

Model-View-Update-Communicate: Session Types meet the Elm Architecture

Simon Fowler 

University of Edinburgh, Scotland
simon.fowler@ed.ac.uk

Abstract

Session types are a type discipline for communication channel endpoints which allow conformance to protocols to be checked statically. Safely implementing session types requires linearity, usually in the form of a linear type system. Unfortunately, linear typing is difficult to integrate with graphical user interfaces (GUIs), and to date most programs using session types are command line applications.

In this paper, we propose the first principled integration of session typing and GUI development by building upon the Model-View-Update (MVU) architecture, pioneered by the Elm programming language. We introduce λ_{MVU} , the first formal model of the MVU architecture, and prove it sound. By extending λ_{MVU} with *commands* as found in Elm, along with *linearity* and *model transitions*, we show the first formal integration of session typing and GUI programming. We implement our approach in the Links web programming language, and show examples including a two-factor authentication workflow and multi-room chat server.

2012 ACM Subject Classification Software and its engineering → Concurrent programming languages

Keywords and phrases Session types, concurrent programming, Model-View-Update

Digital Object Identifier 10.4230/LIPIcs.ECOOP.2020.14

Related Version An extended version of the paper is available on arXiv (<https://arxiv.org/abs/1910.11108>).

Funding This work was supported by ERC Consolidator Grant Skye (grant no. 682315) and an ISCF Metrology Fellowship grant provided by the UK government's Department for Business, Energy and Industrial Strategy (BEIS).

Acknowledgements I thank Jake Browning for sparking my interest in Elm and for his help with an early prototype of the Links MVU library; Sára Decova for a previous version of the multi-room chat server example; Sam Lindley for many useful discussions and suggestions; and James Cheney, April Gonçalves, and the anonymous ECOOP PC and AEC reviewers for detailed comments.

1 Introduction

Modern applications are necessarily concurrent and distributed. Along with concurrency and distribution naturally comes communication, but communication protocols are typically informally described, resulting in costly runtime failures and code maintainability issues.

Session types [23, 24] are a type discipline for communication channel endpoints which allow conformance to a protocol to be checked statically rather than after an application is deployed. Many distributed GUI applications, such as chat applications or multiplayer games, would benefit from session-typed communication with a server. Unfortunately, safely implementing session types requires a linear type system, but safely integrating linear resources and GUIs is nontrivial. As a consequence, to date most programs using session types are batch-style applications run on the command line.

The lack of a principled integration of GUI applications and session types is a significant barrier to their adoption. In this paper, we bridge this gap by extending the Model-View-Update (MVU) architecture, pioneered by the Elm programming language, to support linear



© Simon Fowler;

licensed under Creative Commons License CC-BY

34th European Conference on Object-Oriented Programming (ECOOP 2020).

Editors: Robert Hirschfeld and Tobias Pape; Article No. 14; pp. 14:1–14:29

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

14:2 Model-View-Update-Communicate

resources. We present λ_{MVU} , a core formalism of the MVU architecture, and an extended version of λ_{MVU} which supports session-typed communication. Informed by the formal development, we provide a practical implementation in the Links programming language [10].

Session types by example. Let us consider a two-factor authentication workflow, introduced by Fowler et al. [20]. A user first enters their credentials. If correct, the server can then either grant access, or send a challenge key. If challenged, the user enters the challenge code into a hardware token, which generates a response to be entered into the web page. The server then either authenticates the user or denies access.

We can describe the two-factor authentication example as a session type as follows:

$$\begin{array}{ll} \text{TwoFactorServer} \triangleq & \text{TwoFactorClient} \triangleq \\ ?(\text{Username}, \text{Password}).\oplus\{ & !(\text{Username}, \text{Password}).\&\{ \\ \quad \text{Authenticated} : \text{ServerBody}, & \quad \text{Authenticated} : \text{ClientBody}, \\ \quad \text{Challenge} : !\text{ChallengeKey}.\text{?Response}. & \quad \text{Challenge} : \text{?ChallengeKey}!\text{Response}. \\ \quad \oplus\{\text{Authenticated} : \text{ServerBody}, & \quad \&\{\text{Authenticated} : \text{ClientBody}, \\ \quad \quad \text{AccessDenied} : \text{End}\}, & \quad \text{AccessDenied} : \text{End}\}, \\ \quad \text{AccessDenied} : \text{End}\} & \quad \text{AccessDenied} : \text{End}\} \end{array}$$

The `TwoFactorServer` type shows the session type for the server, which firstly receives (?) the credentials from the client, and then chooses (\oplus) whether to authenticate, deny access, or issue a challenge. If the server issues a challenge, it sends (!) the challenge string, awaits the response, and then chooses whether to accept or reject the request. The `ServerBody` type abstracts over the actions performed in the remainder of the application, for example taking out a loan. The `TwoFactorClient` type is the *dual* of the `TwoFactorServer` type: where the server sends, the client receives, and where the client sends, the server receives. The $\&$ construct denotes offering a choice of branches. Suppose we have constructs for sending along, receiving from, and closing an endpoint:

$$\text{send} : (A \times !A.S) \rightarrow S \qquad \text{receive} : ?A.S \rightarrow (A \times S) \qquad \text{close} : \text{End} \rightarrow \mathbf{1}$$

Let us also suppose we have constructs for selecting and offering a choice:

$$\begin{array}{ll} \text{select } \ell_j M : S_j & \text{where } M \text{ has session type } \oplus\{\ell_i : S_i\}_{i \in I}, \text{ and } j \in I \\ \text{offer } M \{\ell_i(x_i) \mapsto N_i\}_{i \in I} : A & \text{where } M \text{ has session type } \&\{\ell_i : S_i\}_{i \in I}, \text{ each } x_i \text{ binds an} \\ & \text{endpoint with session type } S_i, \text{ and each } N_i \text{ has type } A \end{array}$$

We can write a server implementation as follows:

```
twoFactorServer : TwoFactorServer → 1
twoFactorServer(s) ≜ let ((username, password), s) = receive s in
  if checkDetails(username, password) then
    let s = select Authenticated s in serverBody(s)
  else let s = select AccessDenied s in close s
```

To implement session-typed communication safely, we require a linear type system [44] to ensure each communication endpoint is used exactly once: as an example, without linearity it would be possible to attempt to receive the credentials twice.

Linearity and GUIs. We can also write a client application:

```
twoFactorClient : (Username × Password × TwoFactorClient) → 1
twoFactorClient(username, password, s) ≜
  let s = send ((username, password), s) in
  offer s {Authenticated(s) ↦ clientBody(s)
          Challenge(s) ↦ let (key, s) = receive s in
                        let s = send (generateResponse(key), s) in
                        offer s {Authenticated(s) ↦ clientBody(s)
                                AccessDenied(s) ↦ close s; loginFailed}
          AccessDenied(s) ↦ close s; loginFailed}
```

However, such a client is of little use, as it sends only a pre-defined set of credentials, and the step where a user enters the response to the challenge is replaced by a function `generateResponse`. Ideally, we would like the credentials to be entered into a GUI, and for a button press to trigger the session communication with the server.

Let us attempt to write a GUI for the first stage of the two-factor authentication example; as HTML is well-understood, we concentrate on web pages in the remainder of the paper.

```
render(c) ≜
  <html>
    <body>
      <input id = "username"></input>
      <input id = "password"></input>
      <button onClick = login(c)>Submit</button>
    </body>
  </html>

login(c) ≜ λ().
  let user = getContents("username") in
  let pass = getContents("password") in
  let c = send ((user, pass), c) in
  handleResponse(c)
```

Given a channel c of type `TwoFactorClient`, the `render` function generates a web page with input boxes for the username and password, and a button to submit the credentials. The `login` function, triggered when the button is clicked, retrieves the username and password from the two input boxes, and sends the credentials along c . The `handleResponse` function, which we omit, receives the response from the server and updates the web page.

On first inspection, this implementation seems sound since the endpoint c is used linearly. However, the above attempt is unsound due to the asynchronous nature of GUI programming: there is nothing stopping the user pressing the button twice and sending the credentials twice along c , in contravention of the session type. As a further complication, suppose we augmented the protocol with a “forgotten password” branch, triggered by another button. This would require two instances of c in the GUI, again violating linearity:

```
<button onClick = login(c)>Submit</button>
<button onClick = reset(c)>Reset password</button>
```

It is clear that directly embedding linear resources into a GUI is a non-starter. A more successful approach involves spawning a separate process which contains the linear resource, and which receives *non-linear* messages from the GUI. Upon receiving a GUI message, the process can then perform the session communication, while ignoring duplicate GUI messages:

```
render(c) ≜
  let pid = spawn handler(c) in
  <html>
    <body>
      <input id = "username"></input>
      <input id = "password"></input>
      <button onClick = login(pid)>Submit</button>
    </body>
  </html>

login(pid) ≜ λ().
  let user = getContents("username") in
  let pass = getContents("password") in
  pid! SubmitLogin(user, pass)

handler(c) ≜
  case (get ()) {
    SubmitLogin(user, pass) ↦
      let c = send ((user, pass), c) in
      handleResponse(c)
  }
```

The `render` function begins by spawning `handler(c)` as a separate process with an incoming message queue (or *mailbox*), returning the process ID pid . As before, the `login` function is triggered by pressing the button, and retrieves the credentials from the web page. Instead of communicating on the channel directly, it sends a `SubmitLogin` message containing the credentials to the process ID of handler process, written $pid! `SubmitLogin(user, pass)`. The handler process retrieves the message from its mailbox (`get ()`), and can then communicate with the server over the linear endpoint. Such an approach also scales to the “forgotten password” extension, by adding another GUI message.$

The above approach is used by Fowler et al. [20], who provide the first integration of session types and web application development, including the ability to gracefully handle failures such as the user closing their browser mid-session. Unfortunately, the approach is

14:4 Model-View-Update-Communicate

brittle and ad-hoc. All interaction with the web page occurs using imperative operations such as `getContents` and `setContents`; contrary to best practices such as the Model-View-Controller (MVC) [30] pattern, the state of the web page is not derived directly from the data contained by the application. Furthermore, there is no connection between the state of the handler process and what is displayed on the web page: this can easily lead to mismatches between the possible GUI messages which can be sent and which can be handled.

Model-View-Update. This paper is about doing better. Our approach is to formalise Model-View-Update, an architectural pattern for GUI development popularised by the Elm programming language [1], and extend it to support linear resources. MVU is an appealing starting point as it is particularly suited to functional programming. Furthermore, MVU has directly inspired popular technologies such as Redux [5] and the Flux architecture [4], which are used with the popular React [2] frontend web framework for JavaScript.

The Elm programming language [1] is a functional programming language designed for writing web applications. Elm was originally designed to use *functional reactive programming* (FRP) [14], where time-varying *signals* can be used to construct reactive web applications. A paper describing Elm, and its core formal semantics, was published at PLDI 2013 [12].

For many languages, that would be the end of the story. But unusually for a research language, Elm gained a user community, and a standard architectural pattern known as *The Elm Architecture* grew organically to such a point that Elm abandoned FRP altogether [11]. At its core, The Elm Architecture is a descendant of MVC where a *model* contains the state of the application; a *view function* renders the model; and the rendered model produces *messages* which are handled by an *update* function to produce a new model. More generally, this pattern has been referred to as *Model-View-Update*, or MVU for short [3, 40].

Consider the following web application, where a user enters text into a text box, and the application displays the text, reversed:

```
hello
olleh
```

We can write this example using MVU as follows:

```
Model ≜ (contents : String)
Message ≜ UpdateBox(String)

model : Model
model ≜ (contents = "")

update : (Message × Model) → Model
update ≜ λ(UpdateBox(str), m).(contents = str)

view : Model → Html(Message)
view ≜ λmodel.html
  <input type = "text" value = {model.contents}
    onInput = {λstr.UpdateBox(str)}></input>
  <div>
    {htmlText (reverseString (model.contents))}
  </div>
(model, view, update)
```

We define two type aliases: the `Model` captures the state of the application and is defined as a record with a single `String` field, `contents`. `Messages` are produced as a result of user interaction. The `Message` type is defined as a singleton variant type with constructor `UpdateBox`, containing the updated value of the text box.

The view function renders a model. It has the type `Model → Html(Message)`, which is a function taking a `Model` as its argument, and returning HTML which may produce messages of type `Message`. The `value = {model.contents}` attribute of the `input` box states that the contents of the text box should reflect the `contents` field of the model. The `onInput` attribute is an *event handler*: its body is a function taking the current value of the input box (`str`) and producing an `UpdateBox` message containing the updated contents of the box. The contents of the div tag are derived from the reversed contents.

Syntax		
Types	$A, B, C ::= \mathbf{1} \mid A \rightarrow B \mid A \times B \mid A + B \mid \mathbf{String} \mid \mathbf{Int}$ $\mid \mathbf{Html}(A) \mid \mathbf{Attr}(A)$	
String literals	s	
Integers	n	
Terms	$L, M, N ::= x \mid \lambda x.M \mid \mathbf{rec} f(x). M \mid M N \mid () \mid s \mid n$ $\mid (M, N) \mid \mathbf{let} (x, y) = M \mathbf{in} N$ $\mid \mathbf{inl} x \mid \mathbf{inr} x \mid \mathbf{case} L \{ \mathbf{inl} x \mapsto M; \mathbf{inr} y \mapsto N \}$ $\mid \mathbf{htmlTag} t M N \mid \mathbf{htmlText} M \mid \mathbf{htmlEmpty}$ $\mid \mathbf{attr} ak M \mid \mathbf{attrEmpty} \mid M \star N$	
Tag names	t	
Attribute names	at	
Attribute keys	$ak ::= at \mid h$	
Event handler names	h	
Typing rules for terms $\Gamma \vdash M : A$		
$\frac{\mathbf{T-HtmlTAG} \quad \Gamma \vdash M : \mathbf{Attr}(A) \quad \Gamma \vdash N : \mathbf{Html}(A)}{\Gamma \vdash \mathbf{htmlTag} t M N : \mathbf{Html}(A)}$	$\frac{\mathbf{T-HtmlTEXT} \quad \Gamma \vdash M : \mathbf{String}}{\Gamma \vdash \mathbf{htmlText} M : \mathbf{Html}(A)}$	$\frac{\mathbf{T-HtmlEMPTY}}{\Gamma \vdash \mathbf{htmlEmpty} : \mathbf{Html}(A)}$
$\frac{\mathbf{T-Attr} \quad \Gamma \vdash M : \mathbf{String}}{\Gamma \vdash \mathbf{attr} at M : \mathbf{Attr}(A)}$	$\frac{\mathbf{T-EvtAttr} \quad \Gamma \vdash M : \mathbf{ty}(h) \rightarrow A}{\Gamma \vdash \mathbf{attr} h M : \mathbf{Attr}(A)}$	$\frac{\mathbf{T-AttrEmpty}}{\Gamma \vdash \mathbf{attrEmpty} : \mathbf{Attr}(A)}$
$\frac{\mathbf{T-HtmlAppend} \quad \Gamma \vdash M : \mathbf{Html}(A) \quad \Gamma \vdash N : \mathbf{Html}(A)}{\Gamma \vdash M \star N : \mathbf{Html}(A)}$	$\frac{\mathbf{T-AttrAppend} \quad \Gamma \vdash M : \mathbf{Attr}(A) \quad \Gamma \vdash N : \mathbf{Attr}(A)}{\Gamma \vdash M \star N : \mathbf{Attr}(A)}$	

■ **Figure 1** Syntax and typing rules for λ_{MVU} terms

The `update` function takes a message and previous model as its arguments, and produces a new model. In this case, the `update` function constructs a new model where the `contents` field is set to the payload of the `UpdateBox` message. Finally, the program is a 3-tuple containing the initial model, and the view and update functions.

To achieve our goal of a formal integration of session typing and GUI programming, we first formalise MVU, and then generalise the architecture to support linear models and messages. Supporting linearity poses some challenges, as we will see in §3.

1.1 Contributions.

The overarching contribution of this paper is the first principled integration of session-typed communication with a GUI framework. Concretely, we make three contributions:

1. We introduce the first formal model of the MVU architecture, λ_{MVU} (§2). We prove (§2.3) that λ_{MVU} satisfies preservation and event progress properties.
2. We extend λ_{MVU} with *commands*, *linearity*, and *model transitions* (§3), which allow λ_{MVU} to support GUIs incorporating session-typed communication, and we prove the soundness of the extended calculus.
3. We implement the architecture in the Links web programming language. We show an extended example of a chat application where client code uses the linear MVU framework, and where client-server communication happens over session-typed channels (§4).

The implementation and examples are available in the paper’s companion artifact.

Event name ev	Event Handler h ($handler(ev)$)	Payload type ($ty(ev)$, $ty(h)$)	Payload Description
click	onClick	1	Unit value
input	onInput	String	Updated contents of a text field
keyUp	onKeyUp	Int	Key code
keyDown	onKeyDown	Int	Key code

■ **Figure 2** Example event signatures

2 Model-View-Update, Formally

In this section, we formalise MVU as a core calculus, λ_{MVU} , an extension of the simply-typed λ -calculus with products, sums, HTML, and event handling. Even without extensions, λ_{MVU} is expressive enough to support many common applications such as form handling.

2.1 Syntax

Types. Figure 1 shows the syntax and typing rules for λ_{MVU} . Types are ranged over by A, B, C , and consist of the unit type **1**, functions $A \rightarrow B$, products $A \times B$, sums $A + B$, and string and integer types. Types $Html(A)$ and $Attr(A)$ are the type of HTML elements and attributes which can produce messages of type A .

Terms. Terms, ranged over by L, M, N , include variables, λ abstractions, anonymous recursive functions, function application, the unit value, string literals, integers, and sum and pair introduction and elimination. The remaining terms encode HTML *elements* and *attributes*. The **htmlTag** $t M N$ construct denotes an HTML element with tag name t (for example, `div`), attributes M , and children N ; the **htmlText** M construct describes a text node with text M ; and **htmlEmpty** defines an empty HTML node.

The **attr** $ak M$ construct describes an attribute with key ak and body M , where the key ak is either an attribute name at or an event handler name h . The **attrEmpty** construct defines an empty attribute.

The $M \star N$ operator appends two HTML elements or attributes. Since both HTML elements and attributes support a unit element (**htmlEmpty** and **attrEmpty** respectively), elements and attributes together with \star form two monoids.

Events. We model interaction with the Document Object Model (DOM) through *events*, which model those dispatched by a browser. An *event signature* is a 3-tuple (ev, h, A) consisting of an event name ev , handler name h , and payload type A . We require a bijective mapping between event and handler names. Figure 2 describes example event signatures used in the remainder of the paper. We consider four primitive events: **click**, which is fired when an element is clicked; **input**, which is fired when the contents of a text field are changed; and **keyUp** and **keyDown**, which are fired when a key is pressed while focused on an element.

Event handlers are attached to elements as attributes, and generate a message in response to an event. We write $handler(ev)$ to refer to the handler for ev : for example, $handler(click) = onClick$. We write $ty(ev)$ to refer to the payload type of ev and write $ty(h)$ for the payload type of an event handled by h . As an example, both $ty(click) = \mathbf{1}$ and $ty(onClick) = \mathbf{1}$.

Term typing. Term typing rules for λ -calculus constructs are standard, so are omitted. Rule T-HTMLTAG states that **htmlTag** $t M N$ can be given type $Html(A)$ if its attributes M have type $Attr(A)$ and children have type $Html(A)$. Text nodes **htmlText** M do not produce

Values	$U, V, W ::= \lambda x.M \mid \mathbf{rec} f(x).M \mid () \mid (V, W) \mid \mathbf{inl} V \mid \mathbf{inr} V \mid s \mid n$ $\mid \mathbf{htmlTag} \ t \ V \ W \mid \mathbf{htmlEmpty} \mid \mathbf{htmlText} V$ $\mid \mathbf{attr} \ ak \ V \mid \mathbf{attrEmpty} \mid V \star W$
Events	$e ::= \mathbf{ev}(V)$
DOM Pages	$D ::= \mathbf{domTag}(\vec{e}) \ t \ V \ D \mid \mathbf{domText} V \mid \mathbf{domEmpty} \mid D \star D'$
Active thread	$T ::= \mathbf{idle} V_m \mid M$
Function state	$F ::= (V_v, V_u)$
Processes	$P, Q ::= \mathbf{run} M \mid \langle T \mid F \rangle \mid ((M)) \mid P \parallel Q$
Configurations	$\mathcal{C} ::= P \S D$
Process contexts	$\mathcal{P} ::= [] \mid \mathcal{P} \parallel P$
DOM contexts	$\mathcal{D} ::= [] \mid \mathbf{domTag}(\vec{e}) \ t \ V \ \mathcal{D} \mid \mathcal{D} \star D \mid D \star \mathcal{D}$
Thread contexts	$\mathcal{T} ::= \mathbf{run} E \mid \langle E \mid F \rangle \mid ((E))$

■ **Figure 3** Runtime syntax for λ_{MVU}

any messages, and so have type $\mathbf{Html}(A)$ if M has type \mathbf{String} (T-HTMLTEXT); similarly, $\mathbf{htmlEmpty}$ has type $\mathbf{Html}(A)$ (T-HTMLEMPTY).

Rule T-ATTR assigns attributes $\mathbf{attr} \ at \ M$ type $\mathbf{Attr}(A)$ for any A if M has type \mathbf{String} . Rule T-EVTATTR types event handler attributes $\mathbf{attr} \ h \ M$: if the event handler M has type $\mathbf{ty}(h) \rightarrow A$ (i.e., it *produces messages of type A*), then the attribute can be given type $\mathbf{Attr}(A)$. Finally, T-ATREEMPTY states that the empty attribute $\mathbf{attrEmpty}$ has type $\mathbf{Attr}(A)$ for any type A . We overload the \star operator to append both HTML elements and attributes (T-HTMLAPPEND and T-ATTRAPPEND).

Syntactic sugar. We assume the usual encodings of records as pairs and variant types as binary sums, and use pattern matching notation. It is useful to be able to write HTML using XML-like notation, where an *antiquoted expression* $\{M\}$ allows a term M to be embedded within an HTML tree. The view function from the introduction desugars to:

$$\lambda model.$$

$$(\mathbf{htmlTag} \ \mathbf{input}$$

$$((\mathbf{attr} \ \mathbf{type} \ \text{"text"}) \star (\mathbf{attr} \ \mathbf{value} \ model.\mathbf{contents}) \star$$

$$(\mathbf{attr} \ \mathbf{onInput} \ (\lambda str.\mathbf{UpdateBox}(str)))) \ \mathbf{htmlEmpty}) \star$$

$$\mathbf{htmlTag} \ \mathbf{div} \ \mathbf{attrEmpty} \ (\mathbf{htmlText} \ \mathbf{reverseString} \ (model.\mathbf{contents}))$$

The formal definitions and desugaring translations are unsurprising; the details can be found in the extended version [18].

2.2 Operational Semantics

We can now provide λ_{MVU} with a small-step operational semantics.

2.2.1 Runtime Syntax

Figure 3 describes the runtime syntax of λ_{MVU} . Values, ranged over by U, V, W , are standard. An event $\mathbf{ev}(V)$ consists of event name \mathbf{ev} and payload V . We write ϵ for an empty meta-level sequence, and use \cdot for sequence concatenation. DOM pages, ranged over by D , are the runtime representation of HTML, where tags $\mathbf{domTag}(\vec{e}) \ t \ V \ D$ contain an event queue \vec{e} of events dispatched to the element.

Concurrency. Concurrency is vital when modelling GUI applications as event handling is asynchronous: computation triggered by a user interaction should not block the UI.

Meta-level definitions	
$\text{handle}(m, (v, u), \text{msg}) \triangleq$ $\text{let } m' = u(\text{msg}, m) \text{ in}$ $(m', v m')$	$\text{handlers}(\text{ev}, \mathbf{attrEmpty}) = \epsilon$ $\text{handlers}(\text{ev}, V \star W) = \text{handlers}(\text{ev}, V) \cdot \text{handlers}(\text{ev}, W)$ $\text{handlers}(\text{ev}, \mathbf{attr } at V) = \epsilon$ $\text{handlers}(\text{ev}, \mathbf{attr } h V) = \begin{cases} V & \text{if handler}(\text{ev}) = h \\ \epsilon & \text{otherwise} \end{cases}$
Process reduction	
$\text{EP-HANDLE } \langle \mathbf{idle } V_m \mid F \rangle \parallel ((V)) \longrightarrow \langle \text{handle}(V_m, F, V) \mid F \rangle$ $\text{EP-PAR } P_1 \parallel P_2 \longrightarrow P'_1 \parallel P'_2 \quad \text{if } P_1 \longrightarrow P'_1$ $\text{EP-LIFTT } \mathcal{T}[M] \longrightarrow \mathcal{T}[N] \quad \text{if } M \longrightarrow_M N$	<div style="border: 1px solid black; padding: 2px; display: inline-block;">$P \longrightarrow P'$</div>
Configuration reduction	
$\text{E-RUN } \mathcal{P}[\mathbf{run}(V_m, V_v, V_u)] \S D \longrightarrow \mathcal{P}[\langle (V_m, V_v V_m) \mid (V_v, V_u) \rangle] \S D$ $\text{E-UPDATE } \mathcal{P}[\langle (V_m, U) \mid F \rangle] \S D \longrightarrow \mathcal{P}[\langle \mathbf{idle } V_m \mid F \rangle] \S D' \quad \text{where } \text{diff}(U, D) = D'$ $\text{E-INTERACT } P \S \mathcal{D}[\mathbf{domTag}(\vec{e}) \text{ t } U D] \longrightarrow P \S \mathcal{D}[\mathbf{domTag}(\vec{e} \cdot \text{ev}(V)) \text{ t } U D]$ <p style="text-align: center; margin-left: 100px;">for some ev, V such that $\vdash \text{ev}(V)$</p> $\text{E-EVT } P \S \mathcal{D}[\mathbf{domTag}(\text{ev}(W) \cdot \vec{e}) \text{ t } U D] \longrightarrow P \parallel ((V_1 W)) \parallel \dots \parallel ((V_n W)) \S \mathcal{D}[\mathbf{domTag}(\vec{e}) \text{ t } U D]$ <p style="text-align: center; margin-left: 100px;">where $\text{handlers}(\text{ev}, U) = \vec{V}$</p> $\text{E-STRUCT } \mathcal{C} \longrightarrow \mathcal{C}' \quad \text{if } \mathcal{C} \equiv \mathcal{C}_1, \mathcal{C}_1 \longrightarrow \mathcal{C}_2, \text{ and } \mathcal{C}_2 \equiv \mathcal{C}'$ $\text{E-LIFTP } P \S D \longrightarrow P' \S D \quad \text{if } P \longrightarrow P'$	<div style="border: 1px solid black; padding: 2px; display: inline-block;">$\mathcal{C} \longrightarrow \mathcal{C}'$</div>

■ **Figure 4** Reduction rules for λ_{MVU} terms and configurations

Concurrency is also essential when considering session-typed communication. We therefore formulate the calculus as a concurrent λ -calculus in the style of Niehren et al. [36], by augmenting the simply-typed λ -calculus with processes and concurrent reduction.

Processes. An *initialisation process* $\mathbf{run } M$ evaluates the initial system state written by a user, where M is a 3-tuple containing the initial model, view function, and update function. An *event loop process* $\langle T \mid F \rangle$ consists of an active thread T and function state F comprising the view and update functions. The thread can either be $\mathbf{idle } V_m$, meaning the process has current model V_m and is waiting for another message to process, or evaluating a term M . An *event handler process* $((M))$ is spawned to generate a message in response to an event.

Configurations. Concurrent and event-driven reduction happens in the context of a *system configuration* $P \S D$, where P is the concurrent fragment of the system and D is the current DOM page. An MVU program as written by a user is a term M specifying the initial model, view function, and update function, of type $(A \times (A \rightarrow \text{Html}(B)) \times ((B \times A) \rightarrow A))$. A program is evaluated in the context of an *initial configuration*:

► **Definition 1** (Initial configuration). *An initial configuration for a term M is of the form $\mathbf{run } M \S \mathbf{domEmpty}$.*

Evaluation contexts. Term evaluation contexts E (omitted) are set up for call-by-value, left-to-right evaluation. Process contexts \mathcal{P} allow reduction under parallel composition. Thread contexts \mathcal{T} allow reduction inside threads. DOM contexts \mathcal{D} allow us to focus on each element of a DOM forest; note that they deliberately allow non-unique decomposition in order to support nondeterministic reduction.

2.2.2 Reduction Rules

Figure 4 shows the reduction rules for λ_{MVU} processes and configurations; reduction on terms is standard β -reduction. Reduction on configurations is defined modulo the associativity and commutativity of parallel composition.

Diffing. As DOM pages include event queues, they contain strictly more information than HTML. To avoid losing pending events, we require a diffing operation. Define $\text{erase}(D)$ as the operation $\text{erase}(\mathbf{domTag}(\vec{e}) \text{ t } U \ D) = \mathbf{htmlTag} \text{ t } U (\text{erase}(D))$, with the other cases defined recursively. DOM pages can be modified by adding a node with an empty queue, removing a node, or updating a node's attributes. We define operation $\text{diff}(U, D) = D'$ if $\text{erase}(D') = U$, and D' is obtained from D by the minimum number of insertions and deletions.

Semantics by example. Let us return to our original example from §1: a box and a text node displaying the reversed box contents. We reuse the `view` and `update` functions and let $V_m = (\text{contents} = \text{""})$, $V_v = \text{view}$, and $V_u = \text{update}$. We extend the HTML syntactic sugar to pages, letting $\llbracket - \rrbracket$ be a desugaring function and $\llbracket \langle \text{t } \vec{a} \ @ \ \vec{e} \ > \vec{D}_H \langle / \text{t} \rangle \rrbracket = \mathbf{domTag}(\vec{e}) \text{ t } \llbracket \vec{a} \rrbracket \llbracket \vec{D}_H \rrbracket$.

We write \mathcal{R}^+ for the transitive closure of a relation \mathcal{R} . We begin by supplying the model, view, and update parameters to an initial configuration. By E-RUN, we get an event loop process, and then term $V_v \ V_m$ reduces to the initial rendered HTML. By diffing against the empty page, we display the initial DOM page (E-UPDATE).

$$\begin{aligned}
 & \text{run } (V_m, V_v, V_u) \ ; \ \mathbf{domEmpty} \\
 \rightarrow & \text{(E-RUN)} \ \langle (V_m, V_v \ V_m) \mid (V_v, V_u) \rangle \ ; \ \mathbf{domEmpty} \\
 \rightarrow & \text{M}^+ \ \langle (V_m, \ \langle \text{input type = "text" value = ""} \\
 & \quad \text{onInput = } \{\lambda str. \text{UpdateBox}(str)\} \rangle \langle / \text{input} \rangle \ \rangle \mid (V_v, V_u) \rangle \ ; \ \mathbf{domEmpty} \\
 & \ \langle \text{div} \rangle \langle / \text{div} \rangle \\
 \rightarrow & \text{(E-UPDATE)} \ \langle \text{idle } V_m \mid (V_v, V_u) \rangle \ ; \ \langle \text{input type = "text" value = ""} \\
 & \quad \text{onInput = } \{\lambda str. \text{UpdateBox}(str)\} \ @ \ \epsilon \rangle \langle / \text{input} \rangle \\
 & \ \langle \text{div } @ \ \epsilon \rangle \langle / \text{div} \rangle
 \end{aligned}$$

The system now does not reduce until a user interacts with the text box and presses the k key, modelled by E-INTERACT. At this point, the event queue for the input box receives four events: `click`, `keyDown`, `keyUp`, and `input`, which are processed by rule E-EVT. The input element does not have handlers for the `click`, `keyDown`, and `keyUp` events, so no processes are spawned, but *does* contain an `onInput` handler, which handles the `input` event by spawning $((\text{UpdateBox}(\text{"k"})))$.

$$\begin{aligned}
 \rightarrow & \text{+ (E-INTERACT)} \ \langle \text{idle } V_m \mid (V_v, V_u) \rangle \ ; \ \langle \text{input type = "text" value = ""} \\
 & \quad \text{onInput = } \{\lambda str. \text{UpdateBox}(str)\} \ @ \ \text{click}(\ ()) \cdot \\
 & \quad \text{keyDown}(75) \cdot \text{keyUp}(75) \cdot \text{input}(\text{"k"}) \rangle \langle / \text{input} \rangle \\
 & \ \langle \text{div } @ \ \epsilon \rangle \langle / \text{div} \rangle \\
 \rightarrow & \text{+ (E-EVT)} \ \langle \text{idle } V_m \mid (V_v, V_u) \rangle \ ; \ \langle \text{input type = "text" value = ""} \\
 & \quad \text{onInput = } \{\lambda str. \text{UpdateBox}(str)\} \ @ \ \text{input}(\text{"k"}) \rangle \\
 & \ \langle / \text{input} \rangle \\
 & \ \langle \text{div } @ \ \epsilon \rangle \langle / \text{div} \rangle \\
 \rightarrow & \text{(E-EVT)} \ \langle \text{idle } V_m \mid (V_v, V_u) \rangle \ \parallel \ ((\text{UpdateBox}(\text{"k"}))) \ ; \ \langle \text{input type = "text" value = ""} \\
 & \quad \text{onInput = } \{\lambda str. \text{UpdateBox}(str)\} \\
 & \quad \ @ \ \epsilon \rangle \langle / \text{input} \rangle \\
 & \ \langle \text{div } @ \ \epsilon \rangle \langle / \text{div} \rangle
 \end{aligned}$$

Since $\text{UpdateBox}(\text{"k"})$ is already a value and the event loop process is idle, we can process the message (E-HANDLE). The `handle` meta-function calculates a new model m' by applying

Typing rules for events $\vdash e$	Typing rules for active threads $\vdash T : \text{EvtLoop}(A, B)$
$\frac{\text{TE-EVT}}{\cdot \vdash V : \text{ty}(\text{ev})} \quad \vdash \text{ev}(V)$	$\frac{\text{TS-IDLE}}{\cdot \vdash V_m : A} \quad \vdash \text{idle } V_m : \text{EvtLoop}(A, B)$
$\frac{\text{TS-PROCESSING}}{\cdot \vdash M : (A \times \text{Html}(B))} \quad \vdash M : \text{EvtLoop}(A, B)$	
Typing rules for processes and configurations $\vdash^\phi P : A$ $\vdash C$	
$\frac{\text{TP-RUN}}{\cdot \vdash M : (A \times (A \rightarrow \text{Html}(B)) \times ((B \times A) \rightarrow A))} \quad \vdash^\bullet \text{run } M : B$	$\frac{\text{TP-EVENTLOOP}}{\cdot \vdash V_v : A \rightarrow \text{Html}(B) \quad \cdot \vdash V_u : (B \times A) \rightarrow A} \quad \vdash^\bullet \langle T \mid (V_v, V_u) \rangle : B$
$\frac{\text{TP-THREAD}}{\cdot \vdash M : A} \quad \vdash^\circ ((M)) : A$	$\frac{\text{TP-PAR}}{\vdash^{\phi_1} P_1 : A \quad \vdash^{\phi_2} P_2 : A} \quad \vdash^{\phi_1 + \phi_2} P_1 \parallel P_2 : A$
$\frac{\text{TC-SYSTEM}}{\vdash^\bullet P : A \quad \vdash D : \text{Page}(A)} \quad \vdash P \S D$	
Combination of flags $\phi_1 + \phi_2$	
$\circ + \circ = \circ \quad \circ + \bullet = \bullet \quad \circ + \bullet = \bullet \quad \bullet + \bullet \text{ undefined}$	

■ **Figure 5** Runtime typing for λ_{MVU}

the `update` function to a pair of the previous model and the message, calculates a new HTML value v' by applying the view function to m' , and returns the pair (m', v') . Finally, the page is diffed against the previous DOM page to generate a new DOM page D' , and the event loop process reverts to being idle:

$$\begin{aligned} &\longrightarrow (\text{EP-HANDLE}) && \langle \text{handle}(V_m, (V_v, V_u), \text{UpdateBox}("k")) \mid (V_v, V_u) \rangle \S && \begin{aligned} &\langle \text{input type} = \text{"text"} \text{ value} = \text{""} \\ &\quad \text{onInput} = \{\lambda \text{str}. \text{UpdateBox}(\text{str})\} \\ &\quad @ \epsilon > \langle \text{/input} \rangle \\ &\langle \text{div } @ \epsilon > \langle \text{/div} \rangle \end{aligned} \\ &\longrightarrow_M^+ && \begin{aligned} &(\text{contents} = \text{"k"}), \\ &\langle \langle \text{input type} = \text{"text"} \text{ value} = \text{"k"} \\ &\quad \text{onInput} = \{\lambda \text{str}. \text{UpdateBox}(\text{str})\} \rangle \rangle \mid (V_v, V_u) \rangle \S \\ &\quad \langle \text{/input} \rangle \\ &\quad \langle \text{div} > k \langle \text{/div} \rangle \end{aligned} && \begin{aligned} &\langle \text{input type} = \text{"text"} \text{ value} = \text{""} \\ &\quad \text{onInput} = \{\lambda \text{str}. \text{UpdateBox}(\text{str})\} \\ &\quad @ \epsilon > \langle \text{/input} \rangle \\ &\quad \langle \text{div } @ \epsilon > \langle \text{/div} \rangle \end{aligned} \\ &\longrightarrow (\text{E-UPDATE}) && \langle \text{idle}(\text{contents} = \text{"k"}) \mid (V_v, V_u) \rangle \S && \begin{aligned} &\langle \text{input type} = \text{"text"} \text{ value} = \text{"k"} \\ &\quad \text{onInput} = \{\lambda \text{str}. \text{UpdateBox}(\text{str})\} @ \epsilon > \langle \text{/input} \rangle \\ &\quad \langle \text{div } @ \epsilon > k \langle \text{/div} \rangle \end{aligned} \end{aligned}$$

2.3 Metatheory

Runtime typing. To reason about the metatheory, we require runtime typing rules, shown in Figure 5. Judgement $\vdash e$ states that the payload of an event e has the payload type specified by its signature. Judgement $\vdash T : \text{EvtLoop}(A, B)$ can be read “Active thread T has model type A and message type B ”. An idle thread `idle` V_m has type $\text{EvtLoop}(A, B)$ if V_m has type A (TS-IDLE). An active thread M currently processing a message has type $\text{EvtLoop}(A, B)$ if M has type $(A \times \text{Html}(B))$, i.e., computes a pair of a new model with type A and HTML which produces messages of type B (TS-PROCESSING).

Judgement $\vdash^\phi P : A$ states that process P is well typed and produces or consumes messages of type A . The parallel composition of two processes $P_1 \parallel P_2$ has message type A if both P_1 and P_2 have message type A (TP-PAR). An event handler process $((M))$ has message type A if term M has type A (TP-THREAD).

An initialisation process $\text{run } M$ is well-typed if M is a product type where each component has the correct model, view, and update types. An event loop process $\langle T \mid (V_v, V_u) \rangle$ has message type B if its active thread T has model type A and message type B ; its view function V_v has type $A \rightarrow \text{Html}(B)$; and its update function has type $(B \times A) \rightarrow A$ (TP-EVENTLOOP). Thread flags ϕ ensure that there is precisely one initialisation process or event loop process in a process typeable under flag \bullet .

Judgement $\vdash \mathcal{C}$ states that configuration \mathcal{C} is well-typed: a system configuration $P \wp D$ is well-typed if process P has precisely one event loop process with message type A and page D has type $\text{Page}(A)$. The omitted typing rules for pages (of shape $\vdash D : \text{Page}(A)$) are similar to those for terms of type $\text{Html}(A)$.

Note that we consider only closed configurations and processes since it makes little sense for DOM pages D to contain free variables, and because processes do not bind variables.

We are now well-placed to state some formal results. We omit proofs in the main body of the paper; full proofs can be found in the extended version [18].

Preservation. Reduction preserves typing.

► **Theorem 2** (Preservation). *If $\vdash \mathcal{C}$ and $\mathcal{C} \longrightarrow \mathcal{C}'$, then $\vdash \mathcal{C}'$.*

Progress. The system vacuously satisfies a progress property as it can always reduce by E-INTERACT due to user input. It is more interesting to consider the *event progress* property enjoyed by the system *without* E-INTERACT: either there are no events to process and the system is idle, or the system can reduce. Functional reduction satisfies progress.

► **Lemma 3** (Progress (Terms)). *If $\cdot \vdash M : A$, then either M is a value, or there exists some N such that $M \longrightarrow_M N$.*

Let \longrightarrow_E be the relation \longrightarrow without rule E-INTERACT. The concurrent fragment of the language will reduce until all event handler threads have finished evaluating, and there are no more messages to process. By appeal to Lemma 3, we can show event progress.

► **Theorem 4** (Event Progress). *If $\vdash \mathcal{C}$, either:*

1. *there exists some \mathcal{C}' such that $\mathcal{C} \longrightarrow_E \mathcal{C}'$; or*
2. *$\mathcal{C} = \langle \text{idle } V_m \mid (V_v, V_u) \rangle \wp D$ where D cannot be written $\mathcal{D}[\text{dom Tag}(\vec{e}) \ t \ V \ W]$ for some non-empty \vec{e} .*

3 λ_{MVU} with Session Types

In this section, we extend λ_{MVU} to support session types. We require three extensions: *commands*, to perform side-effects; *linearity*, to implement session types safely; and *transitions*, to allow multiple model and message types. We begin by showing each extension by example, and show the extended formalism in §3.4.

3.1 Commands

Real-world applications require side-effects. To this end, Elm supports *commands* which describe side-effects to be performed in the event loop. Although commands in Elm are more general, for our purposes, it is particularly useful to be able to spawn a process which will run concurrently and eventually return a message. As an example, we may want to await the result of an expensive computation, and display the result when the computation completes.

14:12 Model-View-Update-Communicate

Letting $\text{naiveFib}(x)$ be the naïve Fibonacci function and assuming an intToString function, we can write:

```

Model  $\triangleq$  Maybe(Int)    Message  $\triangleq$  StartComputation | Result(Int)
view : Model  $\rightarrow$  Html(Message)
view =  $\lambda model.$ html
  <html>
  <body>
    {case model {
      Just(result)  $\mapsto$  htmlText intToString(result);
      Nothing  $\mapsto$  htmlText "waiting ..."} }
    <button onClick = { $\lambda().$ StartComputation}>Start!</button>
  </body>
</html>

update : (Message  $\times$  Model)  $\rightarrow$  (Model, Cmd(Message))
update =  $\lambda(message, model).$ 
  case message {
    StartComputation  $\mapsto$  (Nothing, cmdSpawn Result(naiveFib(1000)))
    Result(x)  $\mapsto$  (Just(x), cmdEmpty)
  }

```

The model is of type $\text{Maybe}(\text{Int})$, with value $\text{Just}(V)$ for some integer value V if the result has been computed, or Nothing if the application is awaiting the result. The Message type is a variant type consisting of StartComputation which is sent to start the computation, and $\text{Result}(\text{Int})$, which is sent to return a result. The view function renders either the result, or “waiting...” if no result is available.

The type of the `update` function is changed to return a *pair* of an updated model and a command. In our case, the StartComputation message results in a pair of Nothing and $\text{cmdSpawn Result(naiveFib(1000))}$, which spawns $\text{Result(naiveFib(1000))}$ to evaluate in a separate thread. When the function (eventually) completes, the thread will have evaluated to a Result message, which can be processed by the `update` function to update the model and display the result.

3.2 Linearity

As we showed in §1, safely implementing session types requires linearity: we therefore require linear model and message types. Linearity would also prove useful for other linear resources such as functional arrays with in-place update [44]. Unfortunately, λ_{Mvu} as defined so far does not support linearity. Consider `handle`:

$$\text{handle}(m, (v, u), \text{msg}) \triangleq \text{let } m' = u(\text{msg}, m) \text{ in } (m', v \text{ } m')$$

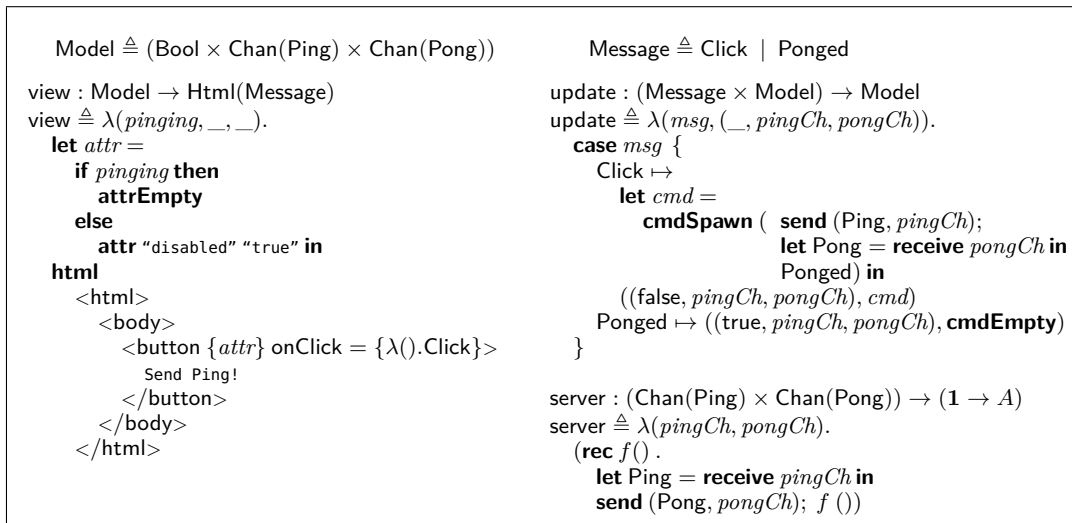
The updated model, m' , is used non-linearly as it is returned for use in subsequent requests, and also used to render the model as HTML.

Extraction. Linear resources are needed only when *updating* the model—not when rendering the webpage—as we will not need to communicate on session channels when rendering. If the developer implements a function:

$$\text{extract} : A \rightarrow (A \times B)$$

where A is the type of a model, and B is the *unrestricted* fragment of the model, we can restore linear usage of the model (letting e be the extraction function):

$$\text{handle}(m, (v, u, e), \text{msg}) \triangleq \text{let } m' = u(\text{msg}, m) \text{ in } \text{let } (m', \text{unrM}) = e \text{ } m' \text{ in } (m', v \text{ } \text{unrM})$$



■ **Figure 6** PingPong application using simply-typed channels

An alternative approach would be to assign the view function type $A \rightarrow (A \times \text{Html}(B))$, returning the linear model and allowing it to be re-bound. We would need to modify handle:

$$\text{handle}(m, (v, u), \text{msg}) \triangleq \text{let } m' = u(m, \text{msg}) \text{ in } v m'$$

A key disadvantage of this approach is that rendering is no longer a read-only operation, breaking an important abstraction barrier.

Example. We can now write our first session-typed λ_{MVU} application. Our web client consists of a button which, when clicked, triggers the sending of a Ping message to the server. Once clicked, the button is disabled. The server then receives the Ping message and responds with a Pong message; upon receiving the response, the client then re-enables the button.



Simply-typed channels. Before considering a session-typed version of the application, it is instructive to consider a version *without* session typing, shown in Figure 6. Let $\text{Chan}(A)$ be the type of a simply-typed channel over which one can send and receive values of type A . The model is a 3-tuple containing a Boolean value which is true when waiting for the user to click the “Send Ping!” button, and false when waiting for a response; a channel for Ping messages; and a channel for Pong messages. There are two types of UI message: Click denotes that the button has been clicked, and Ponged denotes that a Pong message has been received along the Pong channel.

The view function displays the page, adding the `disabled` attribute to the button if we are waiting for a Pong message. The `update` function case-splits on the UI message: in the case of a Click message raised by the button, the model is updated to set the `pinging` flag to `false`, and the function creates a command to send a Ping message along `pingCh`, receive a Pong message from `pongCh`, and return a Ponged UI message. In the case of a Ponged message, the model is updated to set the `pinging` flag to `true`, enabling the button again. The server function models a server thread, which repeatedly receives Ping messages from `pingCh` and sends Pong messages to `pongCh`.

<pre> PingPong \triangleq $\mu t. !\text{Ping}. ?\text{Pong}. t$ UModel \triangleq UPinging UWaiting view : UModel \rightarrow Html(Message) view \triangleq $\lambda uModel.$ let attr = case uModel { UPinging \mapsto attrEmpty UWaiting \mapsto attr "disabled" "true" } in html <html> <body> <button {attr} onClick = {$\lambda()$.Click}> Send Ping! </button> </body> </html> handleClick(model) \triangleq case model { Pinging(c) \mapsto let cmd = cmdSpawn (let $c = \text{send}$ (Ping, c) in let ($pong, c$) = receive c in Ponged(c) in (Waiting, cmd) Waiting \mapsto (Waiting, cmdEmpty) } </pre>	<pre> Model \triangleq Pinging(PingPong) Waiting Message \triangleq Click Ponged(PingPong) update : (Message \times Model) \rightarrow (Model \times Cmd(Message)) update \triangleq $\lambda(msg, model).$ case msg { Click \mapsto handleClick(model) Ponged(c) \mapsto handlePonged(model, c) } extract : Model \rightarrow (Model \times UModel) extract \triangleq $\lambda model.$ case model { Pinging(c) \mapsto (Pinging(c), UPinging) Waiting \mapsto (Waiting, UWaiting) } handlePonged(model, c) \triangleq case model { Pinging(c') \mapsto cancel c'; (Pinging(c), cmdEmpty) Waiting \mapsto (Pinging(c), cmdEmpty) } </pre>
---	--

■ Figure 7 PingPong application

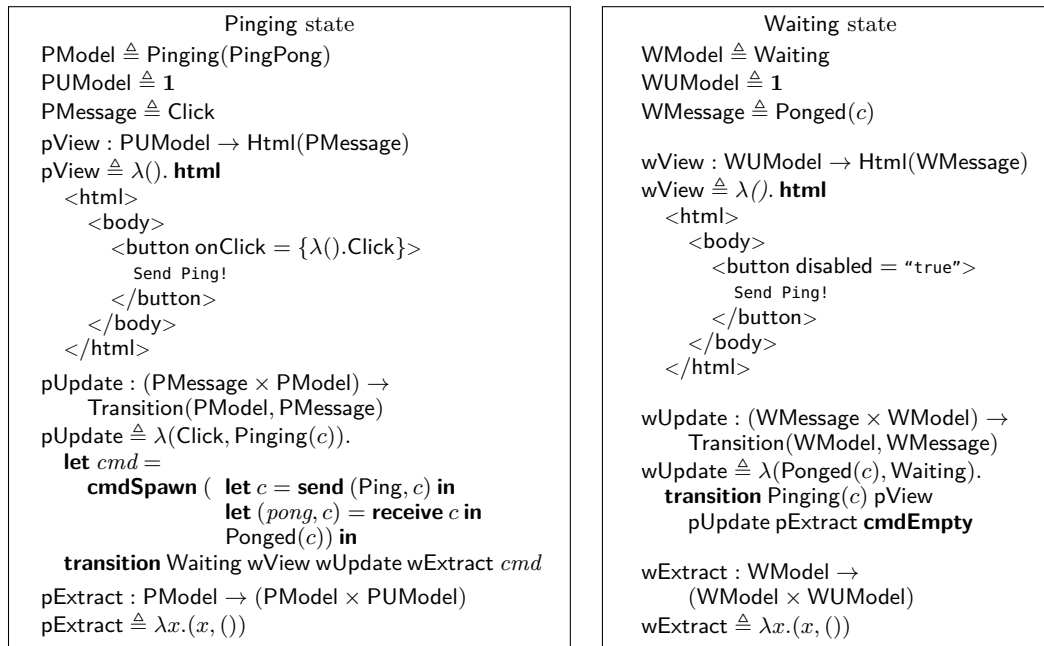
Even in this simple example, it is very easy to communicate incorrectly: if the client neglected to send a **Ping** message before trying to receiving a **Pong** message along $pongCh$, then the command would hang forever and the GUI would never re-enable the button. A similar situation would arise if the server received the **Ping** message but failed to respond.

Session types. Session types S range over output $!A.S$, input $?A.S$, the completed session **End**, recursive session types $\mu t.S$, and (possibly dualised) recursive type variables t . We take an equi-recursive treatment of recursive session types, identifying a recursive session type with its unfolding. We omit types and constructs for branching and selection as they can be encoded [13, 28]. The **send** constant sends a value of type A over an endpoint of type $!A.S$ and returns the continuation of the session, S . The **close** constant closes a completed session endpoint. The **receive** constant takes an endpoint of type $?A.S$ and receives a pair of a value of type A and endpoint of type S . The **cancel** constant allows an endpoint to be discarded safely [20, 35].

$$\text{Session types } S ::= !A.S \mid ?A.S \mid \mu t.S \mid t \mid \bar{t} \mid \text{End}$$

$$\text{send} : (A \times !A.S) \rightarrow S \quad \text{receive} : ?A.S \rightarrow (A \times S) \quad \text{close} : \text{End} \rightarrow \mathbf{1} \quad \text{cancel} : S \rightarrow \mathbf{1}$$

Figure 7 shows the PingPong client written in λ_{MVU} . We can encode the PingPong protocol as a session type, $\mu t. !\text{Ping}. ?\text{Pong}. t$. The Model type encodes the two states of the application: **Pinging**(c) is the state where the “Send Ping!” button is enabled and the user can send a **Ping** message along session channel c , whereas **Waiting** is the state where the button is disabled and awaiting a **Pong** message from the other party. The UModel type is the unrestricted model type which does not include the session channel. Again, the Message



■ **Figure 8** PingPong application using model transitions

type encodes the UI messages in the application: the `Click` UI message is produced when the button is pressed, whereas the `Ponged(PingPong)` UI message is produced when a `Pong` session message has been received. Note that the `Ponged` UI message now contains a session channel of type `PingPong` as a parameter.

The view function takes an unrestricted model and displays a button, which is disabled in the `Waiting` state but enabled in the `Pinging` state. The extract function pairs the linear model with the associated unrestricted model.

The update function case-splits on the message. The `handleClick` function handles the `Click` message, and case-splits on the model. If the model is in the `Pinging(c)` state, then the function creates a command to spawn a process which will send a `Ping` message along `c`, receive a `Pong` message along `c`, and generate a `Ponged` UI message when the `Pong` message is received. The function finally updates the model to the `Waiting` state. If the model is in the `Waiting` state—which should not occur, since the button is disabled—then the model remains the same and no command is created.

The `handlePonged` function handles a `Ponged(c)` message. Again, we must case split on the model. If the model is in the `Waiting` state, then we can change to the `Pinging` state, given endpoint `c`. However, if the model is in the `Pinging(c')` state and a `Ponged` message is received—which should not occur, since according to the session type, there is no way for the peer to send a `Pong` message while we are waiting to send a `Ping`—we now have *two* linear resources. We choose to discard `c'` using **cancel**, and change the model to `Pinging(c')`, but this is an arbitrary choice to satisfy a code path that must exist, but should never be used.

3.3 Model transitions

Our proposal is still not quite satisfactory: as we saw with the `PingPong` example, we need to include cases in the `update` function which cannot arise. We highlight these in red. This is even more pronounced when dealing with linear resources, such as needing to handle a `Ponged` message when waiting to send a `Ping`.

Kinds	$\kappa ::= \mathbf{L} \mid \mathbf{U}$
Types	$A, B, C ::= \mathbf{1} \mid A \rightarrow^\kappa B \mid A \times B \mid A + B \mid \mathbf{String} \mid \mathbf{Int} \mid S$ $\mid \mathbf{Html}(A) \mid \mathbf{Attr}(A) \mid \mathbf{Cmd}(A) \mid \mathbf{Transition}(A, B)$
Session types	$S ::= !A.S \mid ?A.S \mid \mu t.S \mid t \mid \bar{t} \mid \mathbf{End}$
Terms	$L, M, N ::= x \mid \lambda x.M \mid M N \mid K M \mid () \mid s \mid n$ $\mid (M, N) \mid \mathbf{let} (x, y) = M \mathbf{in} N$ $\mid \mathbf{inl} x \mid \mathbf{inr} x \mid \mathbf{case} L \{ \mathbf{inl} x \mapsto M; \mathbf{inr} y \mapsto N \}$ $\mid \mathbf{htmlTag} \ t \ M \ N \mid \mathbf{htmlText} \ M$ $\mid \mathbf{attr} \ ak \ M \mid \mathbf{attrEmpty}$ $\mid \mathbf{cmdSpawn} \ M \mid \mathbf{cmdEmpty} \mid M \star N$ $\mid \mathbf{transition} \ M_m \ M_v \ M_u \ M_e \ M_c \mid \mathbf{noTransition} \ M_m \ M_c$ $\mid \mathbf{raise} \mid \mathbf{try} \ L \ \mathbf{as} \ x \ \mathbf{in} \ M \ \mathbf{otherwise} \ N$
Constants	$K ::= \mathbf{send} \mid \mathbf{receive} \mid \mathbf{new} \mid \mathbf{cancel} \mid \mathbf{close}$

■ **Figure 9** Syntax of extended calculus

The problem is that we are encoding the Model type using a sum type, whereas in fact we require *multiple* model types, and a way to *transition* between them.

Example. Figure 8 shows how we can modify PingPong to use multiple model types. The left-hand side of the figure shows the Pinging state: the model type consists of the singleton variant tag Pinging(PingPong) containing an endpoint of type PingPong, the unrestricted model is the unit type, and the only message that the Pinging state can receive is Click. The pView function is similar to before, and the pExtract function returns a pair of the current state and the unit value. The pUpdate function is more interesting. Given the current state and a Click message, the function constructs a command which will send the Ping session message, receive the Pong session message, and then generate a Ponged(*c*) UI message containing the session channel. The function *transitions* into the Waiting state using the **transition** primitive, which allows the developer to specify new model, view, update, extract functions, and a command to evaluate. The functions for the Waiting state follow a similar pattern. Session types rule out the communication errors besetting the example in Figure 6, and model transitions eliminate the redundant code paths arising due to illegal states.

3.4 λ_{MVU} with Commands, Linearity, and Transitions

Commands, linearity, and transitions are the three key ingredients needed to extend MVU to support models which include session-typed channels. In this section, we introduce a calculus which combines all three extensions, and prove that the extended calculus is sound.

3.4.1 Syntax and Typing

Figure 9 shows the syntax of λ_{MVU} extended with commands, linearity, and transitions.

Types and kinds. To support linearity, types are assigned *kinds*, ranged over by κ . Types can either be *linear* (L) or *unrestricted* (U). A value of linear type must be used precisely once, whereas a value of unrestricted type can be used many times.

We modify function types to include a kind annotation: linear functions may close over linear variables and so must be used once. To support commands, we introduce type $\mathbf{Cmd}(A)$ which is the type of a command which produces messages of type A . To support transitions, we introduce type $\mathbf{Transition}(A, B)$ which is parameterised by the *current* model type A and message type B . Finally, we extend types to include session types S as described in §3.2.

Terms. Term **cmdSpawn** M is a command which can spawn term M as a thread, and is monoidally composable using \star and **cmdEmpty**.

There are two terms for transitions: the **noTransition** $M_m M_c$ term denotes that no transition is to occur, and that the model should be updated to M_m and command M_c should be evaluated; and **transition** $M_m M_v M_u M_e M_c$ denotes that a transition should occur, with new model M_m , view function M_v , update function M_u , extraction function M_e , and command M_c to be run once the transition has taken place.

To support session typing, we introduce session typing constants, ranged over by K , as described in §3.2. We also introduce an application form for constants, $K M$.

Finally, as discussed in §3.2, it is useful to be able to explicitly discard (or *cancel*) a session channel. In particular, cancellation is crucial in order to handle the interplay between linearity and transitions, as all unprocessed messages (which may contain linear resources) must be safely discarded when a transition occurs.

Following Mostrous and Vasconcelos [35] and Exceptional GV (EGV) by Fowler et al. [20], if a thread tries to receive from an endpoint whose peer has been cancelled, an exception is raised (**raise**). Exceptions can be handled using the **try** L **as** x **in** M **otherwise** N construct, which tries to evaluate term L , and binds the result to x in M if the term evaluates to a value, and evaluates N if the term raises an exception.

Kinding and subkinding. The kinding relation $A :: \kappa$ assigns kind κ to type A ; our formulation is inspired by that of Padovani [38]. Base types and HTML and attribute types are unrestricted. The kind of a function type is determined by its kind annotation. Session types are linear. The kinds of product, sum, command and transition types are determined by the kinds of their type parameters. The reflexive *subkinding* rule $\mathsf{U} \leq \mathsf{L}$ combined with the kinding subsumption rule states that if a value can be used many times, then it can also be treated as only being used once. We write $\Gamma :: \kappa$ if $A :: \kappa$ for each $x : A \in \Gamma$.

► **Definition 5** (Kinding and subkinding). *We define the subkinding relation as the reflexive relation on kinds \leq such that $\mathsf{U} \leq \mathsf{L}$. We define the kinding relation $A :: \kappa$ as the largest relation between types and kinds such that:*

- $A :: \kappa'$ if $A :: \kappa$ and $\kappa \leq \kappa'$
- $S :: \mathsf{L}$
- $A :: \mathsf{U}$ if $A \in \{\mathbf{1}, \mathbf{String}, \mathbf{Int}, \mathbf{Html}(B), \mathbf{Attr}(B)\}$
- $A \rightarrow^\kappa B :: \kappa$
- $\mathbf{Cmd}(A) :: \kappa$ if $A :: \kappa$
- $C :: \kappa$ if $C \in \{A \times B, A + B, \mathbf{Transition}(A, B)\}$ and both $A :: \kappa$ and $B :: \kappa$

Term typing. Figure 10 shows the typing rules for the extended calculus. The splitting relation $\Gamma = \Gamma_1 + \Gamma_2$ [8] splits a typing context Γ into two subcontexts which may share only unrestricted variables. We support linearity by changing T-VAR to only type a variable in an unrestricted context, and by using the context splitting judgement when typing subterms. The adaptation of the remaining rules to use context splitting is standard, so we omit them.

The constant application rule T-APPK types term $K M$ and makes use of the type schema function $\Sigma(K)$ to ensure that the argument M is of the correct type. Rule T-CMDSPAWN assigns term **cmdSpawn** M type $\mathbf{Cmd}(A)$ if term M has type A , and rules T-CMDEEMPTY and T-CMDAPPEND allow commands to be composed monoidally.

Rule T-TRANSITION types a **transition** term. The typing rule ensures that the types of the new model, and view, update and extract functions are compatible. Note that the type parameters of the $\mathbf{Transition}(A', B')$ need not match the types of the new model and

Context splitting		$\Gamma = \Gamma_1 + \Gamma_2$
$\cdot = \cdot + \cdot$	$\frac{A :: \mathbf{U}}{\Gamma, x : A = (\Gamma_1, x : A) + (\Gamma_2, x : A)}$	$\frac{}{\Gamma_1 + \Gamma_2, x : A = (\Gamma_1, x : A) + \Gamma_2}$
Modified typing rules for terms		$\Gamma \vdash M : A$
$\frac{}{\Gamma, x : A \vdash x : A}$	$\frac{}{\Gamma \vdash \lambda x. M : A \rightarrow^\kappa B}$	$\frac{}{\Gamma \vdash K M : B}$
$\frac{}{\Gamma \vdash \mathbf{cmdSpawn} M : \mathbf{Cmd}(A)}$	$\frac{}{\Gamma \vdash \mathbf{cmdEmpty} : \mathbf{Cmd}(A)}$	$\frac{}{\Gamma_1 + \Gamma_2 \vdash M \star N : \mathbf{Cmd}(A)}$
$\frac{}{\Gamma_1 \vdash M_m : A \quad \Gamma_2 \vdash M_v : A \rightarrow^{\mathbf{U}} \mathbf{Html}(B) \quad \Gamma_3 \vdash M_u : (B \times A) \rightarrow^{\mathbf{U}} \mathbf{Transition}(A, B) \quad \Gamma_4 \vdash M_e : A \rightarrow^{\mathbf{U}} (A \times C) \quad \Gamma_5 \vdash M_c : \mathbf{Cmd}(A) \quad C :: \mathbf{U}}{\Gamma_1 + \dots + \Gamma_5 \vdash \mathbf{transition} M_m M_v M_u M_e M_c : \mathbf{Transition}(A', B')}$		
$\frac{}{\Gamma \vdash \mathbf{attr} h M : \mathbf{Attr}(A)}$	$\frac{}{\Gamma_1 + \Gamma_2 \vdash \mathbf{noTransition} M N : \mathbf{Transition}(A, B)}$	
$\frac{}{\Gamma_1 + \Gamma_2 \vdash \mathbf{try} L \mathbf{as} x \mathbf{in} M \mathbf{otherwise} N : B}$	$\frac{}{\Gamma \vdash \mathbf{raise} : A}$	(other rules modified to split contexts)
Typing of constants		$\Sigma(c)$
$\Sigma(\mathbf{send}) = (A \times !A.S) \rightarrow^{\mathbf{U}} S$	$\Sigma(\mathbf{receive}) = ?A.S \rightarrow^{\mathbf{U}} (A \times S)$	$\Sigma(\mathbf{new}) = \mathbf{1} \rightarrow^{\mathbf{U}} (S \times \bar{S})$
$\Sigma(\mathbf{cancel}) = S \rightarrow^{\mathbf{U}} \mathbf{1}$	$\Sigma(\mathbf{close}) = \mathbf{End} \rightarrow^{\mathbf{U}} \mathbf{1}$	\bar{S}
Duality		\bar{S}
$\bar{!A.S} = ?A.\bar{S} \quad \bar{?A.S} = !A.\bar{S}$		$\bar{\mu t. S} = \mu t. \bar{S}\{\bar{t}/t\} \quad \bar{\bar{t}} = t \quad \bar{\mathbf{End}} = \mathbf{End}$

■ **Figure 10** Term typing for extended calculus.

functions. Rule T-NOTRANSITION assigns term **noTransition** $M N$ type $\mathbf{Transition}(A, B)$ if new model M has type A , and N is a command of type $\mathbf{Cmd}(B)$. Note that in this way, the **noTransition** $M N$ term replaces the standard result of the `update` function.

Rule T-TRY types an exception handler: the continuations share a typing environment, but the success continuation is augmented with the a variable of the type of the possibly-failing continuation. Finally, **raise** can have any type as is it does not return (T-RAISE).

The type and kind system ensures that the kind of type A determines the kind of the typing environment needed to type a term of type A .

► **Lemma 6.** *If $\Gamma \vdash M : A$ and $A :: \kappa$, then $\Gamma :: \kappa$.*

Duality. The duality relation for session types is standard: output types are dual to input types; we use a self-dual **End** type; and we use the formulation of the duality of recursive session types advocated by Lindley and Morris [32].

3.4.2 Operational Semantics

Runtime syntax. Figure 11 shows the runtime syntax for the combined calculus. We introduce *runtime names* c, d which identify session channel endpoints.

Runtime syntax	
Runtime names	c, d
Values	$U, V, W ::= \dots \mid c \mid \mathbf{cmdSpawn} M \mid \mathbf{noTransition} V W$ $\mid \mathbf{transition} V_m V_v V_u V_e V_c$
Active thread	$T ::= \mathbf{idle} V_m \mid \mathbf{updating} M \mid \mathbf{extracting}[V_c] M$ $\mid \mathbf{extractingT}[F, V_c] M \mid \mathbf{rendering}[V_m, V_c] M$ $\mid \mathbf{transitioning}[V_m, F, V_c] M$
Versions	ι
Processes	$P, Q ::= \mathbf{run} M \mid \langle T \mid F \rangle_\iota \mid ((M))_\iota \mid P \parallel Q$ $\mid (\nu cd)P \mid [M] \mid \zeta c \mid \mathbf{halt}$
Function state	$F ::= (V_v, V_u, V_e)$
Configurations	$\mathcal{C} ::= P \S D$
Process contexts	$\mathcal{P} ::= [] \mid \mathcal{P} \parallel P \mid (\nu cd)\mathcal{P}$
Evaluation contexts	$E ::= \dots \mid K E \mid \mathbf{noTransition} E M \mid \mathbf{noTransition} V E$ $\mid \mathbf{transition} E M_v M_u M_e M_c \mid \dots \mid \mathbf{transition} V_m V_v V_u V_e E$ $\mid \mathbf{try} E \mathbf{as} x \mathbf{in} M \mathbf{otherwise} N$
Pure contexts	$E_P ::=$ (as E , but without exception handling frames)
Active thread contexts	$\mathcal{T}_A ::= \mathbf{updating} E \mid \mathbf{rendering}[V_m, V_c] E \mid \mathbf{extracting}[V_c] E$ $\mid \mathbf{extractingT}[V_c, F'] E \mid \mathbf{transitioning}[V_m, F', V_c] E$
Pure active thread contexts	$\mathcal{T}_P ::=$ (as \mathcal{T}_A , but for pure contexts)
Thread contexts	$\mathcal{T} ::= \mathbf{run} E \mid \langle \mathcal{T}_A \mid F \rangle_\iota \mid ((E))_\iota \mid [E]$
Active thread state machine	
<pre> graph LR idle([idle]) -- "(Model transition)" --> updating([updating]) updating -- "(No model transition)" --> extracting([extracting]) updating -- "(Model transition)" --> extractingT([extractingT]) extracting --> rendering([rendering]) extractingT --> transitioning([transitioning]) </pre>	

■ **Figure 11** Runtime syntax for extended calculus

The biggest departure is that we require a richer structure on active threads, which form a state machine based on whether a model transition occurs. The **idle** state is as before, and the **updating** state evaluates the update function. If there is no model transition, then the thread moves to the **extracting** state to extract the unrestricted model, and the **rendering** state to render the new HTML. If there is a model transition, then the thread moves to the **extractingT** state followed by the **transitioning** state to calculate the new HTML to be displayed after the transition. Each state records values which are required in later states: for example, the **rendering** $[V_m, V_c] M$ state records the new model V_m and the command to be executed upon updating the page V_c .

We introduce four new types of process. To model client-server communication, we introduce server processes $[M]$ to model a process M running on the server; the thread to spawn is given as an argument to **run**. As an example, we could write a Ponger server process for the PingPong example, which immediately responds with a Pong message:

let $(c, s) = \mathbf{new} ()$ in	ponger $(s) \triangleq \lambda().$
$(\mathbf{Pinging}(c), \mathbf{pView}, \mathbf{pUpdate}, \mathbf{pExtract},$	$(\mathbf{rec} f(s) .$
$\mathbf{cmdEmpty}, \mathbf{ponger}(s))$	$\mathbf{let} (\mathbf{Ping}, s) = \mathbf{receive} s$ in
	$\mathbf{let} s = \mathbf{send} (\mathbf{Pong}, s)$ in} f s)

A name restriction $(\nu cd)P$ binds runtime names c and d in process P , following the double-binder formulation due to Vasconcelos [43]. A *zapper thread* ζc denotes an endpoint c

<p>Additional term reduction rule $M \rightarrow_M N$</p> <p>E-TRY $\text{try } V \text{ as } x \text{ in } M \text{ otherwise } N \rightarrow_M M\{V/x\}$</p> <p>Equivalence of processes</p> <p>$(\nu cd)(\nu c' d')P \equiv (\nu c' d')(\nu cd)P$ $P \parallel ((\nu cd)Q) \equiv (\nu cd)(P \parallel Q)$ if $c, d \notin \text{fn}(P)$ $(\nu cd)P \equiv (\nu dc)P$</p> <p>$P_1 \parallel P_2 \equiv P_2 \parallel P_1$ $P_1 \parallel (P_2 \parallel P_3) \equiv (P_1 \parallel P_2) \parallel P_3$ $(\nu cd)(\zeta c \parallel \zeta d) \parallel P \equiv P$ $(\cdot) \parallel P \equiv P$</p> <p>Reduction of processes $P \rightarrow P'$</p>	<p>Additional meta-level definitions</p> <p>$\text{procs}(\text{cmdEmpty}) = \epsilon$ $\text{procs}(\text{cmdSpawn } M) = M$ $\text{procs}(V \star W) = \text{procs}(V) \cdot \text{procs}(W)$</p> <p>$P \equiv P'$</p>
<p>MVU reduction rules</p>	
<p>E-DISCARD $\langle T \mid F \rangle_\iota \parallel ((V))_{\iota'} \rightarrow \langle T \mid F \rangle_\iota \parallel \zeta V$ if $\iota \neq \iota'$</p> <p>E-DISCARDHALT $\text{halt} \parallel ((V))_\iota \rightarrow \text{halt} \parallel \zeta V$</p> <p>E-HANDLE $\langle \text{idle } V_m \mid (V_v, V_u, V_e) \rangle_\iota \parallel ((V))_\iota \rightarrow \langle \text{updating } V_u(V, V_m) \mid (V_v, V_u, V_e) \rangle_\iota$</p> <p>E-EXTRACT $\langle \text{updating (noTransition } V_m V_c) \mid F \rangle_\iota \rightarrow \langle \text{extracting}[V_c](V_e V_m) \mid F \rangle_\iota$ where $F = (V_v, V_u, V_e)$</p> <p>E-EXTRACTT $\langle \text{updating (transition } V_m V_v V_u V_e V_c) \mid F \rangle_\iota \rightarrow \langle \text{extractingT}[(V_v, V_u, V_e), V_c](V_e V_m) \mid F \rangle_\iota$</p> <p>E-RENDER $\langle \text{extracting}[V_c](V_m, V_{um}) \mid F \rangle_\iota \rightarrow \langle \text{rendering}[V_m, V_c](V_v V_{um}) \mid F \rangle_\iota$ where $F = (V_v, V_u, V_e)$</p> <p>E-RENDERT $\langle \text{extractingT}[F', V_c](V_m, V_{um}) \mid F \rangle_\iota \rightarrow \langle \text{transitioning}[V_m, F', V_c](V_v V_{um}) \mid F \rangle_\iota$ where $F' = (V_v, V_u, V_e)$</p>	
<p>Session reduction rules</p>	
<p>E-NEW $\mathcal{T}[\text{new}()] \rightarrow (\nu cd)(\mathcal{T}[(c, d)])$ where c, d fresh</p> <p>E-COMM $(\nu cd)(\mathcal{T}[\text{send}(V, c)] \parallel \mathcal{T}'[\text{receive } d]) \rightarrow (\nu cd)(\mathcal{T}[c] \parallel \mathcal{T}'[(V, d)])$</p> <p>E-CLOSE $(\nu cd)(\mathcal{T}[\text{close } c] \parallel \mathcal{T}'[\text{close } d]) \rightarrow \mathcal{T}[\cdot] \parallel \mathcal{T}'[\cdot]$</p> <p>E-CANCEL $\mathcal{T}[\text{cancel } c] \rightarrow \mathcal{T}[\cdot] \parallel \zeta c$</p> <p>E-SENDZAP $(\nu cd)(\mathcal{T}[\text{send}(V, c)] \parallel \zeta d) \rightarrow (\nu cd)(\mathcal{T}[\text{raise}] \parallel \zeta c \parallel \zeta V \parallel \zeta d)$</p> <p>E-RECVZAP $(\nu cd)(\mathcal{T}[\text{receive } c] \parallel \zeta d) \rightarrow (\nu cd)(\mathcal{T}[\text{raise}] \parallel \zeta c \parallel \zeta d)$</p> <p>E-CLOSEZAP $(\nu cd)(\mathcal{T}[\text{close } c] \parallel \zeta d) \rightarrow (\nu cd)(\mathcal{T}[\text{raise}] \parallel \zeta c \parallel \zeta d)$</p>	
<p>Exception reduction rules</p>	
<p>E-RAISEH $\mathcal{T}[\text{try } E_P[\text{raise}] \text{ as } x \text{ in } M \text{ otherwise } N] \rightarrow \mathcal{T}[N] \parallel \zeta E_P$</p> <p>E-RAISEURUN $\text{run}(E_P[\text{raise}]) \rightarrow \text{halt} \parallel \zeta E_P$</p> <p>E-RAISEUMAIN $\langle \mathcal{T}_P[\text{raise}] \mid F \rangle_\iota \rightarrow \text{halt} \parallel \zeta \mathcal{T}_P$</p> <p>E-RAISEUTHREAD $((E_P[\text{raise}]))_\iota \rightarrow \zeta E_P$</p> <p>E-RAISEUSERVER $[E_P[\text{raise}]] \rightarrow \zeta E_P$</p>	
<p>Administrative reduction rules</p>	
<p>E-LIFTT $\mathcal{T}[M] \rightarrow \mathcal{T}[N]$ if $M \rightarrow_M N$</p> <p>E-NU $(\nu ab)P \rightarrow (\nu ab)P'$ if $P \rightarrow P'$</p> <p>E-PAR $P_1 \parallel P_2 \rightarrow P'_1 \parallel P_2$ if $P_1 \rightarrow P'_1$</p>	

■ **Figure 12** Reduction rules for extended calculus (1)

that has been cancelled and cannot be used in future communications; we write ζV to mean $\zeta c_1 \parallel \dots \parallel \zeta c_n$ for $c_i \in \text{fn}(V)$, where $\text{fn}(V)$ enumerates the free runtime names in a value V , and extend this sugar to evaluation contexts. The **halt** process denotes that the event loop process has terminated due to an unhandled exception.

We extend evaluation contexts in the standard way, and introduce a class of *pure contexts* E_P , which are evaluation contexts which do not contain any exception handling frames.

Versions. *Versions* ι ensure that threads spawned in a previous state do not deliver incompatible messages. We annotate event loop processes and event handler threads with versions: given an event loop $\langle T \mid F \rangle_\iota$, a thread $((M))_{\iota'}$ where $\iota \neq \iota'$ can be of arbitrary type as it will be discarded. We write $\text{version}(P) = \iota$ if P contains a subprocess $\langle T \mid F \rangle_\iota$.

Reduction. Figures 12 and 13 show the extended process equivalence and reduction rules. Rule E-TRY handles evaluation of the success continuation of an exception handler, and

Reduction of configurations	$\boxed{C \longrightarrow C'}$
$\begin{array}{l} \text{E-RUN} \\ \mathcal{P}[\text{run}(V_m, V_v, V_u, V_e, V_c, \lambda().M)] \S D \longrightarrow \langle \text{extracting}[V_c](V_e V_m) \mid (V_v, V_u, V_e)_0 \mid [M] \S D \\ \text{E-UPDATE} \\ \mathcal{P}[\langle \text{rendering}[V'_m, V_c] U \mid F \rangle_\iota] \S D \longrightarrow \mathcal{P}[\langle \text{idle } V'_m \mid F \rangle_{\iota'} \mid ((M_1))_{\iota'} \mid \dots \mid ((M_n))_{\iota'}] \S D' \\ \text{where } \text{diff}(U, D) = D' \text{ and } \text{procs}(V_c) = \vec{M} \\ \text{E-TRANSITION} \\ \mathcal{P}[\langle \text{transitioning}[V_m, F', V_c] U \mid F \rangle_\iota] \S D \longrightarrow \mathcal{P}[\langle \text{idle } V_m \mid F' \rangle_{\iota'} \mid ((M_1))_{\iota'} \mid \dots \mid ((M_n))_{\iota'}] \S D' \\ \text{where } \iota' = \iota + 1, \text{diff}(U, D) = D' \\ \text{and } \text{procs}(V_c) = \vec{M} \\ \text{E-EVT} \\ P \S \mathcal{D}[\text{domTag}(\text{ev}(W) \cdot \vec{e}) \text{t } U D] \longrightarrow P \mid ((V_1 W))_\iota \mid \dots \mid ((V_n W))_\iota \S \mathcal{D}[\text{domTag}(\vec{e}) \text{t } U D] \\ \text{where } \text{handlers}(\text{ev}, U) = \vec{V} \text{ and } \text{version}(P) = \iota \\ \\ \text{(E-INTERACT, E-STRUCT, E-LIFTP unchanged)} \end{array}$	
Cancellation of pure active thread contexts	$\boxed{\not\leq \mathcal{T}_P}$
$\begin{array}{l} \not\leq \text{updating } E_P = \not\leq E_P \quad \not\leq \text{rendering}[V_m, V_c] E_P = \not\leq V_m \mid \not\leq V_c \mid \not\leq E_P \quad \not\leq \text{extracting}[V_c] E_P = \not\leq V_c \mid \not\leq E_P \\ \not\leq \text{extractingT}[V_c, F] E_P = \not\leq V_c \mid \not\leq E_P \quad \not\leq \text{transitioning}[V_m, F, V_c] E_P = \not\leq V_m \mid \not\leq V_c \mid \not\leq E_P \end{array}$	

■ **Figure 13** Reduction rules for extended calculus (2)

the `procs` meta-definition returns a sequence of processes to be spawned by a command. Process equivalence is extended to allow commutativity of name restrictions, reordering of names in a binder, and scope extrusion. The final “garbage collection” equivalences $(\nu cd)(\not\leq c \mid \not\leq d) \mid P \equiv P$ and $\lfloor () \rfloor \mid P \equiv P$ allow us to discard a channel where both endpoints have been cancelled, and a completed server thread, respectively.

Figure 12 details the extended MVU process reduction rules.

MVU reduction. MVU reduction rules are specific to MVU. Central to safely integrating linearity and transitions are rules E-DISCARD, E-DISCARDHALT, and E-HANDLE. Rule E-HANDLE is modified so that the event loop process only handles a message if the message has the same version. If the versions do not match, then E-DISCARD safely discards any channel endpoints in the discarded message by generating zipper threads. Rules E-EXTRACT, E-EXTRACTT, E-RENDER, and E-RENDERT handle the state machine transitions described in Figure 11 and are used to calculate the new model and HTML.

Session reduction. Session reduction rules encode session-typed communication and are mostly standard: E-NEW generates a name restriction and returns two fresh endpoints; E-COMM handles synchronous communication; and E-CLOSE discards the endpoints of a completed session. The remaining session communication rules handle session cancellation, and are a synchronous variant of Exceptional GV described by Fowler et al. [20]. Rule E-CANCEL discards an endpoint. Rules E-SENDZAP, E-RECVZAP, and E-CLOSEZAP raise an exception if a thread tries to communicate along an endpoint whose peer is cancelled, ensuring resources are discarded safely.

Exception reduction. Rule E-RAISEH describes exception handling: as **raise** occurs in a pure context, the exception is handled by the innermost handler; the rule evaluates the failure continuation and discards all linear resources in the aborted context. Rules E-RAISEURUN and E-RAISEUMAIN apply to unhandled exceptions in a main thread, generating the

Typing rules for names, events, and function state				$\Gamma \vdash M : A$	$\vdash e$	$\Psi \vdash F : \text{State}(A, B, C)$	
$\frac{\Gamma :: \mathbf{U}}{\Gamma, c : S \vdash c : S}$	$\frac{\text{TE-EVT} \quad \vdash V : \text{ty}(\text{ev}) \quad \text{ty}(\text{ev}) :: \mathbf{U}}{\vdash \text{ev}(V)}$	$\frac{\text{TF-STATE} \quad \Psi_1 \vdash V_v : A \rightarrow^{\mathbf{U}} \text{Html}(B) \quad \Psi_2 \vdash V_u : (B \times A) \rightarrow^{\mathbf{U}} \text{Transition}(A, B) \quad \Psi_3 \vdash V_e : A \rightarrow^{\mathbf{U}} (A \times C) \quad C :: \mathbf{U}}{\Psi_1, \Psi_2, \Psi_3 \vdash (V_m, V_v, V_u) : \text{State}(A, B, C)}$					
Typing rules for active threads				$\Psi \vdash T : \text{EvtLoop}(A, B, C)$			
$\frac{\text{TT-IDLE} \quad \Psi \vdash V_m : A}{\Psi \vdash \text{idle } V_m : \text{EvtLoop}(A, B, C)}$		$\frac{\text{TT-UPDATING} \quad \Psi \vdash M : \text{Transition}(A, B)}{\Psi \vdash \text{updating } M : \text{EvtLoop}(A, B, C)}$					
$\frac{\text{TT-RENDERING} \quad \Psi_1 \vdash V_m : A \quad \Psi_2 \vdash V_c : \text{Cmd}(B) \quad \Psi_3 \vdash M : \text{Html}(B)}{\Psi_1, \Psi_2, \Psi_3 \vdash \text{rendering}[V_m, V_c] M : \text{EvtLoop}(A, B, C)}$		$\frac{\text{TT-EXTRACTING} \quad \Psi_1 \vdash V_c : \text{Cmd}(B) \quad \Psi_2 \vdash M : (A \times C)}{\Psi_1, \Psi_2 \vdash \text{extracting}[V_c] M : \text{EvtLoop}(A, B, C)}$					
$\frac{\text{TT-EXTRACTINGT} \quad \Psi_1 \vdash F : \text{State}(A, B, C) \quad \Psi_2 \vdash V_c : \text{Cmd}(B) \quad \Psi_3 \vdash M : (A \times C)}{\Psi_1, \Psi_2, \Psi_3 \vdash \text{extractingT}[F, V_c] M : \text{EvtLoop}(A', B', C')}$							
$\frac{\text{TT-TRANSITIONING} \quad \Psi_1 \vdash V_m : A \quad \Psi_2 \vdash F : \text{State}(A, B, C) \quad \Psi_3 \vdash V_c : \text{Cmd}(B) \quad \Psi_4 \vdash M : \text{Html}(B)}{\Psi_1, \dots, \Psi_4 \vdash \text{transitioning}[V_m, F, V_c] M : \text{EvtLoop}(A', B', C')}$							
Typing rules for processes				$\Psi \vdash_i^\phi P : A$			
$\frac{\text{TP-RUN} \quad \Psi \vdash M : (A \times (A \rightarrow^{\mathbf{U}} \text{Html}(B)) \times ((B \times A) \rightarrow^{\mathbf{U}} \text{Transition}(A, B)) \times (A \rightarrow^{\mathbf{U}} (A \times C)) \times \text{Cmd}(B) \times (\mathbf{1} \rightarrow^{\mathbf{L}} \mathbf{1})) \quad C :: \mathbf{U}}{\Psi \vdash_i^\bullet \text{run } M : B}$							
$\frac{\text{TP-EVENTLOOP} \quad \Psi_1 \vdash T : \text{EvtLoop}(A, B, C) \quad \Psi_2 \vdash F : \text{State}(A, B, C)}{\Psi_1, \Psi_2 \vdash_i^\bullet \langle T \mid F \rangle_i : B}$		$\frac{\text{TP-THREAD} \quad \Psi \vdash M : A}{\Psi \vdash_i^\circ ((M))_i : A}$		$\frac{\text{TP-OLDTHREAD} \quad \Psi \vdash M : B \quad i \neq i'}{\Psi \vdash_i^\circ ((M))_{i'} : A}$		$\frac{\text{TP-SERVERTHREAD} \quad \Psi \vdash M : \mathbf{1}}{\Psi \vdash_i^\circ [M] : A}$	
$\frac{\text{TP-PAR} \quad \Psi_1 \vdash_i^{\phi_1} P_1 : A \quad \Psi_2 \vdash_i^{\phi_2} P_2 : A}{\Psi_1, \Psi_2 \vdash_i^{\phi_1 + \phi_2} P_1 \parallel P_2 : A}$		$\frac{\text{TP-ZAP} \quad c : S \vdash_i^\circ \not\vdash c : A}{\vdash_i^\bullet \text{halt} : A}$		$\frac{\text{TP-NU} \quad \Psi, c : S, d : \bar{S} \vdash_i^\phi P : A}{\Psi \vdash_i^\phi (\nu cd)P : A}$			

■ **Figure 14** Runtime typing for extended calculus

halt configuration and cancelling any linear resources in the aborted context. Rules E-RAISEUPTHREAD and E-RAISEUSERSERVER apply to unhandled exceptions in event loop thread and server threads respectively, by cancelling any channels in the aborted continuation.

Configuration reduction. Figure 13 shows the modified configuration reduction rules. We modify E-RUN to take into account the new arguments, and spawn the given server thread. We modify E-UPDATE to spawn threads described by the returned command; E-TRANSITION is similar but changes the function state and increments the version. We modify E-EVT to tag each spawned event handler thread with the version of the event handler process.

3.4.3 Metatheory

Runtime typing. Figure 14 shows the runtime typing rules for the extended calculus. Rule T-NAME types channel endpoints, and rule TE-EVT mandates that event payload types are unrestricted. The rules for active threads ensure that the types of the terms being evaluated

correspond to the state in the state machine (for example, that the **updating** state returns a term of type $\text{Transition}(A, B)$), and that any recorded values have the correct types.

Let Ψ range over environments containing only runtime names: $\Psi ::= \cdot \mid \Psi, c : S$. We write Ψ_1, Ψ_2 for the disjoint union of environments Ψ_1 and Ψ_2 .

We modify the shape of the process typing judgement to $\Psi \vdash_l^\phi P : A$, which can be read “under typing environment Ψ and thread flag ϕ , process P has type A and version l ”. We modify rule TP-EVENTLOOP to include the extraction function, and mandate that the unrestricted model type C has kind U. We modify rule T-THREAD to state that type of an event handler thread $((M))_l$ has type A if term M has type A and the version matches that of the event handler process. Rule TP-OLDTHREAD allows a thread to have a mismatching type to the event handler process if the versions are incompatible. Finally, TP-ZAP and TP-HALT type zipper threads and the **halt** thread, and TP-NU types a name restriction $(\nu cd)P$ by adding c and d with dual session types into the typing environment.

Properties. The extended calculus satisfies preservation.

► **Theorem 7** (Preservation). *If $\vdash C$ and $C \longrightarrow C'$, then $\vdash C'$.*

Although session types rule out deadlock within a single session, without imposing a tree-like structure on processes [31, 45] (which is too inflexible for our purposes) or using techniques such as channel priorities [29, 37, 39], it is not possible to rule out deadlocks when considering multiple sessions. Since communication over multiple sessions can introduce deadlocks, we begin by proving an error-freedom property, similar to that of Gay and Vasconcelos [21]. An *error process* involves a communication mismatch.

► **Definition 8** (Error process). *A process P is an error process if it contains one of the following processes as a subprocess:*

1. $(\nu cd)(\mathcal{T}[\mathbf{send}(V, c)] \parallel \mathcal{T}'[\mathbf{send}(W, d)])$
2. $(\nu cd)(\mathcal{T}[\mathbf{send}(V, c)] \parallel \mathcal{T}'[\mathbf{close} d])$
3. $(\nu cd)(\mathcal{T}[\mathbf{receive} c] \parallel \mathcal{T}'[\mathbf{receive} d])$
4. $(\nu cd)(\mathcal{T}[\mathbf{receive} c] \parallel \mathcal{T}'[\mathbf{close} d])$

Configuration typing ensures error-freedom.

► **Theorem 9** (Error-freedom). *If $\Psi \vdash_l^\phi P : A$, then P is not an error process.*

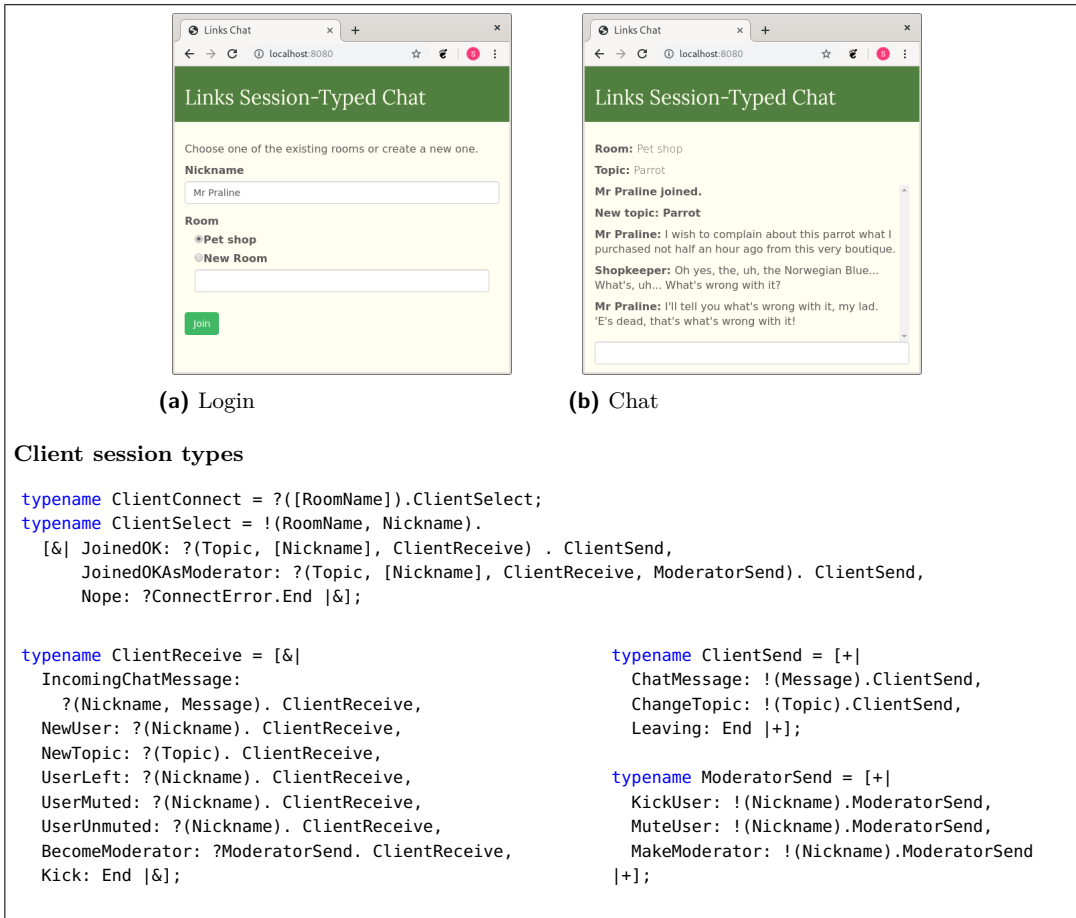
Error-freedom shows that session typing ensures the absence of communication mismatches. What remains is to show that, apart from the possibility of deadlock, the additional features do not interfere with the progress property enjoyed by λ_{MVU} . We begin by classifying the notion of a *blocked* thread, which is a thread blocked on a communication action.

► **Definition 10** (Blocked thread). *We say a thread $\mathcal{T}[M]$ is blocked if either $M = \mathbf{send}(V, W)$, $M = \mathbf{receive} V$, or $M = \mathbf{close} V$.*

Let us refer to **halt**, $\langle T \mid F \rangle_l$, and **run** M as *main threads*, and $((M))_l$, $[M]$, and ζc as *auxiliary threads*. Each well-typed configuration has precisely one main thread.

We can now classify the notion of progress enjoyed by the extended calculus. Either the configuration can reduce; is waiting for an event; has halted due to an unhandled exception; or is deadlocked. Again, let \longrightarrow_E be the \longrightarrow relation without E-INTERACT.

► **Theorem 11** (Weak Event Progress). *Suppose $\vdash C$. Either there exists some C' such that $C \longrightarrow C'$, or there exists some C' such that $C \equiv C'$ and:*



■ **Figure 15** Chat server application

1. D cannot be written $D[\mathbf{domTag}(\vec{e}) \ t \ V \ D]$ for a non-empty \vec{e} .
2. If the main thread of C' is **halt**, then all auxiliary threads are blocked or zipper threads.
3. If the main thread of C' is **run** M , then M is blocked, and all auxiliary threads are either blocked, values, or zipper threads.
4. If the main thread of C' is $\langle T \mid F \rangle_t$, then:
 - a. if $T = \mathbf{idle} \ V_m$, then each auxiliary thread is either blocked or a zipper thread; or
 - b. if $T = \mathcal{T}_A[L]$ then L is blocked, and each auxiliary thread is either blocked, a value, or a zipper thread.

4 Implementation and Example Application

We have implemented an MVU library for the Links tierless web programming language, which includes all extensions in the paper; Links already has a linear type system and distributed session types, so is an ideal fit.

We now describe a chat application, extending the application presented by Fowler et al. [20]. The application (Figure 15) has two main stages shown to the user: on the first, the user is presented with a list of rooms, and enters a username and selects a room. If a user with the given nickname is not already in the selected room, then the user joins the room, receiving the current topic, a list of other nicknames, and a channel used to receive messages from the server. The user can then send chat messages, change the topic, and leave the

room. If the user is the first user in the room, then they join as a moderator and receive an additional channel which can be used to kick, mute, or promote other users to moderators. Users can receive incoming chat messages, and system messages detailing changes such as a new topic or a user joining the room.

We can encode these interaction patterns using session types. Links session type notation for offering a choice is $[&| \dots |&]$, and making a choice is $[+| \dots |+]$. Type `clientConnect` describes the client receiving the room list. Type `clientSelect` describes the client sending the room name and nickname, and receiving the response from the server: either joining as a regular user (`JoinedOK`); joining as a moderator (`JoinedOKAsModerator`); or an error (`Nope`). Types `clientSend` and `clientReceive` detail the messages that the client can send to, and receive from the server, respectively. Type `ModeratorSend` details privileged moderator actions.

Although the original version of Links [10] ran as a CGI script, modern Links applications run as a persistent webserver. Upon execution, the chat application creates an *access point* for sessions of type `clientConnect`, which supports session establishment, and spawns an acceptor thread to accept incoming requests on the access point. Each chat room is represented as a process on the server. When an HTTP request is made, the response contains the MVU application and the access point ID which can be used to establish a session of type `clientConnect`. After the initial HTTP response, further communication between the client and server happens over a WebSocket [16].

The application has three states: connection, chatting, and a “waiting” state shown while waiting for a response. For the purposes of the paper, we consider the connection state.

```

typename SelectedRoom =
  [| NewRoom | SelectedRoom: String |];
typename NotConnectedModel =
  (nickname: String, rooms: [RoomName],
   selectedRoom: SelectedRoom,
   newRoomText: RoomName, error: Maybe(Error));
typename NCMModel =
  (ClientSelect, NotConnectedModel);
typename NCMMessage = [|
  [| UpdateNickname: Nickname
   | UpdateSelectedRoom: SelectedRoom
   | UpdateNewRoom: RoomName | SubmitJoinRoom |];

```

The `NotConnectedModel` is the unrestricted part of the model, and contains the current nickname (`nickname`), list of rooms (`rooms`), selected room (`selectedRoom`), value of the “new room” text box (`newRoomText`), and an optional error message to display (`error`). The model, `NCMModel`, is a pair of a session endpoint of type `clientSelect` and a `NotConnectedModel`. The UI messages are described by the `NCMMessage` type: for example, the `updateNickname` message is generated by the `onInput` event of the nickname input box.

Upon receiving the `submitJoinRoom` UI message when the form is submitted, the application can send the nickname and selected room along the `clientSelect` channel, all of which are contained in the model, without requiring ad-hoc messaging or imperative updates.

5 Related work

Flapjax [34] was the first web programming language to use *functional reactive programming* (FRP) [14] in the setting of web applications. Flapjax provides *behaviours*, which are variables whose contents change over time, and *event streams*, which are an infinite stream of discrete events which change a behaviour. ScalaLocs [46] is a multi-tier reactive programming framework written in Scala, where changes in reactive *signals* are propagated across tiers, rather than using explicit message passing. Ur/Web [9] and WebSharper UI [19] store data in mutable variables, and allow views of the data to be combined using monadic combinators.

Felleisen et al. [15] describe an earlier approach similar to MVU written in the DrScheme [17] system. Similar to the MVU update function, events such as key presses and mouse movements are handled using functions of type $(\text{Model} \times \text{Event}) \rightarrow \text{Model}$. The approach

handles “environment” events rather than events dispatched by individual elements, and the approach is not formalised. Environment events can be handled using *subscriptions* in Elm, which can be added to λ_{MVU} (see the extended version of the paper [18]).

React [2] is a popular JavaScript UI framework. In React, a user defines data models and rendering functions, and similar to Elm, updates are propagated to the DOM by diffing. Differently to MVU, there is no notion of a message, and a page consists of multiple components rather than being derived from a single model. We expect some technical machinery from λ_{MVU} (e.g., event queues, DOM contexts, and diffing) could be reused when formalising React. Redux [5] is a state container for JavaScript applications: to modify the state, one dispatches an *action*, and a function takes the previous state and an action and produces a new state. In combination with React, the approach strongly resembles MVU.

Hop.js [41] is a multi-tier web framework written in JavaScript. Hop.js *services* allow remote function invocation, and the framework supports client-side message-passing concurrency using Web Workers [22], but there is no cross-tier message-passing concurrency.

Session types were introduced by Honda [23] and were first considered in a linear functional language by Gay and Vasconcelos [21]; Wadler [45] later introduced a session-typed functional language GV and a logically-grounded session-typed calculus CP (following Caires and Pfenning [7]), and translated GV into CP. Lindley and Morris [31] introduced an operational semantics for GV, and showed type- and semantics-preserving translations between GV and CP. GV inspires FST [33], which is the core calculus for Links’ treatment of session typing.

Fowler et al. [20] extend GV with failure handling, and extend Links with cross-tier session-typed communication. They do not formally consider GUI development, and their approach to frontend web programming using session types (described in Section 1) leads to a disconnect between the state of the page and the application logic. We build upon their approach to session-typed web programming, while also allowing idiomatic GUI development.

King et al. [27] present a toolchain for writing web applications which respect multiparty session types [25]. Protocols are compiled to PureScript [42] using a parameterised monad [6] to guarantee linearity, and the authors integrate their encoding of session types with the Concur UI framework [26]. Each application may only have a single session connecting the client and server, whereas in our system there may be multiple; our approach supports first-class linearity and cross-tier typechecking; our approach is formalised; and our approach supports failure handling. Links does not yet support multiparty session types.

6 Conclusion

Session types allow conformance to protocols to be checked statically. The last few years have seen a flurry of activity in implementing session types in a multitude of programming languages, but linearity—a vital prerequisite for implementing session types safely—is difficult to reconcile with the asynchronous nature of graphical user interfaces. Consequently, the vast majority of implementations using session types are command line applications, and the few implementations which do integrate session types and GUIs do so in an ad-hoc manner.

In this paper, we have addressed this problem by extending the Model-View-Update architecture, pioneered by the Elm programming language. We have presented the first formal study of MVU by introducing a core calculus, λ_{MVU} . Leveraging our formal characterisation of MVU, we have introduced three extensions: commands, linearity, and model transitions, enabling us to present the first formal integration of session-typed communication with a GUI framework. Informed by our formal model, we have implemented our approach in Links. As future work, we will investigate how to encode allowed transitions as a behavioural type.

References

- 1 Elm: A delightful language for reliable webapps, 2019. URL <http://www.elm-lang.org>. Accessed on 2019-07-04.
- 2 React – a JavaScript library for building user interfaces, 2019. URL <http://www.reactjs.org>. Accessed on 2019-09-02.
- 3 WebSharper.Mvu, 2019. URL <https://github.com/dotnet-websharper/mvu>. Accessed on 2019-07-04.
- 4 Flux, 2020. URL <https://facebook.github.io/flux/>. Accessed on 2020-01-08.
- 5 Redux - a predictable state container for JavaScript apps, 2020. URL <https://redux.js.org/>. Accessed on 2020-01-08.
- 6 Robert Atkey. Parameterised notions of computation. *J. Funct. Program.*, 19(3-4): 335–376, 2009.
- 7 Luís Caires and Frank Pfenning. Session types as intuitionistic linear propositions. In *CONCUR*, volume 6269 of *Lecture Notes in Computer Science*, pages 222–236. Springer, 2010.
- 8 Iliano Cervesato and Frank Pfenning. A linear logical framework. In *LICS*, pages 264–275. IEEE Computer Society, 1996.
- 9 Adam Chlipala. Ur/Web: A simple model for programming the web. In *POPL*, pages 153–165. ACM, 2015.
- 10 Ezra Cooper, Sam Lindley, Philip Wadler, and Jeremy Yallop. Links: Web programming without tiers. In *FMCO*, volume 4709 of *Lecture Notes in Computer Science*, pages 266–296. Springer, 2006.
- 11 Evan Czaplicki. Farewell to FRP, 2016. URL <https://elm-lang.org/news/farewell-to-frp>. Accessed on 2019-09-02.
- 12 Evan Czaplicki and Stephen Chong. Asynchronous functional reactive programming for GUIs. In Hans-Juergen Boehm and Cormac Flanagan, editors, *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '13, Seattle, WA, USA, June 16-19, 2013*, pages 411–422. ACM, 2013. doi: 10.1145/2491956.2462161. URL <https://doi.org/10.1145/2491956.2462161>.
- 13 Ornela Dardha, Elena Giachino, and Davide Sangiorgi. Session types revisited. *Inf. Comput.*, 256:253–286, 2017.
- 14 Conal Elliott and Paul Hudak. Functional reactive animation. In *ICFP*, pages 263–273. ACM, 1997.
- 15 Matthias Felleisen, Robert Bruce Findler, Matthew Flatt, and Shriram Krishnamurthi. A functional I/O system or, fun for freshman kids. In *ICFP*, pages 47–58. ACM, 2009.
- 16 Ian Fette and Alexey Melnikov. The WebSocket protocol, 2011.
- 17 Robert Bruce Findler, John Clements, Cormac Flanagan, Matthew Flatt, Shriram Krishnamurthi, Paul Steckler, and Matthias Felleisen. DrScheme: a programming environment for Scheme. *J. Funct. Program.*, 12(2):159–182, 2002.
- 18 Simon Fowler. Model-View-Update-Communicate: Session types meet the Elm architecture (Extended version), 2019. URL <https://arxiv.org/abs/1910.11108>.
- 19 Simon Fowler, Loïc Denuzière, and Adam Granicz. Reactive single-page applications with dynamic dataflow. In *PADL*, volume 9131 of *Lecture Notes in Computer Science*, pages 58–73. Springer, 2015.
- 20 Simon Fowler, Sam Lindley, J. Garrett Morris, and Sára Decova. Exceptional asynchronous session types: session types without tiers. *PACMPL*, 3(POPL):28:1–28:29, 2019. doi: 10.1145/3290341. URL <https://doi.org/10.1145/3290341>.

- 21 Simon J. Gay and Vasco Thudichum Vasconcelos. Linear type theory for asynchronous session types. *J. Funct. Program.*, 20(1):19–50, 2010.
- 22 Ido Green. *Web Workers - Multithreaded Programs in JavaScript*. O’Reilly, 2012.
- 23 Kohei Honda. Types for dyadic interaction. In *CONCUR*, volume 715 of *Lecture Notes in Computer Science*, pages 509–523. Springer, 1993.
- 24 Kohei Honda, Vasco Thudichum Vasconcelos, and Makoto Kubo. Language primitives and type discipline for structured communication-based programming. In *ESOP*, volume 1381 of *Lecture Notes in Computer Science*, pages 122–138. Springer, 1998.
- 25 Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. *J. ACM*, 63(1):9:1–9:67, 2016.
- 26 Anupam Jain. Concur, 2019. URL <https://ajnsit.github.io/concur>. Accessed on 2019-09-02.
- 27 Jonathan King, Nicholas Ng, and Nobuko Yoshida. Multiparty session type-safe web development with static linearity. In *PLACES@ETAPS*, volume 291 of *EPTCS*, pages 35–46, 2019.
- 28 Naoki Kobayashi. Type systems for concurrent programs. In *10th Anniversary Colloquium of UNU/IIST*, volume 2757 of *Lecture Notes in Computer Science*, pages 439–453. Springer, 2002.
- 29 Naoki Kobayashi. A new type system for deadlock-free processes. In *CONCUR*, volume 4137 of *Lecture Notes in Computer Science*, pages 233–247. Springer, 2006.
- 30 Glenn E. Krasner and Stephen T. Pope. A Cookbook for using the Model-view Controller user interface paradigm in Smalltalk-80. *J. Object Oriented Program.*, 1(3):26–49, August 1988. ISSN 0896-8438. URL <http://dl.acm.org/citation.cfm?id=50757.50759>.
- 31 Sam Lindley and J. Garrett Morris. A semantics for propositions as sessions. In *ESOP*, volume 9032 of *Lecture Notes in Computer Science*, pages 560–584. Springer, 2015.
- 32 Sam Lindley and J. Garrett Morris. Talking bananas: structural recursion for session types. In *ICFP*, pages 434–447. ACM, 2016.
- 33 Sam Lindley and J. Garrett Morris. Lightweight functional session types. *Behavioural Types: from Theory to Tools*. River Publishers, pages 265–286, 2017.
- 34 Leo A. Meyerovich, Arjun Guha, Jacob P. Baskin, Gregory H. Cooper, Michael Greenberg, Aleks Bromfield, and Shriram Krishnamurthi. Flapjax: a programming language for Ajax applications. In *OOPSLA*, pages 1–20. ACM, 2009.
- 35 Dimitris Mostrous and Vasco T. Vasconcelos. Affine sessions. *Logical Methods in Computer Science*, 14(4), 2018. doi: 10.23638/LMCS-14(4:14)2018. URL [https://doi.org/10.23638/LMCS-14\(4:14\)2018](https://doi.org/10.23638/LMCS-14(4:14)2018).
- 36 Joachim Niehren, Jan Schwinghammer, and Gert Smolka. A concurrent lambda calculus with futures. *Theor. Comput. Sci.*, 364(3):338–356, 2006.
- 37 Luca Padovani. Deadlock and lock freedom in the linear π -calculus. In *CSL-LICS*, pages 72:1–72:10. ACM, 2014.
- 38 Luca Padovani. Context-free session type inference. In *ESOP*, volume 10201 of *Lecture Notes in Computer Science*, pages 804–830. Springer, 2017.
- 39 Luca Padovani and Luca Novara. Types for deadlock-free higher-order programs. In *FORTE*, volume 9039 of *Lecture Notes in Computer Science*, pages 3–18. Springer, 2015.
- 40 Adam Pedley. Functional Model-View-Update Architecture for Flutter, 2019. URL <https://buildflutter.com/functional-model-view-update-architecture-for-flutter/>. Accessed on 2019-09-24.
- 41 Manuel Serrano and Vincent Prunet. A glimpse of Hopjs. In *ICFP*, pages 180–192. ACM, 2016.

- 42 The PureScript Contributors. PureScript, 2019. URL <http://www.purescript.org/>. Accessed on 2019-09-02.
- 43 Vasco T. Vasconcelos. Fundamentals of session types. *Inf. Comput.*, 217:52–70, 2012.
- 44 Philip Wadler. Linear types can change the world! In *Programming Concepts and Methods*, page 561. North-Holland, 1990.
- 45 Philip Wadler. Propositions as sessions. *J. Funct. Program.*, 24(2-3):384–418, 2014. doi: 10.1017/S095679681400001X. URL <https://doi.org/10.1017/S095679681400001X>.
- 46 Pascal Weisenburger, Mirko Köhler, and Guido Salvaneschi. Distributed system development with ScalaLoc. *PACMPL*, 2(OOPSLA):129:1–129:30, 2018.