within a single active object. This allows for splitting control and state among numerous objects, resembling the way our facets and fields individually contribute to the actor’s overall behavior.

Much like with conventional object-oriented programming, communication between active objects forms a graph structure. Each active object is a node in the graph and has an edge to each object it communicates with via asynchronous method invocation. Indeed, many active object languages support both the communication-level graph and the traditional object-oriented “refers to” relationship with support for passive objects (Boer et al., 2017). Each active object (node in the communication graph) may implement its behavior using any number of such passive objects following traditional object-oriented design (Gamma et al., 1994).

Dataspaces and faceted actors possess the same graph-like communication structure. Each actor (node) is connected to each other actor that it communicates with via assertions (edge). Thinking in terms of the graph abstracts away the differences between the message-passing paradigms employed by active objects (point-to-point) and dataspaces (publish/subscribe). Just as how an active object may implement its behavior using any number of passive objects with their own fields and method definitions, a facet-based actor defines its behavior across a collection of facets, fields, and endpoints. The primary difference is that passive objects form a graph, while facets form a (dynamically changing) tree. The tree structure of facets enables the ability to program directly to the concerns of startup and, especially, shutdown (via the stop statement and on stop event handlers).

One interpretation of this analysis is that it may be fruitful to adapt the facet notation to the domain of active objects. Combining facets for defining *behavior* with passive objects for defining *state* has the potential to synthesize the best of both worlds.

ConGoLog. While communication-by-assertion in dataspaces is inspired by Prolog fact databases, ConGoLog (De Giacomo et al., 2000) is a more direct adaptation of logic programming to concurrency. In ConGoLog, a programmer first models a particular domain—typical applications include robotic control systems—with a collection of axioms describing relevant properties and the effects of performing particular actions. The main program then specifies some number of concurrent agents—as logic programs—emitting effects, each of which may update the stated properties, and reacting to the truthiness of such properties.

ConGoLog and dataspaces share a deep connection but also differ in significant ways with respect to their conceptions of state and agent. As a variant of the actor model, dataspace provide each actor with its own isolated memory, while the dataspace itself contains

information shared with the rest of the program. Information shared in the dataspace is tied to the identity of some particular actor, imposing a rigid structure on updates and a connection to component failure. Meanwhile, ConGoLog lacks the distinction between local state of an agent and globally visible state information of the situation. Concurrent ConGoLog agents come without the structure provided by facets for grouping related behavior and state, as well as orchestrating startup, shutdown, and component evolution—partly because these concerns are not a goal of the design.

Sparrow. The Sparrow DSL (Avila et al., 2020) extends Erlang with the ability to react to complex event patterns. That is, rather than just handling one message at a time from its mailbox, a programmer may utilize Sparrow to specify an actor’s behavior dependent on particular combinations of messages, timing constraints, and so on. Reactions to a given pattern may be dynamically added to and removed from an actor’s behavior. A reaction is like a facet with a single event-handler endpoint. They lack the other features of our language, like the hierarchical structure, start-up and shut-down behavior, and grouping of facets, as well as localized state of fields.

Dataflow. The simple dataflow system described here is most similar to the simple dependency tracking approach to object-oriented reactive programming described by Salvaneschi and Mezini (2014, section 2.3), and was in fact directly inspired by the dependency tracking of JavaScript frameworks such as Knockout¹² (Sanderson, 2010) and Meteor.¹³

11.2 Behavioral Type Systems

In general, the design and theory of a type system is intimately tied to the underlying language. This is doubly true of behavioral type systems, which express control-flow aspects of the program, not just the static relationships found in purely structural type systems. Consequently, applying ideas around the actor model to dataspace and facet-oriented actors poses a serious challenge. Dataspace provide a form of publish/subscribe communication (Eugster et al., 2003) that has received little attention in the context of behavioral type systems. While encoding dataspace in a traditional message-passing (or channel-based) setting is possible, the encoding would obscure a program’s communication patterns too much for the targeted behavioral type system to be of use. Application-specific information could potentially be recovered by utilizing more powerful type-level reasoning, such as dependent session types (Toninho et al., 2011), but presently bringing such machinery to bear in a usable programming language remains unsolved.

Our work specifically draws inspiration from studying the approach to type checking more so than the specifics of the systems. In that sense, our design follows that of Igarashi and Kobayashi (2001) and Chaki et al. (2002), in which the type checker constructs type-level processes as abstractions for term-level behavior. The type-level processes serve as the basis for behavioral analysis. In the case of Igarashi and Kobayashi (2001), the checks are generic. Subtyping and other parts of the type checker may be instantiated to yield checks for races, deadlocks, etc. Chaki et al. (2002) use the SPIN model checker

¹² <http://knockoutjs.com/>

¹³ <https://docs.meteor.com/api/tracker.html>

2209 for simulation-based subtyping and analyzing conformance to LTL specifications. In their
2210 case, LTL formulae state properties of channels in the system.

2211 *Effpi*. The Effpi message-passing framework (Scalas et al., 2019) reflects core pro-
2212 cess and communication operations to the type level and uses dependent function types
2213 to precisely track which channels are used. By model checking these type-level descrip-
2214 tions, it becomes possible to verify certain properties stated in a temporal logic, including
2215 deadlock-freedom and some communication patterns. By contrast, we seek to verify any
2216 LTL formula, allowing application-specific correctness properties. The biggest difference
2217 to our work is the communication paradigm. While Effpi uses message-passing along
2218 channels in a process calculus, our actors share knowledge via a dataspace.

2219 *Conversation Types*. Conversation types (Caires and Vieira, 2009) add behavioral check-
2220 ing to the Conversation Calculus (Vieira et al., 2008). A conversation type describes
2221 a sequence of message exchanges within a particular context, i.e., conversation. Like
2222 the global types of Honda et al. (2008), a conversation type may be decomposed to
2223 type(s) describing the actions of the individual processes participating in the conversa-
2224 tion. Crucially, the decomposition is flexible. A given conversation type may be realized as
2225 the composition of numerous different combinations of participant types. This flexibility
2226 meshes well with the anonymous communication in the conversation calculus, allowing a
2227 degree of agnosticism with regard to the number of participants in a conversation and their
2228 individual roles when looking from the global perspective. At first glance, this approach
2229 may work for dataspace actors, but because every communication is between a single
2230 sender and a single receiver, the model remains similar to channel-based models rather
2231 than publish/subscribe communication. In addition to structural message safety, their types
2232 ensure a degree of deadlock-freedom.

2233 *Active Object Languages*. There are several notable behavioral type systems and efforts
2234 to perform static verification on active object languages. Behavioral type systems for
2235 active objects tend to focus on the problem of detecting deadlocks (Henrio et al., 2017).
2236 They employ a similar types-as-processes technique, capturing key information, such as
2237 dependencies between futures, as effect types. Proof rules (i.e. type checking) can then
2238 determine whether the program may deadlock. Our behavioral checks of facet programs
2239 instead treat effect types as abstracted, but still executable, processes, with the goal of both
2240 (model)checking properties of programs and validating the design of the DSL.

2241 The Rebeca modeling language (Sirjani et al., 2004) is the closest to our own behavioral
2242 checking. With Rebeca, a programmer may use communication and abstraction facili-
2243 ties of active objects to define a modes of their program or system. The Rebeca model
2244 checker, Modere (Jaghoori et al., 2006), can then verify properties of the model specified as
2245 temporal logic formulae. Rebeca has numerous extensions, notably broadcasting (Yousefi
2246 et al., 2015). The goal of our work is to check properties of programs versus models of
2247 programs and to thus validate the design of the facet DSL. It might be possible to tar-
2248 get Rebeca as a backend for the implementation, rather than SPIN, and benefit from its
2249 message-passing-specific optimizations.

2250 *ConGoLog*. In bounded contexts, it is possible to decide some temporal properties of
2251 ConGoLog programs (De Giacomo et al., 2016). Thus, it ought to be possible to develop a
2252 model checker based on these decision procedures that could express and check temporal
2253

2254

properties similar to the ones that we explore for dataspace programs. No such model-checker implementation appears in the literature.

Types for the Join Calculus. The join calculus (Fournet and Gonthier, 2000) shares some similarities with dataspace actors. It has a soup of messages versus the table of assertions in a dataspace. The event handler endpoints of facets resemble join patterns, especially in the objective join calculus (Fournet et al., 2000). Recent work (Crafa and Padovani, 2017) has developed behavioral types for the join calculus around the notion of type state (Strom and Yemini, 1986). They describe the types of messages understood by an object and use connectives such as conjunction, disjunction, and repetition to determine how a reference to an object can and must be used. The resulting design is able to elegantly track the dynamic interface of e.g. a lock object. Unfortunately, they note the challenge of implementing such connectives, as well as the necessary substructural support.

Mailbox Types. Similarly to the behavioral type systems for the Join Calculus, mailbox types (de'Liguoro and Padovani, 2018) use a commutative regular expression of messages to describe the inputs to an actor. Though the description of messages is unordered, sequencing is still expressible using mailbox-passing. Following a familiar technique in behavioral systems, type analysis collects information from the program as the basis for analysis. In this case it is a dependency graph among mailboxes. The dependency graphs are vital for detecting and preventing programs that may deadlock.

Session Types. Type systems for the π -calculus and related process calculi have been widely studied. Session type systems (Honda et al., 1998) in particular have been utilized to describe and verify the communication properties of many kinds of systems. Our methods are closely related to the generic π -calculus type system of Igarashi and Kobayashi (2001). They relate processes with abstract process types, including the primary process constructors: channel send and receive, parallel composition, and so on. Their system is parametric over constraints on process types, allowing different instantiations targeted for different properties, such as deadlock prevention.

Multiparty session types (Honda et al., 2008; Scalas and Yoshida, 2019) come closer to describing the group oriented communications between dataspace actors. Global types provide a perspective for describing the protocol among a group of participants. Though efforts towards adapting session type theories to actor systems have found some success (Mostrous and Vasconcelos, 2011; Crafa, 2012; Neykova and Yoshida, 2014), much work remains to be done to handle the full spectrum of actor programming and its relatives, such as the dataspace model.

12 Conclusion

The notation of facets for dataspace actors is good for concurrent actor programming. Our examples show how the notation directly addresses the most common concerns: engaging and disengaging in behavior, managing local state and resources, sharing information during a conversation, etc. Furthermore, the language design captures the key components of behavior, allowing a types-as-processes approach to analyzing temporal behavior.

The facet notation is also good for reasoning about programs. A type system can ensure basic soundness and provides the basis of behavioral verification. The evaluation on realistic programs illustrates the system’s usefulness. Our formal work proves that this approach is well-founded; our evaluation supplies initial evidence of the practicality of the approach.

The work itself suggests several directions for future improvements. Concerning the facet notation, the design could benefit from Avila et al. (2020)’s language of communication patterns. Concerning the behavioral type system, precision could be improved by using a dependency notation for correlation information in assertions similar to the work of Scalas et al. (2019) for channel names. Finally, an additional validation could come from porting our work to a widely used actor system, such as Akka. The broadened context would introduce additional research challenges, such as accounting for the possibility of communication failures, as well as shed additional light on the usefulness and usability of our design.

Acknowledgments We would like to thank Mitch Gamburg for his assistance in developing programs for the evaluation.

Financial Support This work was partially supported by NSF grant SHF 2315884.

Conflicts of Interest None.

References

- Agha, G. (1986) *Actors: a Model of Concurrent Computation in Distributed Systems*. MIT Press. Massachusetts.
- Akka Project. (2022) Akka actors. <https://akka.io/>. Accessed: 2022-06-24.
- Armstrong, J. (2003) *Making Reliable Distributed Systems in the Presence of Software Errors*. PhD dissertation. Royal Institute of Technology, Stockholm.
- Avila, H. R., De Koster, J. & De Meuter, W. (2020) Advanced join patterns for the actor model based on CEP techniques. *The Art, Science, and Engineering of Programming*. **5**.
- Berry, G. & Gonthier, G. (1992) The Esterel synchronous programming language: design, semantics, implementation. *Science of Computer Programming*. **19**(2).
- Boer, F. d., Serbanescu, V., Hähnle, R., Henrio, L., Rochas, J., Din, C. C., Johnsen, E. B., Sirjani, M., Khamespanah, E., Fernandez-Reyes, K. & Yang, A. M. (2017) A survey of active object languages. *ACM Computing Surveys*. **50**(5).
- Caires, L. & Vieira, H. T. (2009) Conversation types. In *Proceedings of the European Symposium on Programming (ESOP)*.
- Caldwell, S. (2023) *Reasoning About Actors That Share State*. PhD dissertation. Northeastern University.
- Caldwell, S., Garnock-Jones, T. & Felleisen, M. (2020) Typed dataspace actors. *Journal of Functional Programming*. **30**, e27.
- Carbone, M. & Montesi, F. (2013) Deadlock-freedom-by-design: multiparty asynchronous global programming. Proceedings of the ACM SIGPLAN Symposium on Principles of Programming Languages (POPL).
- Cardelli, L. & Gordon, A. D. (2000) Mobile ambients. *Theoretical Computer Science*. **240**(1).
- Carriero, N. J., Gelernter, D., Mattson, T. G. & Sherman, A. H. (1994) The Linda alternative to message-passing systems. *Parallel Computing*.
- Chaki, S., Rajamani, S. K. & Rehof, J. (2002) Types as models: model checking message-passing programs. In *Proceedings of the ACM SIGPLAN Symposium on Principles of Programming Languages (POPL)*.

- 2347 Chang, S., Knauth, A. & Greenman, B. (2017) Type systems as macros. In *Proceedings of the ACM*
 2348 *SIGPLAN Symposium on Principles of Programming Languages (POPL)*.
- 2349 Clocksin, W. & Mellish, C. (1981) *Programming in Prolog*. Springer.
- 2350 Cooper, G. H. & Krishnamurthi, S. (2006) Embedding dynamic dataflow in a call-by-value language.
 2351 In *Proceedings of the European Symposium on Programming (ESOP)*.
- 2352 Crafa, S. (2012) Behavioural types for actor systems. Technical report. [https://arxiv.org/abs/](https://arxiv.org/abs/1206.1687)
 2353 [1206.1687](https://arxiv.org/abs/1206.1687).
- 2354 Crafa, S. & Padovani, L. (2017) The chemical approach to typestate-oriented programming. In *ACM*
 2355 *Transactions on Programing Languages and Systems*. vol. 39. A preliminary version appeared in
 2356 the proceedings of OOPSLA'15.
- 2357 Culpepper, R. & Felleisen, M. (2010) Fortifying macros. In *Proceedings of the ACM SIGPLAN*
 2358 *International Conference on Functional Programming (ICFP)*.
- 2359 Culpepper, R. & Felleisen, M. (2012) Fortifying macros. *Journal of Functional Programming*. **22**.
- 2360 de Boer, F. S., Clarke, D. & Johnsen, E. B. (2007) A complete guide to the future. In *Proceedings of*
 2361 *the European Symposium on Programming (ESOP)*.
- 2362 De Giacomo, G., Lespérance, Y. & Levesque, H. J. (2000) ConGolog, a concurrent programming
 2363 language based on the situation calculus. *Artificial Intelligence*. **121**(1–2).
- 2364 De Giacomo, G., Lespérance, Y., Patrizi, F. & Sardina, S. (2016) Verifying ConGolog programs
 2365 on bounded situation calculus theories. In *Proceedings of the AAAI Conference on Artificial*
 2366 *Intelligence*.
- 2367 De Koster, J., Van Cutsem, T. & De Meuter, W. (2016) 43 years of Actors: a taxonomy of Actor
 2368 models and their key properties. In *Proceedings of the International Workshop on Programming*
 2369 *Based on Actors, Agents, and Decentralized Control (AGERE)*.
- 2370 de'Liguoro, U. & Padovani, L. (2018) Mailbox types for unordered interactions. In *Proceedings of*
 2371 *the European Conference on Object-Oriented Programming (ECOOP)*.
- 2372 Demetrescu, C., Finocchi, I. & Ribichini, A. (2011) Reactive imperative programming with dataflow
 2373 constraints. In *Proceedings of the ACM SIGPLAN Symposium on Object-Oriented Programming,*
 2374 *Systems, Languages & Applications (OOPSLA)*.
- 2375 Eugster, P. T., Felber, P. A., Guerraoui, R. & Kermarrec, A.-M. (2003) The many faces of
 2376 publish/subscribe. *ACM Computing Surveys*. **35**(2).
- 2377 Felleisen, M. (1988) The theory and practice of first-class prompts. Proceedings of the ACM
 2378 SIGPLAN Symposium on Principles of Programming Languages (POPL).
- 2379 Flatt, M., Yu, G., Findler, R. & Felleisen, M. (2007) Adding delimited and composable control
 2380 to a production programming environment. Proceedings of the ACM SIGPLAN International
 2381 Conference on Functional Programming (ICFP).
- 2382 Floyd, R. W. (1967) Assigning meaning to programs. *Proceedings of Symposium on Applied*
 2383 *Mathematics*. **19**.
- 2384 Fournet, C. & Gonthier, G. (2000) The join calculus: a language for distributed mobile programming.
 2385 Applied Semantics Summer School.
- 2386 Fournet, C., Laneve, C., Maranget, L. & Rémy, D. (2000) Inheritance in the join calculus. In
 2387 *Foundations of Software Technology and Theoretical Computer Science*.
- 2388 Gamma, E., Helm, R., Johnson, R. & Vlissides, J. (1994) *Design patterns: elements of reusable*
 2389 *object-oriented software*. Addison Wesley, Massachusetts.
- 2390 Garnock-Jones, T. (2017) *Conversational Concurrency*. PhD dissertation. Northeastern University.
- 2391 Garnock-Jones, T. & Felleisen, M. (2016) Coordinated concurrent programming in Syndicate. In
 2392 *Proceedings of the European Symposium on Programming (ESOP)*.
- Garnock-Jones, T., Tobin-Hochstadt, S. & Felleisen, M. (2014) The network as a language construct.
 In *Proceedings of the European Symposium on Programming (ESOP)*.
- González Boix, E., Scholliers, C., De Meuter, W. & D'Hondt, T. (2014) Programming mobile
 context-aware applications with TOTAM. *Journal of Systems and Software*. **92**(1).
- Henrio, L., Laneve, C. & Mastandrea, V. (2017) Analysis of synchronisations in stateful active
 objects. *Integrated Formal Methods*.
- Hewitt, C., Bishop, P. & Steiger, R. (1973) A universal modular actor formalism for artificial

intelligence. In *International Joint Conference on Artificial Intelligence*.

2393 Hoare, C. A. R. (1969) An axiomatic basis for computer programming. *Communications of the ACM*.
2394 **12**.

2395 Holzmann, G. (1997) The model checker SPIN. *IEEE Transactions on Software Engineering*.

2396 Honda, K., Vasconcelos, V. T. & Kubo, M. (1998) Language primitives and type discipline for
2397 structured communication-based programming. In *Proceedings of the European Symposium on*
Programming (ESOP).

2398 Honda, K., Yoshida, N. & Carbone, M. (2008) Multiparty asynchronous session types. In
2399 *Proceedings of the ACM SIGPLAN Symposium on Principles of Programming Languages (POPL)*.

2400 Igarashi, A. & Kobayashi, N. (2001) A generic type system for the Pi-calculus. In *Proceedings of the*
2401 *ACM SIGPLAN Symposium on Principles of Programming Languages (POPL)*.

2402 Jaghoori, M. M., Movaghgar, A. & Sirjani, M. (2006) Modere: the model-checking engine of Rebeca.
2403 *Proceedings of the ACM Symposium on Applied Computing (SAC)*.

2404 Lucassen, J. M. & Gifford, D. K. (1988) Polymorphic effect systems. In *Proceedings of the ACM*
SIGPLAN Symposium on Principles of Programming Languages (POPL).

2405 Manna, Z. & Pnueli, A. (1991) *The temporal logic of reactive and concurrent systems: specification*.
2406 Springer Verlag.

2407 Miller, M. S., Tribble, E. D. & Shapiro, J. (2005) Concurrency among strangers. In *International*
Symposium on Trustworthy Global Computing. Edinburgh, Scotland.

2408 Mostinckx, S., Lombide Carreton, A. & De Meuter, W. (2008) Reactive context-aware programming.
2409 In *Workshop on Context-aware Adaptation Mechanisms for Pervasive and Ubiquitous Services*
2410 *(CAMPUS)*. vol. 10 of *Electronic Communications of the EASST, DisCoTec*.

2411 Mostinckx, S., Scholliers, C., Philips, E., Herzeel, C. & De Meuter, W. (2007) Fact spaces: coordi-
2412 nation in the face of disconnection. In *International Conference on Coordination Languages and*
2413 *Models*.

2414 Mostrous, D. & Vasconcelos, V. T. (2011) Session typing for a featherweight Erlang. *International*
Conference on Coordination Languages and Models.

2415 Murphy, A. L., Picco, G. P. & Roman, G.-C. (2006) Lime: A coordination model and middle-
2416 ware supporting mobility of hosts and agents. *ACM Transactions on Software Engineering and*
2417 *Methodology*. **15**(3).

2418 Neykova, R. & Yoshida, N. (2014) Multiparty session actors. *International Conference on*
Coordination Languages and Models.

2419 Pnueli, A. (1977) The temporal logic of programs. *Symposium on Foundations of Computer Science*
2420 *(SFCS)*.

2421 Rozier, K. Y. (2011) Linear temporal logic symbolic model checking. *Computer Science Review*.
2422 **5**(2).

2423 Salvaneschi, G. & Mezini, M. (2014) Towards reactive programming for object-oriented applica-
2424 tions. In *Transactions on Aspect-Oriented Software Development*.

2425 Sanderson, S. (2010) Introducing Knockout, a UI library for
2426 JavaScript. [http://blog.stevensanderson.com/2010/07/05/
introducing-knockout-a-ui-library-for-javascript/](http://blog.stevensanderson.com/2010/07/05/introducing-knockout-a-ui-library-for-javascript/). Accessed: 2022-07-25.

2427 Scalas, A. & Yoshida, N. (2019) Less is more: multiparty session types revisited. In *Proceedings of*
2428 *the ACM SIGPLAN Symposium on Principles of Programming Languages (POPL)*.

2429 Scalas, A., Yoshida, N. & Benussi, E. (2019) Verifying message-passing programs with dependent
2430 behavioural types. In *Proceedings of the ACM SIGPLAN Conference on Programming Language*
Design and Implementation (PLDI).

2431 Schäfer, J. & Poetsch-Heffter, A. (2008) CoBoxes: unifying active objects and structured heaps. In
2432 *Formal Methods for Open Object-Based Distributed Systems*. Lecture Notes in Computer Science.

2433 Schäfer, J. & Poetsch-Heffter, A. (2010) JCoBox: generalizing active objects to concurrent compo-
2434 nents. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*.
2435 Lecture Notes in Computer Science.

2436 Scholliers, C., González Boix, E. & De Meuter, W. (2009) TOTAM: scoped tuples for the Ambient.
2437 *Workshop on Context-aware Adaptation Mechanisms for Pervasive and Ubiquitous Services*

(CAMPUS).

- 2439 Scholliers, C., González Boix, E., De Meuter, W. & D'Hondt, T. (2010) Context-aware tuples for the
2440 Ambient. In *Proceedings of On the Move to Meaningful Internet Systems. OTM'10*.
- 2441 Sirjani, M., Movaghar, A., Shali, A. & de Boer, F. S. (2004) Modeling and verification of reactive
2442 systems using Rebeca. *Fundamenta Informaticae*. **63**(4).
- 2443 Strom, R. E. & Yemini, S. (1986) Typestate: A programming language concept for enhancing
2444 software reliability. *IEEE Transactions on Software Engineering*.
- 2445 Tasharofi, S., Dinges, P. & Johnson, R. E. (2013) Why do Scala developers mix the actor model
2446 with other concurrency models? In *Proceedings of the European Conference on Object-Oriented
Programming (ECOOP)*.
- 2447 Toninho, B., Caires, L. & Pfenning, F. (2011) Dependent session types via intuitionistic linear
2448 type theory. In *International Symposium on Principles and Practice of Declarative Programming
(PPDP)*.
- 2449 Van Cutsem, T., Gonzalez Boix, E., Scholliers, C., Lombide Carreton, A., Harnie, D., Pinte, K. & De
2450 Meuter, W. (2014) AmbientTalk: programming responsive mobile peer-to-peer applications with
2451 actors. *Computer Languages, Systems & Structures*. **40**(3-4).
- 2452 Van Cutsem, T., Mostinckx, S., González Boix, E., Dedecker, J. & De Meuter, W. (2007)
2453 AmbientTalk: object-oriented event-driven programming in mobile ad hoc networks. International
2454 Conference of the Chilean Society of Computer Science.
- 2455 Vardi, M. Y. (2001) Branching vs. linear time: final showdown. Tools and Algorithms for the
2456 Construction and Analysis of Systems (TACAS).
- 2457 Vieira, H. T., Caires, L. & Seco, J. C. (2008) The conversation calculus: a model of service-oriented
2458 computation. In *Proceedings of the European Symposium on Programming (ESOP)*.
- 2459 Wright, A. K. & Felleisen, M. (1994) A syntactic approach to type soundness. *Information and
2460 Computation*. **115**(1).
- 2461 Yi, K. & Ryu, S. (2002) A cost-effective estimation of uncaught exceptions in standard ML programs.
2462 *Theoretical Computer Science*. **277**.
- 2463 Yousefi, B., Ghassemi, F. & Khosravi, R. (2015) Modeling and efficient verification of broadcasting
2464 actors. *Fundamentals of Software Engineering*.
- 2465
- 2466
- 2467
- 2468
- 2469
- 2470
- 2471
- 2472
- 2473
- 2474
- 2475
- 2476
- 2477
- 2478
- 2479
- 2480
- 2481
- 2482
- 2483
- 2484

A Comparing Facet-based and Functional Actors

This section provides additional motivation for the facet notation via a comparison to the results of implementing actor behaviors in a procedural manner.

Figure 22 implements the behavior of the `light-actor` (figure 4) using a procedural interface to the Racket dataspace library. In the procedural interface, the behavior of an individual actor is specified as a state transducer function.

The `BEHAVIOR` function (lines 3–8) implements this interface. Its inputs are a representation of an event and the `current-state` of the actor. The body of the `BEHAVIOR` function dispatches to one of several sub-behavior functions for the actor. The `current-state` is an instance of the `light-actor-state` struct containing two pieces of information. The first is the actor's current control state, dictating which conversation(s) to engage in response to an event. The actor's interactions are simple enough that its control may be encoded as a switch.¹⁴ The second holds any additional data relevant to the active sub-behavior.

Placing the `light-actor` into a dataspace as an instance of an actor struct (lines 9–11) requires specifying the initial `light-actor-state` as well as the actor's initial assertions, establishing interest in receiving `wall-switch` information. Note that the actor's `ctrl-state`—`UNPOWERED`—is expressed separately but must be in sync with the initial assertion(s).

The `unpowered`, `powered`, and `located` sub-routines implement the actor's various sub-behaviors. Each routine computes a `transition-to` struct as its result. The first slot of the `transition-to` specifies a new private state for the actor, while the second requests updates to the actor's current assertions. The updates are a list, with each item calling for the introduction of a new assertion (`assert`) or withdrawal of an existing one (`retract`).

The `unpowered` routine (lines 13–19) responds to notification that the `wall-switch` is `ON` with a `transition-to` the `POWERED` state. In the `POWERED` state, the actor uses the `powered-state` struct as its private state to keep track of whether the light is currently `ON` or `OFF`. Additionally, it introduces the `light` assertion and expresses interest in a `room-assignment`.

The `powered` routine (lines 21–32) implements two possibilities. The first is the event that the wall switch is turned off (line 23). The resulting transition returns to the "unpowered" state, `retract`-ing the `powered`-specific assertions. The second possibility is the appearance of a `room-assignment` assertion (line 28). In that case, the actor transitions to the `LOCATED` state. The private state of the actor must then keep track of both the current `ON/OFF` status of the light as well as the assigned room. The `located-state` struct serves this purpose. Finally, the actor introduces a new interest in presence sensor information for the assigned room.

¹⁴ An actor engaged in truly complex conversations requires a more sophisticated encoding or the use of (delimited) continuations. Note that while it may seem attractive to represent the control state as a function, allowing the body of `BEHAVIOR` to simply be the application of the `ctrl-state` to `event`, that approach makes composing the behaviors of different states more difficult, as when the `located` behavior inspects the desired `ctrl-state` of the `powered` behavior (lines 40–41). While additional programming patterns may be employed to address this deficiency, the facet notation is the result of considering such patterns and designing a DSL around them.

The `located` routine (lines 34-54) must still react to the withdrawal of `wall-switch` assertions, in the manner of `powered`. To accomplish this goal, it invokes the `powered` function on the event and a synthesized instance of `powered-state`. A procedure-based actor must determine how to combine the state-changes and actions from each of its conversations on a case-by-case basis. Here, it inspects the resulting transition. If the transition is to the "unpowered" state, the `located` function augments the transition with the retraction of interest in presence sensor assertions. Otherwise, the actor ignore the transition and proceeds by analyzing the event. Depending on the presence sensor information, the actor either updates the light's private state and public assertion to ON or OFF.

Analysis. The facet-based implementation enjoys a number of advantages.

Spatial locality. The procedural actor describes its initial assertion in one place and updates to its assertions throughout different locations in the code. The actor manipulates its `light` assertion in four different locations (lines 18, 26, 46, and 50). The problem is not unique to assertion-manipulating dataspace actors. Consider the code for a networked actor. The code for connecting is typically in one place, while code for using the connection or tearing it down is somewhere else. By contrast, in the facet language, an actor accomplishes the same behavior with a single endpoint—which is less error-prone than when functionality is distributed over many locations.

Conversational structure. The organization of the facet code directly reflects the relationships among the actor's conversations. The interaction with the presence sensor takes place as a sub-conversation in the context of a `room-assignment` assertion, which is dependent on the conversation with the wall switch. Recovering the same relationships from the implementation of the procedural actor requires a careful inspection of its code.

Localized State and Automatic Demultiplexing. Facet-based actors define state item-by-item and in the relevant context, and dataflow keeps public and private state automatically in sync. Likewise, the control state of the facet-based actor is maintained by the language implementation, which automatically routes incoming assertion patches to the (pattern)matching endpoints.

The contrast between the code in a facet notation and a procedural language is large, even when the latter is enriched with imperative assignment or objects. For simple conversations, switches like those in the light actor combined with assignments or objects will simplify the procedural code. Complications arise for an actor conducting parallel conversations. While each conversation's state could be encapsulated in an object, doing so would have the unfortunate consequence that the code for manually demultiplexing incoming events and the code for maintaining a local view would be at separate places.

Similarly, a procedural language forces an actor to maintain its control state explicitly. In our concrete example, the `ctrl-state` field is the symptom of this problem. Continuations (delimited or unconstrained) (Felleisen, 1988; Flatt et al., 2007) do not solve the problem either, as code (or patterns) for managing the continuations must still be deployed to keep track of the state.

Brevity. Finally, the facet-based implementation requires far fewer lines of code.


```

1 (define (spawn-light wall-switch-id)
2577 2 (define my-id (generate-unique-id "light"))
3 (define (BEHAVIOR event current-state)
2578 4 (match-define (light-actor-state ctrl-state data) current-state)
2579 5 (match ctrl-state
6 [UNPOWERED (unpowered event current-state)]
2580 7 [POWERED (powered event current-state)]
2581 8 [LOCATED (located event current-state)]))
9 (define STATE0 (light-actor-state UNPOWERED (unpowered-state)))
2582 10 (define ASSERTIONS0 (assertions (? (wall-switch wall-switch-id ON))))
2583 11 (actor BEHAVIOR STATE0 ASSERTIONS0))
12
2584 13 (define (unpowered event current-state)
2585 14 (match event
15 [(asserted (wall-switch _ ON))
16 (transition-to (light-actor-state POWERED
2587 17 (powered-state ON))
18 #:with-updates (list (assert (light my-id ON))
19 (assert (? (room-assignment my-id *))))))]
20
2590 21 (define (powered event current-state)
2591 22 (match event
23 [(retracted (wall-switch _ ON))
2592 24 (transition-to (light-actor-state UNPOWERED
25 (unpowered-state))
2593 25 #:with-updates (list (retract (light my-id *))
26 (retract (? (room-assignment my-id *)))))]
2594 27 [(asserted (room-assignment my-id $room))
2595 28 (transition-to (light-actor-state LOCATED
2596 29 (located-state (powered-state-on? current-state)
2597 30 room))
2598 31 #:with-updates (list (assert (? (in-room room 0))))))]
2599 32
2600 33 (define (located event current-state)
2601 34 (match-define (located-state on? room) current-state)
2602 35 (define powered-transition (powered event (powered-state on?)))
2603 36 (match powered-transition
2604 37 [(transition-to (light-actor-state UNPOWERED _) $updates)
2605 38 (transition-to (light-actor-state UNPOWERED (unpowered-state))
2606 39 #:with-updates (append updates
2607 40 (list (retract (? (in-room room 0))))))]
2608 41 [_
2609 42 (match event
2610 43 [(asserted (in-room room 0))
2611 44 (transition-to (light-actor-state LOCATED (located-state OFF room))
2612 45 #:with-updates (list (retract (light my-id *))
2613 46 (assert (light my-id OFF))))]
2614 47 [(retracted (in-room room 0))
2615 48 (transition-to (light-actor-state LOCATED (located-state ON room))
2616 49 #:with-updates (list (retract (light my-id *))
2617 50 (assert (light my-id ON))))]]))
2618 51

```

Fig. 22: Procedural Light Actor

B Facet Metafunctions

The semantics of section 5.2 refers to several secondary metafunctions. This section provides their formal definitions.

The *match_D* function matches an event against an event descriptor. Successful matches yield a set of substitutions for the binding variables of the event descriptor's pattern:

2623

2624

2625

2626

2627

2628

2629

2630

2631

2632

2633

2634

2635

2636

2637

2638

2639

2640

2641

2642

2643

2644

2645

2646

2647

2648

2649

2650

2651

2652

2653

2654

2655

2656

2657

2658

2659

2660

2661

2662

2663

2664

2665

2666

2667

2668

$$\begin{aligned}
& \text{match}_D : D \times \text{Evt} \times \pi \times \pi \times \sigma \xrightarrow{\text{partial}} \mathcal{S} \\
\text{match}_D(\text{start}, \text{start}, \pi, \pi', \sigma) &= \{\emptyset\} \\
\text{match}_D(\text{stop}, \text{stop}, \pi, \pi', \sigma) &= \{\emptyset\} \\
\text{match}_D(\text{asserted } e, \pi^+/\pi^-, \pi, \pi', \sigma) &= \text{project}(v, \pi, \pi', \pi^+) \quad \text{if } v = \text{eval}(e, \sigma) \\
\text{match}_D(\text{retracted } e, \pi^+/\pi^-, \pi, \pi', \sigma) &= \text{project}(v, \pi, \pi', \pi^-) \quad \text{if } v = \text{eval}(e, \sigma)
\end{aligned}$$

The *project* function handles matching a pattern against a set of assertions, yielding a set of substitutions. The set is empty when there are no matching assertions in the set:

$$\begin{aligned}
\text{project} &: v \times \pi \times \pi \times \pi \longrightarrow \mathcal{S} \\
\text{project}(v, \pi_i, \pi'_i, \pi) &= \{\text{match}_v(v, I) \mid I \in \{\text{inst}(v, u) \mid u \in \pi\}, \\
&\quad \text{known}(I, \pi_i) \neq \text{known}(I, \pi'_i)\}
\end{aligned}$$

where

$$\text{known}(I, \pi) = 1 \text{ if } \exists u \in \pi. \text{match}_v(I, u) \text{ defined; else, } 0$$

The helper *inst* partially matches a pattern against a value, leaving wildcards and binders in place:

$$\begin{aligned}
\text{inst} &: v \times u \xrightarrow{\text{partial}} v \\
\text{inst}(\star, u) &= \star \\
\text{inst}(x : \tau, u) &= x \\
\text{inst}(v, v) &= v \\
\text{inst}(m(\vec{v}), m(\vec{u})) &= m(\vec{v}') \\
&\quad \text{if } |\vec{v}'| = |\vec{u}| \\
&\quad \vec{v}' = \text{inst}(v, u) \\
\text{inst}(\text{? } v, \text{? } u) &= \text{? } v' \quad \text{if } v' = \text{inst}(v, u)
\end{aligned}$$

The match_v function implements basic value-against-pattern matching. We assume without loss of generality that all binders in a given pattern are unique:

$$\begin{aligned}
\text{match}_v &: v \times u \xrightarrow{\text{partial}} \gamma \\
\text{match}_v(\star, u) &= \emptyset \\
\text{match}_v(x : \tau, u) &= \{x \mapsto u\} \\
\text{match}_v(v, v) &= \emptyset \\
\text{match}_v(m(\vec{v}), m(\vec{u})) &= \bigcup \vec{\gamma} \\
&\quad \text{if } |\vec{v}| = |\vec{u}| \\
&\quad \vec{\gamma} = \text{match}_v(v, u) \\
\text{match}_v(\text{? } v, \text{? } u) &= \text{match}_v(v, u)
\end{aligned}$$

The \oplus operation updates an actor's knowledge of assertions based on a patch:

$$\pi \oplus \pi^+ / \pi^- = (\pi \cup \pi^+) - \pi^-$$

The evaluator for expressions is *eval*. The function is partial due to the possibility of unbound field names. A primitive interpretation δ is a partial function from an operation p and a vector of values the result value:

$$\begin{aligned}
\text{eval} &: e \times \sigma \xrightarrow{\text{partial}} v \\
\text{eval}(b, \sigma) &= b \\
\text{eval}(!x, \sigma) &= \sigma(x) \\
\text{eval}(\star, \sigma) &= \star \\
\text{eval}(x : \tau, \sigma) &= x : \tau \\
\text{eval}(p(\vec{e}), \sigma) &= \delta(p, \vec{v}) \text{ if } \vec{v} = \overline{\text{eval}(e, \sigma)} \\
\text{eval}(m(\vec{e}), \sigma) &= m(\vec{v}) \text{ if } \vec{v} = \overline{\text{eval}(e, \sigma)} \\
\text{eval}(\text{? } e, \sigma) &= \text{? } v \quad \text{if } v = \text{eval}(e, \sigma)
\end{aligned}$$

The *assertions-of* function produces the current assertions made by an actor given its active facet tree and field store:

$$\begin{aligned}
2669 \quad & \text{assertions-of} && : \text{FT} \times \sigma \longrightarrow \pi \\
2670 \quad & \text{assertions-of}(\epsilon, \sigma) &= & \emptyset \\
2671 \quad & \text{assertions-of}(\text{fn}[\vec{e} \overrightarrow{(\text{D Pr})}].\text{FT}, \sigma) &= & \bigcup \overrightarrow{\text{assertions-of}_e(\epsilon, \sigma)} \cup \\
2672 \quad & & & \bigcup \overrightarrow{\text{assertions-of}_D(D, \sigma)} \cup \\
2673 \quad & & & \bigcup \overrightarrow{\text{assertions-of}(\text{FT}, \sigma)}
\end{aligned}$$

2674 It utilizes a family of helper functions for determining the assertions associated with
2675 endpoints. The *assertions-of_e* function does the work for assertion endpoints:

$$\begin{aligned}
2676 \quad & \text{assertions-of}_e && : e \times \sigma \longrightarrow \pi \\
2677 \quad & \text{assertions-of}_e(\star, \sigma) &= & \mathbf{Assertion} \\
2678 \quad & \text{assertions-of}_e(x : \tau, \sigma) &= & \mathbf{Assertion} \\
2679 \quad & \text{assertions-of}_e(m(\vec{e}_i), \sigma) &= & \{m(\vec{u}_i) \mid u_i \in \text{assertions-of}_e(\epsilon_i, \sigma)\} \\
2680 \quad & \text{assertions-of}_e(? e, \sigma) &= & \{? u \mid u \in \text{assertions-of}_k(\epsilon, \sigma)\} \\
2681 \quad & \text{assertions-of}_e(e, \sigma) &= & \{\text{eval}(e, \sigma)\}
\end{aligned}$$

2682 While the *assertions-of_D* metafunction produces the assertion of interest when needed
2683 by an event-handler endpoint:

$$\begin{aligned}
2684 \quad & \text{assertions-of}_D && : D \times \sigma \longrightarrow \pi \\
2685 \quad & \text{assertions-of}_D(\text{start}, \sigma) &= & \emptyset \\
2686 \quad & \text{assertions-of}_D(\text{stop}, \sigma) &= & \emptyset \\
2687 \quad & \text{assertions-of}_D(\text{asserted } e, \sigma) &= & \{? u \mid u \in \text{assertions-of}_e(\epsilon, \sigma)\} \\
2688 \quad & \text{assertions-of}_D(\text{retracted } e, \sigma) &= & \{? u \mid u \in \text{assertions-of}_e(\epsilon, \sigma)\}
\end{aligned}$$

2690 C The Extended Light Actor

2692 The full implementation of the light actor in figure 4 is a good example for illustrating how
2693 the facet machine deals with loading and manipulating states. Here is the fixed version of
2694 the actor description, assuming suitable bindings for wall-switch-id and my-id:
2695

```

2697 start light-facet
2698   0
2699   (asserted wall-switch(wall-switch-id, ON)
2700     field state = ON in
2701       start during<wall-switch>
2702         light(my-id, !state) ∪ 0
2703         (retracted wall-switch(wall-switch-id, ON)
2704           stop during<wall-switch>)
2705         (asserted room-assignment(my-id, room:String)
2706           start during<room-assignment>
2707             0
2708             (retracted room-assignment(my-id, room)
2709               stop during<room-assignment>)
2710             (asserted in-room(room, *)
2711               start during<in-room>
2712                 0
2713                 (retracted in-room(room, *)
2714                   stop during<in-room>)
2715                 (asserted room-empty(room)
2716                   .)
2717                 (asserted room-occupied(room)
2718                   .))))))

```

Figure 23 shows the complete machine transition sequence, starting from the initial state. From there, the machine executes a labeled inject transition to get the start instruction of the facet script. A transfer transition launches the script, which then immediately issues an action patch via another labeled inject transition. At this point, the facet is ready to shut down, which the machine realizes with stop transition.

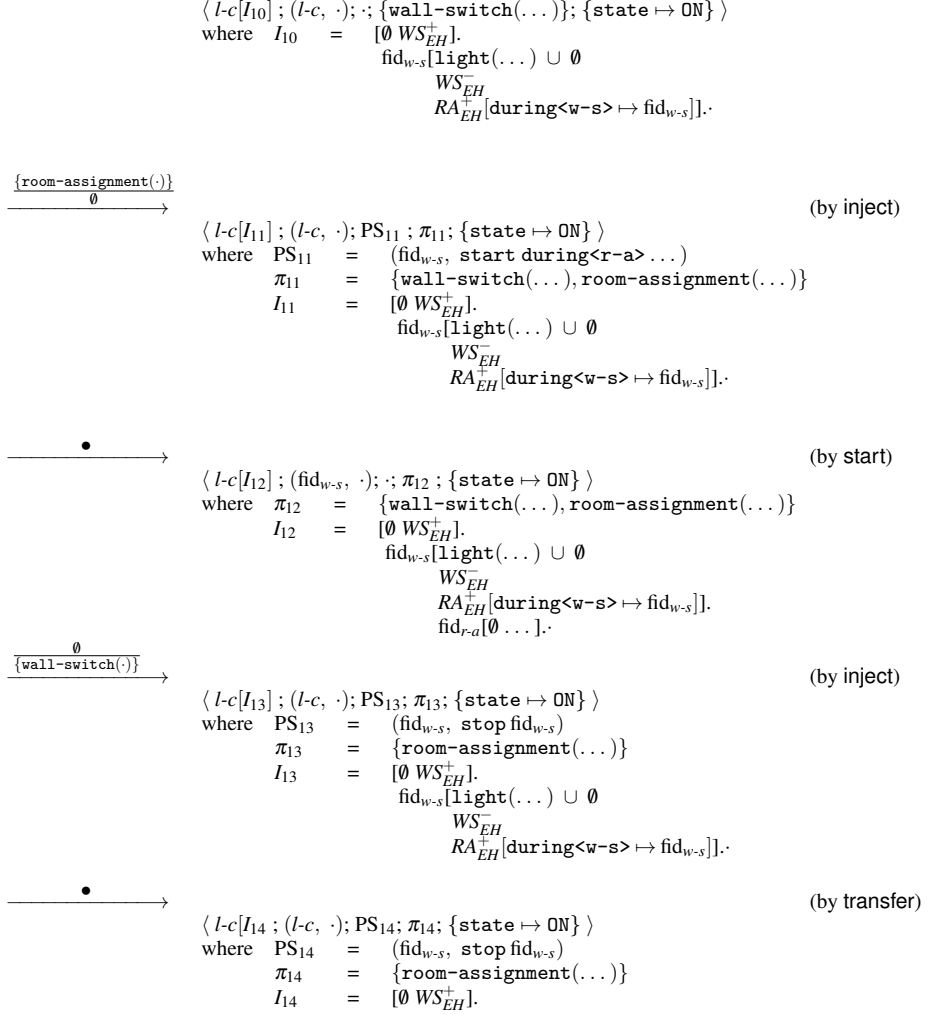


Fig. 23: The light-facet facet as an initial machine state and its transitions

D Type Judgments

The following definitions complete the type judgment for facet actor terms (section 6.1).

The judgment for event descriptions is $\Gamma \vdash_D D : D_T$:

$$\frac{}{\Gamma \vdash_{\text{D}} \text{start} : \text{start}} \boxed{\text{D-START}} \qquad \frac{}{\Gamma \vdash_{\text{D}} \text{stop} : \text{stop}} \boxed{\text{D-STOP}}$$

$$\frac{\Gamma \vdash_e e : \tau \quad \text{pattern}(\tau)}{\Gamma \vdash_{\text{D}} \text{asserted } e : \text{asserted } \tau} \boxed{\text{D-ASSERTED}}$$

$$\frac{\Gamma \vdash_e e : \tau \quad \text{pattern}(\tau)}{\Gamma \vdash_{\text{D}} \text{retracted } e : \text{retracted } \tau} \boxed{\text{D-ASSERTED}}$$

The judgment for expressions is $\Gamma \vdash_e e : \tau$:

$$\frac{}{\Gamma \vdash_e b : \mathbf{B}} \boxed{\text{E-BASE}} \qquad \frac{\Gamma(x) = \tau}{\Gamma \vdash_e x : \tau} \boxed{\text{E-VAR}} \qquad \frac{}{\Gamma \vdash_e * : * } \boxed{\text{E-WILDCARD}}$$

$$\frac{}{\Gamma \vdash_e x : \tau : x : \tau} \boxed{\text{E-BIND}} \qquad \frac{\Gamma(x) = \text{Field } \tau}{\Gamma \vdash_e !x : \tau} \boxed{\text{E-REF}}$$

$$\frac{\Gamma \vdash_e e : \vec{\tau} \quad \delta_{\tau}(p, \vec{\tau}) = \tau'}{\Gamma \vdash_e p(\vec{e}) : \tau'} \boxed{\text{E-PRIMOP}} \qquad \frac{\Gamma \vdash_e e : \tau}{\Gamma \vdash_e ? e : ? \tau} \boxed{\text{E-?}}$$

$$\frac{\Gamma \vdash_e e : \vec{\tau}}{\Gamma \vdash_e m(\vec{e}) : m(\vec{\tau})} \boxed{\text{E-TUPLE}}$$

Types suitable for asserting, $\text{assertable}(\tau)$:

$$\frac{}{\text{assertable}(\mathbf{B})} \boxed{\text{A-BASE}} \qquad \frac{}{\text{assertable}(*)} \boxed{\text{A-WILDCARD}}$$

$$\frac{\text{assertable}(\tau)}{\text{assertable}(\tau)} \boxed{\text{A-WILDCARD}} \qquad \frac{\text{assertable}(\vec{\tau})}{\text{assertable}(m(\vec{\tau}))} \boxed{\text{A-TUPLE}}$$

Types suitable for patterns, $\text{pattern}(\tau)$:

$$\frac{}{\text{pattern}(\mathbf{B})} \boxed{\text{P-BASE}} \qquad \frac{}{\text{pattern}(*)} \boxed{\text{P-WILDCARD}} \qquad \frac{}{\text{pattern}(x : \tau)} \boxed{\text{P-BIND}}$$

$$\frac{\text{pattern}(\tau)}{\text{pattern}(\tau)} \boxed{\text{P-WILDCARD}} \qquad \frac{\text{pattern}(\vec{\tau})}{\text{pattern}(m(\vec{\tau}))} \boxed{\text{P-TUPLE}}$$

The following metafunctions complete the definition of the type judgment.

The $\text{bindings}_{\text{D}}$ metafunction creates a type environment from the binding variables in an event description:

2761
2762
2763
2764
2765
2766
2767
2768
2769
2770
2771
2772
2773
2774
2775
2776
2777
2778
2779
2780
2781
2782
2783
2784
2785
2786
2787
2788
2789
2790
2791
2792
2793
2794
2795
2796
2797
2798
2799
2800
2801
2802
2803
2804
2805
2806

$$\begin{aligned}
2807 \quad & \text{bindings}_{\mathcal{D}} & : & \mathcal{D} \longrightarrow \Gamma \\
2808 \quad & \text{bindings}_{\mathcal{D}}(\text{start}) & = & \cdot \\
& \text{bindings}_{\mathcal{D}}(\text{stop}) & = & \cdot \\
2809 \quad & \text{bindings}_{\mathcal{D}}(\text{asserted } e) & = & \text{bindings}_{\mathcal{P}}(e) \\
2810 \quad & \text{bindings}_{\mathcal{D}}(\text{retracted } e) & = & \text{bindings}_{\mathcal{P}}(e)
\end{aligned}$$

2811 It utilizes a helper function for patterns, $\text{bindings}_{\mathcal{P}}$:

$$\begin{aligned}
2812 \quad & \text{bindings}_{\mathcal{P}} & : & e \longrightarrow \Gamma \\
2813 \quad & \text{bindings}_{\mathcal{P}}(\star) & = & \cdot \\
2814 \quad & \text{bindings}_{\mathcal{P}}(x : \tau) & = & x : \tau \\
2815 \quad & \text{bindings}_{\mathcal{P}}(m(\vec{e})) & = & \overrightarrow{\text{bindings}_{\mathcal{P}}(e)} \\
2816 \quad & \text{bindings}_{\mathcal{P}}(? e) & = & \text{bindings}_{\mathcal{P}}(e) \\
2817 \quad & \text{bindings}_{\mathcal{P}}(e) & = & \cdot
\end{aligned}$$

2818 Two metafunctions prune type environments to help keep track of facet and field names.
 2819 The first, prune , removes all facet and field names from an environment, i.e. variables that
 2820 are not shared between actors:

$$\begin{aligned}
2821 \quad & \text{prune} & : & \Gamma \longrightarrow \Gamma \\
2822 \quad & \text{prune}(\cdot) & = & \cdot \\
2823 \quad & \text{prune}(\Gamma, \text{fn} : \text{FacetName}) & = & \text{prune}(\Gamma) \\
2824 \quad & \text{prune}(\Gamma, x : \tau) & = & \begin{cases} \text{prune}(\Gamma) & \text{if } x = \text{Field } \tau' \\ \text{prune}(\Gamma), x : \tau & \text{otherwise} \end{cases}
\end{aligned}$$

2825 The second, prune-up-to , removes all facet names that are children of a particular facet,
 2826 giving an environment of names of active facets after that facet stops:

$$\begin{aligned}
2827 \quad & \text{prune-up-to} & : & \text{fn} \times \Gamma \longrightarrow \Gamma \\
2828 \quad & \text{prune-up-to}(\text{fn}, \cdot) & = & \cdot \\
2829 \quad & \text{prune-up-to}(\text{fn}, \Gamma, \text{fn}' : \text{FacetName}) & = & \begin{cases} \Gamma & \text{if } \text{fn} = \text{fn}' \\ \text{prune-up-to}(\text{fn}, \Gamma) & \text{otherwise} \end{cases} \\
2830 \quad & \text{prune-up-to}(\text{fn}, \Gamma, x : \tau) & = & x : \tau, \text{prune-up-to}(\Gamma, \text{fn})
\end{aligned}$$

2831 E Dataspace Metafunctions

2832 The semantics of section 7 refers to several metafunctions. For completeness, this section
 2833 provides their formal definitions. But, also see the work of Caldwell et al. (2020) plus
 2834 Garnock-Jones and Felleisen (2016), which fully specify the semantics of dataspace sys-
 2835 tems as well. The primary difference between the semantics presented here and prior work
 2836 is the aggregation of incoming patches performed by bc_{Δ} below.

2837 The $\text{boot}_{\mathcal{DS}}$ metafunction initializes a dataspace program based on a collection of initial
 2838 facet-based actor descriptions:

$$\begin{aligned}
2839 \quad & \text{boot}_{\mathcal{DS}} & : & \overrightarrow{\text{Pr}} \longrightarrow \text{DS} \\
2840 \quad & \text{boot}_{\mathcal{DS}}(\overrightarrow{\text{Pr}}) & = & [(\ell, \overrightarrow{\pi/\theta}); \emptyset; \overrightarrow{A}] \\
2841 \quad & & & \text{where} \\
2842 \quad & & & \overrightarrow{P} = \overrightarrow{\text{boot}_{\mathcal{P}}(\text{Pr})} \\
2843 \quad & & & \overrightarrow{\Sigma}, \overrightarrow{\pi} = \overrightarrow{\text{boot}_{\Sigma}(\text{Pr})} \\
2844 \quad & & & \overrightarrow{\ell} = 0 \cdot 1 \cdot \dots \cdot |\overrightarrow{\text{Pr}}| - 1 \\
2845 \quad & & & \overrightarrow{A} = \overrightarrow{\ell} \mapsto \overrightarrow{\Sigma}
\end{aligned}$$

2846 The boot_{Σ} function boots a process description to an actor state:

2853
2854
2855
2856
2857
2858
2859
2860
2861
2862
2863
2864
2865
2866
2867
2868
2869
2870
2871
2872
2873
2874
2875
2876
2877
2878
2879
2880
2881
2882
2883
2884
2885
2886
2887
2888
2889
2890
2891
2892
2893
2894
2895
2896
2897
2898

$$\begin{aligned} \text{boot}_\Sigma & : \mathbf{P} \longrightarrow \Sigma \times \pi \\ \text{boot}_\Sigma(\text{actor } f \vee \pi) & = \mathbf{0}/\mathbf{0} \triangleright (f, \vee), \pi \end{aligned}$$

The *update* metafunction is the work horse of dataspace event dispatch:

$$\begin{aligned} \text{update} & : (\mathbf{R} \times \vec{\mathbf{A}}) \times (\ell \times \Delta) \longrightarrow \mathbf{R} \times \vec{\mathbf{A}} \\ \text{update}((\mathbf{R}, \vec{\mathbf{A}}), (\ell, \Delta)) & = \mathbf{R} \oplus_{\mathbf{R}} (\Delta, \ell), \overrightarrow{bc_\Delta}(\ell, \mathbf{R}, \Delta, \vec{\mathbf{A}}) \end{aligned}$$

The bc_Δ metafunction updates the pending event for an actor if it has an active interest:

$$\begin{aligned} bc_\Delta & : \ell \times \mathbf{R} \times \text{Evt} \times \mathbf{A}_Q \longrightarrow \mathbf{A}_Q \\ bc_\Delta(\ell_{\text{evt}}, \mathbf{R}_{\text{old}}, \pi_{\text{add}}/\pi_{\text{del}}, \ell \mapsto \Delta \triangleright \mathbf{B} \cdot) & = \begin{cases} \ell \mapsto \Delta \circ \Delta_{fb} \triangleright \mathbf{B} \cdot & \text{if } \ell = \ell_{\text{evt}} \\ \ell \mapsto \Delta \circ \Delta_{\text{other}} \triangleright \mathbf{B} \cdot & \text{if } \ell \neq \ell_{\text{evt}} \end{cases} \end{aligned}$$

where

$$\begin{aligned} \Delta_{fb} & = \{u \mid u \in \pi_{\bullet \text{add}}, (? u, \ell) \in \mathbf{R}_{\text{new}}\} \cup \{u \mid u \in (\pi_{\circ} \cup \pi_{\bullet \text{add}} - \pi_{\bullet \text{del}}), ? u \in \pi_{\text{add}}\} / \\ & \quad \{u \mid u \in \pi_{\bullet \text{del}}, (? u, \ell) \in \mathbf{R}_{\text{old}}\} \cup \{u \mid u \in \pi_{\circ}, ? u \in \pi_{\text{del}}\} \\ \Delta_{\text{other}} & = \{u \mid u \in \pi_{\bullet \text{add}}, (? u, \ell) \in \mathbf{R}_{\text{old}}\} / \{u \mid u \in \pi_{\bullet \text{del}}, (? u, \ell) \in \mathbf{R}_{\text{old}}\} \\ \pi_{\bullet} & = \{u \mid (u, \ell') \in \mathbf{R}_{\text{old}}, \ell' \neq \ell_{\text{evt}}\} \\ \mathbf{R}_{\text{new}} & = \mathbf{R}_{\text{old}} \oplus_{\mathbf{R}} (\pi_{\text{add}}/\pi_{\text{del}}, \ell) \\ \pi_{\bullet \text{add}} & = \pi_{\text{add}} - \pi_{\bullet} \\ \pi_{\circ} & = \{u \mid (u, \ell') \in \mathbf{R}_{\text{old}}\} \\ \pi_{\text{del} \bullet} & = \pi_{\text{del}} - \pi_{\bullet} \end{aligned}$$

The \circ operator applies two patches in sequence, maintaining disjointness:

$$\begin{aligned} \circ & : \Delta \times \Delta \longrightarrow \Delta \\ \frac{\pi_1^+}{\pi_1^-} \circ \frac{\pi_2^+}{\pi_2^-} & = \frac{(\pi_1^+ - \pi_2^-) \cup \pi_2^+}{(\pi_1^- - \pi_2^+) \cup \pi_2^-} \end{aligned}$$

The helper function *patch* calculates the difference in assertions between two machine states:

$$\begin{aligned} \text{patch} & : \mathbf{M} \times \mathbf{M} \longrightarrow \Delta \\ \text{patch}(\mathbf{M}, \mathbf{M}') & = \frac{\text{assertions-of}_{\mathbf{M}}(\mathbf{M}') - \text{assertions-of}_{\mathbf{M}}(\mathbf{M})}{\text{assertions-of}_{\mathbf{M}}(\mathbf{M}) - \text{assertions-of}_{\mathbf{M}}(\mathbf{M}')} \end{aligned}$$

The *label-to-action* function translates a facet machine's internal transition label to a sequence of (zero or one) actions:

$$\begin{aligned} \text{label-to-action} & : l_{\bullet} \longrightarrow \vec{\text{act}} \\ \text{label-to-action}(\bullet) & = \cdot \\ \text{label-to-action}(\overline{\text{Pr}}) & = \text{boot}_{\text{Pr}} \end{aligned}$$

The $\text{assertions-of}_{\mathbf{M}}$ function extends the *assertions-of* family of functions, determining the active assertions of a facet machine:

$$\begin{aligned} \text{assertions-of}_{\mathbf{M}} & : \mathbf{M} \longrightarrow \pi \\ \text{assertions-of}_{\mathbf{M}}(\text{error}) & = \mathbf{0} \\ \text{assertions-of}_{\mathbf{M}}(\langle \text{FT}; \vec{\text{I}}; \overline{\text{PS}}; \pi; \sigma \rangle) & = \text{assertions-of}(\text{FT}, \sigma) \end{aligned}$$

The $\oplus_{\mathbf{R}}$ operator updates the table of active assertions based on an actor's patch:

$$\begin{aligned} \oplus_{\mathbf{R}} & : \mathbf{R} \times (\Delta, \ell) \longrightarrow \mathbf{R} \\ \mathbf{R} \oplus_{\mathbf{R}} (\pi^+/\pi^-, \ell) & = \mathbf{R} \cup \{(u, \ell) \mid u \in \pi^+\} - \{(u, \ell) \mid u \in \pi^-\} \end{aligned}$$

The $\text{dispatch}_{\text{DS}}$ function invokes an actor's behavior on its pending event, when non-empty:

2899	$M_T \in \mathbf{TyMachine} = \langle FT_T; \vec{I}_T; \vec{PS}_T; \pi_\tau; \sigma_\tau \rangle$	$\tau_R \in \mathbf{SpecType} = \{B \mid == R\}$
2900	$FT_T \in \mathbf{FctTreeTy} = \varepsilon$	$\tau_R = \begin{array}{l} \\ \text{ ? } \tau_R \\ \text{ m}(\tau_R) \end{array}$
2901	$ \text{ fn}[A_\tau(\overrightarrow{D_T \ T})].\overrightarrow{FT_T}$	$R \in \mathbf{Refine} = x \mid b$
2902	$C_T \in \mathbf{TyContext} = \square$	$\pi_\tau \in \mathbf{ASetTy} = \mathcal{P}(\tau_R)$
2903	$ \text{ fid}[A_\tau(\overrightarrow{D_T \ T})].\overrightarrow{FT_T} \cdot C_T \cdot \overrightarrow{FT_T}$	$\Delta_\tau \in \mathbf{PatchTy} = \begin{array}{l} \pi_\tau^+ / \pi_\tau^- \\ \text{where } \pi_\tau^+ \cap \pi_\tau^- = \emptyset \end{array}$
2904	$PS_T \in \mathbf{PScriptTy} = (\text{fid}, T)$	$\Delta_\tau \in \mathbf{ExtEventTy} = \Delta_\tau$
2905	$I_T \in \mathbf{InstrTy} = \begin{array}{l} \text{start fn } A_\tau(\overrightarrow{D_T \ T}) @ \text{ fid} \\ \\ \text{stop fn} \\ \\ \text{spawn } \vec{T} \end{array}$	$Evt_\tau \in \mathbf{EventTy} = \begin{array}{l} \Delta_\tau \\ \\ \text{start} \\ \\ \text{stop} \end{array}$
2906	$\sigma_\tau \in \mathbf{StoreTy} = x \xrightarrow{\text{fin}} \tau$	$l_\tau \in \mathbf{LabelTy} = \begin{array}{l} \bullet \\ \\ \Delta_\tau \\ \\ \vec{T} \end{array}$
2907	$\delta_\tau \in \mathbf{TyInterp} = p \times \vec{\tau} \rightarrow \tau$	
2908	$\gamma_\tau \in \mathbf{TySub} = x \xrightarrow{\text{fin}} \tau$	
2909	$\mathcal{S}_\tau \in \mathbf{TySubSets} = \mathcal{P}(\mathbf{TySub})$	
2910		
2911		
2912		
2913		
2914		
2915		

Fig. 24: Type Evaluation Syntax

$$\begin{aligned}
 & \text{dispatch}_{DS} : \Sigma \longrightarrow \Sigma \times \Delta \times \vec{P} \\
 & \text{dispatch}_{DS}(\Delta \triangleright (f, v)) = \begin{cases} \Delta \triangleright (f, v), \emptyset / \emptyset, \cdot & \text{if } \Delta = \emptyset / \emptyset \\ \emptyset / \emptyset \triangleright (f, v'), \Delta, \vec{P} & \text{otherwise} \end{cases} \\
 & \text{where} \\
 & v', \vec{\text{act}} = f(\Delta, v) \\
 & \Delta = \Delta_0 \circ \Delta_1 \circ \dots \text{ for } \Delta_i \in \vec{\text{act}} \\
 & \vec{P} = P_0 \cdot P_1 \dots \text{ for } P_i \in \vec{\text{act}}
 \end{aligned}$$

F Type Level Programs

Section 6.2 sketches a semantics for the behavioral types of facets, appealing to a type-level facet machine. In the same vein, section 7.2 informally presents a semantics for the types of dataspaces. This section supplements these with formal definitions. The definitions of the metafunctions are just like those from appendix B; their definitions are omitted.

Figure 24 defines the syntax for describing the states of a type-level facet machine and figure 25 the transition relation.

F.1 Machine Typing

Theorem 1, soundness for the type-level facet machine, needs matching type judgments.

The type judgment for facet machines, $\vdash_M M : M_T$:

$$\frac{\overrightarrow{\Gamma \vdash_I I : I_T} \quad \overrightarrow{\Gamma \vdash_{PS} PS : PS_T} \quad \Gamma \vdash_{FT} FT : FT_T \quad \vdash_\pi \pi : \pi_\tau \quad \vdash_\sigma \sigma : \sigma_\tau}{\vdash_M \langle FT; \vec{I}; \vec{PS}; \pi; \sigma \rangle : \langle FT_T; \vec{I}_T; \vec{PS}_T; \pi_\tau; \sigma_\tau \rangle} \boxed{\text{T-MACHINE}}$$

2945	$\langle \overrightarrow{\text{FT}}_{\tau}; \cdot; \cdot; \pi_{\tau}; \sigma_{\tau} \rangle \xrightarrow{\Delta_{\tau}} \langle \overrightarrow{\text{FT}}_{\tau}; \cdot; \overrightarrow{\text{PS}}_{\tau}; \pi'_{\tau}; \sigma_{\tau} \rangle$	(injectT)
2946	where	
2947	$\pi'_{\tau} = \pi_{\tau} \oplus_{\tau} \Delta_{\tau}$	
2948	$\overrightarrow{\text{PS}}_{\tau} = \text{dispatch}_{\tau}(\overrightarrow{\text{FT}}_{\tau}, \Delta_{\tau}, \pi_{\tau}, \pi'_{\tau}, \sigma_{\tau}, \langle \cdot \rangle)$	
2949	$\langle \overrightarrow{\text{FT}}_{\tau}; \cdot; (\text{fid}, T) \cdot \overrightarrow{\text{PS}}_{\tau}; \pi_{\tau}; \sigma_{\tau} \rangle \xrightarrow{\bullet} M'_{\tau}$	(transferT)
2950	where	
2951	$M'_{\tau} = \langle \overrightarrow{\text{FT}}_{\tau}; (\text{fid}, \overrightarrow{\text{I}}_{\tau}); \overrightarrow{\text{PS}}_{\tau}; \pi_{\tau}; \sigma'_{\tau} \rangle$	
2952	$\overrightarrow{\text{I}}_{\tau}, \sigma'_{\tau} = p\text{-}e_{\tau}(T, \sigma_{\tau}, \text{fid})$	
2953	$\langle \overrightarrow{\text{FT}}_{\tau}; \text{start fn } A_{\tau}(\overrightarrow{\text{D}}_{\tau} \overrightarrow{\text{T}}) @ \text{fid} \cdot \overrightarrow{\text{I}}_{\tau}; \overrightarrow{\text{PS}}_{\tau}; \pi_{\tau}; \sigma_{\tau} \rangle \xrightarrow{\bullet} M'_{\tau}$	(startT)
2954	where	
2955	$M'_{\tau} = \langle \overrightarrow{\text{FT}}'_{\tau}; \overrightarrow{\text{I}}_{\tau}; \overrightarrow{\text{PS}}_{\tau} \cdot \overrightarrow{\text{PS}}_{\tau \text{ start}} \cdot \overrightarrow{\text{PS}}_{\tau \text{ boot}} \cdot \overrightarrow{\text{PS}}_{\tau \text{ stop}}; \pi_{\tau}; \sigma_{\tau} \rangle$	
2956	$\text{fn}_{\text{new}} \text{ fresh in } \overrightarrow{\text{FT}}_{\tau}$	
2957	$\overrightarrow{\text{FT}}_{\tau \text{ new}} = \text{fn}_{\text{new}}[A_{\tau}(\overrightarrow{\text{D}}_{\tau} \overrightarrow{\text{T}}[\text{fn} \mapsto \text{fn}_{\text{new}}])].\varepsilon$	
2958	$\overrightarrow{\text{PS}}_{\tau \text{ start}} = \text{dispatch}_{\tau}(\overrightarrow{\text{FT}}_{\tau \text{ new}}, \text{start}, \emptyset, \emptyset, \sigma_{\tau}, \text{fid})$	
2959	$\overrightarrow{\text{PS}}_{\tau \text{ boot}} = \text{dispatch}_{\tau}(\overrightarrow{\text{FT}}_{\tau \text{ new}}, \pi_{\tau}/\emptyset, \emptyset, \pi_{\tau}, \sigma_{\tau}, \text{fid})$	
2960	$\overrightarrow{\text{FT}}'_{\tau} = \begin{cases} C_{\tau}[\overrightarrow{\text{FT}}_{\tau \text{ new}}] & \text{if } \text{locate}_{\tau}(\overrightarrow{\text{FT}}_{\tau}, \text{fid}) = C_{\tau} \\ \overrightarrow{\text{FT}}_{\tau} & \text{otherwise} \end{cases}$	
2961	$\overrightarrow{\text{PS}}_{\tau \text{ stop}} = \begin{cases} \cdot & \text{if } \text{locate}_{\tau}(\overrightarrow{\text{FT}}_{\tau}, \text{fid}) \text{ defined} \\ \text{dispatch}_{\tau}(\overrightarrow{\text{FT}}_{\tau \text{ new}}, \text{stop}, \pi_{\tau}, \pi_{\tau}, \sigma_{\tau}, \text{fid}) & \text{otherwise} \end{cases}$	
2962	$\langle \overrightarrow{\text{FT}}_{\tau}; \text{stop fn } \overrightarrow{\text{I}}_{\tau}; \overrightarrow{\text{PS}}_{\tau}; \pi_{\tau}; \sigma_{\tau} \rangle \xrightarrow{\bullet} \langle \overrightarrow{\text{FT}}'_{\tau}; \overrightarrow{\text{I}}_{\tau}; \overrightarrow{\text{PS}}_{\tau} \cdot \overrightarrow{\text{PS}}_{\tau \text{ stop}}; \pi_{\tau}; \sigma_{\tau} \rangle$	(stopT)
2963	where	
2964	$\overrightarrow{\text{FT}}_{\tau} = \begin{cases} C_{\tau}[\varepsilon] & \text{if } \overrightarrow{\text{FT}}_{\tau} = C_{\tau}[\text{fn}[\dots].\overrightarrow{\text{FT}}_{\tau}] \\ \overrightarrow{\text{FT}}_{\tau} & \text{otherwise} \end{cases}$	
2965	$\text{fid} = \text{facet-context}_{\tau}(C_{\tau})$	
2966	$\overrightarrow{\text{PS}}_{\tau \text{ stop}} = \begin{cases} \text{dispatch}_{\tau}(\text{fn}[\dots].\overrightarrow{\text{FT}}_{\tau}, \text{stop}, \pi_{\tau}, \pi_{\tau}, \sigma_{\tau}, \text{fid}) & \text{if } \overrightarrow{\text{FT}}_{\tau} = C_{\tau}[\text{fn}[\dots].\overrightarrow{\text{FT}}_{\tau}] \\ \cdot & \text{otherwise} \end{cases}$	
2967		
2968	$\langle \overrightarrow{\text{FT}}_{\tau}; \text{spawn } \overrightarrow{\text{T}} \cdot \overrightarrow{\text{I}}_{\tau}; \overrightarrow{\text{PS}}_{\tau}; \pi_{\tau}; \sigma_{\tau} \rangle \xrightarrow{\overline{\text{T}}} \langle \overrightarrow{\text{FT}}_{\tau}; \overrightarrow{\text{I}}_{\tau}; \overrightarrow{\text{PS}}_{\tau}; \pi_{\tau}; \sigma_{\tau} \rangle$	(spawnT)

Fig. 25: Type Machine Transitions

The *store-bindings* function creates a type environment from a store:

$$\begin{aligned} \text{store-bindings} & : \sigma_{\tau} \longrightarrow \Gamma \\ \text{store-bindings}(\emptyset) & = \cdot \\ \text{store-bindings}(\sigma_{\tau} \uplus x \mapsto \tau) & = \text{store-bindings}(\sigma_{\tau}), x : \tau \end{aligned}$$

The type judgment for facet trees, $\Gamma \vdash_{\text{FT}} \text{FT} : \overrightarrow{\text{FT}}_{\tau}$:

$$\frac{}{\Gamma \vdash_{\text{FT}} \varepsilon : \varepsilon} \boxed{\text{FT-EMPTY}}$$

$$\frac{\Gamma \vdash_A A : A_{\tau} \quad \overrightarrow{\Gamma} \vdash_D D : \overrightarrow{\text{D}}_{\tau}}{\frac{\Gamma, \text{fn} : \text{FacetName}, \text{bindings}_D(D) \vdash_{\text{Pr}} \text{Pr} : \overrightarrow{\text{T}} \quad \overrightarrow{\Gamma} \vdash_{\text{FT}} \text{FT} : \overrightarrow{\text{FT}}_{\tau}}{\Gamma \vdash_{\text{FT}} \text{fn}[\overrightarrow{e}(D \text{ Pr})].\overrightarrow{\text{FT}} : \text{fn}[A_{\tau}(\overrightarrow{\text{D}}_{\tau} \overrightarrow{\text{T}})].\overrightarrow{\text{FT}}_{\tau}} \boxed{\text{FT-TREE}}}$$

The type judgment for machine instructions, $\Gamma \vdash_{\text{I}} I : \overrightarrow{\text{I}}_{\tau}$:

2991

2992

2993

2994

2995

2996

2997

2998

2999

3000

3001

3002

3003

3004

3005

3006

3007

3008

3009

3010

3011

3012

3013

3014

3015

3016

3017

3018

3019

3020

3021

3022

3023

3024

3025

3026

3027

3028

3029

3030

3031

3032

3033

3034

3035

3036

$$\frac{\Gamma \vdash_A A : A_\tau \quad \overline{\Gamma, \text{bindings}_D(D) \vdash_{Pr} Pr : T}}{\Gamma \vdash_I \text{start fn } \vec{e} \ (D \ \vec{Pr}) \ @ \ \text{fid} : \text{start fn } A_\tau \ (D_T \ \vec{T}) \ @ \ \text{fid}} \boxed{\text{I-START}}$$

$$\frac{}{\Gamma \vdash_I \text{stop fn} : \text{stop fn}} \boxed{\text{I-STOP}} \quad \frac{\text{prune}(\Gamma) \vdash_{Pr} Pr : T}{\Gamma \vdash_I \text{spawn } Pr : \text{spawn } \vec{T}} \boxed{\text{I-SPAWN}}$$

The type judgment for pending scripts, $\Gamma \vdash_{PS} PS : PS_T$:

$$\frac{\Gamma \vdash_{Pr} Pr : T}{\Gamma \vdash_{PS} (\text{fid}, Pr) : (\text{fid}, T)} \boxed{\text{PS-SCRIPT}}$$

The type judgment for assertion sets, $\vdash_\pi \pi : \pi_\tau$:

$$\frac{\pi_\tau = \{\tau_R \mid u \in \pi, \vdash_{\tau_R} u : \tau_R\}}{\vdash_\pi \pi : \pi_\tau} \boxed{\pi\text{-SET}}$$

The type judgment for assertions, $\vdash_{\tau_R} u : \tau_R$:

$$\frac{}{\vdash_{\tau_R} b : B} \boxed{\text{R-BASE}} \quad \frac{\overline{\vdash_{\tau_R} u : \tau_R}}{\vdash_{\tau_R} m(\vec{u}) : m(\vec{\tau}_R)} \boxed{\text{R-TUPLE}} \quad \frac{\vdash_{\tau_R} u : \tau_R}{\vdash_{\tau_R} ? u : ? \tau_R} \boxed{\text{R-?}}$$

The type judgment for patches, $\vdash_\Delta \Delta : \Delta_\tau$:

$$\frac{\vdash_\pi \pi^+ : \pi_\tau^+ \quad \vdash_\pi \pi^- : \pi_\tau^-}{\vdash_\Delta \pi^+ / \pi^- : \pi_\tau^+ / \pi_\tau^-} \boxed{\Delta\text{-PATCH}}$$

The type judgment for events, $\vdash_{Evt} Evt : Evt_\tau$:

$$\frac{\vdash_\Delta \Delta : \Delta_\tau}{\vdash_{Evt} \Delta : \Delta_\tau} \boxed{\text{EVT-PATCH}} \quad \frac{}{\vdash_{Evt} \text{start} : \text{start}} \boxed{\text{EVT-START}}$$

$$\frac{}{\vdash_{Evt} \text{stop} : \text{stop}} \boxed{\text{EVT-STOP}}$$

The type judgment for stores, $\vdash_\sigma \sigma : \sigma_\tau$:

$$\frac{\sigma_\tau = \{x \mapsto \tau \mid x \mapsto v \in \sigma, \cdot \vdash_e v : \tau\}}{\vdash_\sigma \sigma : \sigma_\tau} \boxed{\sigma\text{-SET}}$$

The type judgment for substitutions, $\vdash_\gamma \gamma : \gamma_\tau$:

$$\frac{\gamma_\tau = \{x \mapsto \tau \mid x \mapsto v \in \gamma, \cdot \vdash_e v : \tau\}}{\vdash_\gamma \gamma : \gamma_\tau} \boxed{\gamma\text{-SUB}}$$

The type judgment for transition labels, $\vdash_l l : l_\tau$:

$$\frac{}{\vdash_l \bullet : \bullet} \boxed{\text{L-NONE}} \quad \frac{\vdash_{\text{Evt}} \Delta : \Delta_\tau}{\vdash_l \Delta : \Delta_\tau} \boxed{\text{L-EVT}} \quad \frac{\cdot \vdash_{\text{Pr}} \text{Pr} : \text{T}}{\vdash_l \overline{\text{Pr}} : \overline{\text{T}}} \boxed{\text{L-SPAWN}}$$

F.2 Type-level Dataspaces

A type-level dataspace DS_τ has nearly the same semantics as the one from section 7, with two modifications. First, assertions range over types τ . Second, actor behaviors are defined in terms of facet machine types, defined above, with related metafunctions such as interp_M lifted to the type level, mutatis mutandis.

G Proof Details

The following lemmas establish that the machine transitions for actors and types are related via typing, under certain conditions. The transitions fall roughly in the following categories:

- initialization (lemmas 13);
- transitions in response to an external stimulus (lemmas 15, 16, 17); and
- transitions for performing internal work (lemmas 18, 19).

Finally, lemma 20 establishes that related type and term machines make related assertions.

Lemma 13 (Boot). *If*

- $\Gamma \vdash_{\text{Pr}} \text{Pr} : \text{T}$
- $\text{boot}_{\text{Pr}}(\text{Pr}, \sigma) = \langle \text{FT}; \overrightarrow{\text{I}}; \overrightarrow{\text{PS}}; \pi; \sigma \rangle$

then

- $\text{boot}_{\text{T}}(\text{T}) = \langle \text{FT}_{\text{T}}; \overrightarrow{\text{I}}_{\text{T}}; \overrightarrow{\text{PS}}_{\text{T}}; \pi_{\tau}; \emptyset \rangle$
- $\vdash_M \langle \text{FT}; \overrightarrow{\text{I}}; \overrightarrow{\text{PS}}; \pi; \sigma \rangle : \langle \text{FT}_{\text{T}}; \text{S}_{\text{T}}; \overrightarrow{\text{PS}}_{\text{T}}; \pi_{\tau}; \text{type-store}(\sigma) \rangle$

Proof By induction on the type derivation and the soundness of the *dispatch* metafunction (lemma 16). \square

Definition 14 (Store Typing).

$$\begin{aligned} \text{type-store} & : \sigma \longrightarrow \sigma_\tau \\ \text{type-store}(\emptyset) & = \emptyset \\ \text{type-store}(\sigma[x \mapsto v]) & = \text{type-store}(\sigma)[x \mapsto \tau] \\ & \text{where } \cdot \vdash_e v : \tau \end{aligned}$$

In other words, initialization of related actor terms and types yields related term and type machines.

Lemma 15 (Matching Events). *If*

- $\vdash_e v : \tau$
- $\vdash_e u : \tau'$

then $match_v(v, u) = \gamma$ iff $match_\tau(\tau, \tau') = \gamma_\tau$ with $\vdash_\gamma \gamma : \gamma_\tau$.

Proof By induction on the pattern type derivation. □

Essentially, if a pattern matches an assertion then the pattern's type matches the assertion's type, yielding related substitutions, and vice versa.

Lemma 16 (Dispatch). *If*

- $fid = \langle \overrightarrow{fn} \rangle$
- $\Gamma \vdash_D D : D_T$
- $\Gamma, \overrightarrow{fn} : \text{FacetName}, \text{bindings}_D(D) \vdash_{Pr} Pr : T$
- $\vdash_{Evt} Evt : Evt_\tau$
- $\vdash_\pi \pi : \pi_\tau$
- $\vdash_\pi \pi' : \pi'_\tau$

then

$$dispatch1(fid, D, Pr, Evt, \pi, \pi', \sigma) = \overrightarrow{PS}$$

iff

$$dispatch1_T(fid, D_T, T, Evt_\tau, \pi_\tau, \pi'_\tau) = \overrightarrow{PS}_T$$

with $\Gamma, \overrightarrow{fn} : \text{FacetName} \vdash_{PS} PS : PS_T$

Proof By induction on the type derivations. The dispatch lemma (lemma 16) lifts to events in general, so the term and type applications of $match_D$ yield related substitutions. When applied to related event-handler bodies, the related substitutions yield related scripts. □

In other words, dispatching related events yields related scripts to execute.

Lemma 17 (Partial Evaluation). *If* $\Gamma \vdash_{Pr} Pr : T$ and $\text{safe}(Pr)$ then $p\text{-e}(Pr, \sigma, fid) = \overrightarrow{I}, \sigma'$ iff $p\text{-e}_T(T, \sigma_\tau, fid) = \overrightarrow{I}_T, \sigma'_\tau$ with $\overrightarrow{\Gamma} \vdash_I I : I_T$

Proof By induction on the type derivation. □

That is, partially-evaluating a pending script and its type yields related instructions.

Lemma 18 (Preservation). *If* $\vdash_M M : M_T$ and $\text{safe}(M)$ then $M \rightarrow M' = \langle FT; \overrightarrow{I}; \overrightarrow{PS}; \pi; \sigma \rangle$ iff $M_T \rightarrow \langle FT_T; \overrightarrow{I}_T; \overrightarrow{PS}_T; \pi_\tau; \sigma_\tau \rangle$ with $\vdash_M M' : \langle FT_T; \overrightarrow{I}_T; \overrightarrow{PS}_T; \pi_\tau; \text{type-store}(\sigma) \rangle$

3129 *Proof* Following a standard approach (Wright and Felleisen, 1994), we can show that
 3130 at most one of the machine transition rules can apply. Through the application of related
 3131 lemmas, such as lemmas 16 and 17, the updated parts of the machine state remain related
 3132 via typing. \square

3133 In other words, if a term machine takes a transition, then its type can take a transition to
 3134 a related type machine, up to the type of the store. Since type machines do not manipulate
 3135 or depend on the store at all, the store from the destination term-level machine state is
 3136 translated to a new store type for the type-level.

3137 **Lemma 19** (Progress). *If $\vdash_M M : M_T$ either $\text{inert}(M)$, $\text{inert}(M_T)$ or there exists*
 3138 *M', M'_T such that $M \xrightarrow{l\bullet} M'$ and $M_T \xrightarrow{l\tau\bullet} M'_T$.*

3140 *Proof* By inspection of the machine state. \square

3142 **Lemma 20** (Machine Assertions). *If $\vdash_M M : M_T$ then*

$$3144 \vdash_{\pi} \text{assertions-of}_M(M) : \text{assertions-of}_{M_T}(M_T)$$

3145 *Proof* By induction on the type derivation and via similar properties for the *assertions-of*
 3146 family of functions. \square

3148 This establishes that related term and type machines make the same assertions.

3149
3150
3151
3152
3153
3154
3155
3156
3157
3158
3159
3160
3161
3162
3163
3164
3165
3166
3167
3168
3169
3170
3171
3172
3173
3174