TopHat: A formal foundation for task-oriented programming

Tim Steenvoorden
Software Science
Radboud University
Nijmegen, The Netherlands
tim@cs.ru.nl

Nico Naus
Information and Computing Sciences
Utrecht University
Utrecht, The Netherlands
n.naus@uu.nl

Markus Klinik
Software Science
Radboud University
Nijmegen, The Netherlands
m.klinik@cs.ru.nl

ABSTRACT

Software that models how people work is omnipresent in today's society. Current languages and frameworks often focus on usability by non-programmers, sacrificing flexibility and high level abstraction. Task-oriented programming (TOP) is a programming paradigm that aims to provide the desired level of abstraction while still being expressive enough to describe real world collaboration. It prescribes a declarative programming style to specify multi-user workflows. Workflows can be higher-order. They communicate through typed values on a local and global level. Such specifications can be turned into interactive applications for different platforms, supporting collaboration during execution. TOP has been around for more than a decade, in the forms of iTasks and mTasks, which are tailored for real-world usability. So far, it has not been given a formalisation which is suitable for formal reasoning.

In this paper we give a description of the TOP paradigm and then decompose its rich features into elementary language elements, which makes them suitable for formal treatment. We use the simply typed lambda-calculus, extended with pairs and references, as a base language. On top of this language, we develop TopHat, a language for modular interactive workflows. We describe TopHat by means of a layered semantics. These layers consist of multiple big-step evaluations on expressions, and two labelled transition systems, handling user inputs.

With TopHat we prepare a way to formally reason about TOP languages and programs. This approach allows for comparison with other work in the field. We have implemented the semantic rules of TopHat in Haskell, and the task layer on top of the iTasks framework. This shows that our approach is feasible, and lets us demonstrate the concepts by means of illustrative case studies. TOP has been applied in projects with the Dutch coast guard, tax office, and navy. Our work matters because formal program verification is important for mission-critical software, especially for systems with concurrency.

1 INTRODUCTION

Many applications these days are developed to support workflows in institutions and businesses. Take for example expense declarations, order processing, and emergency management. Some of these workflows occur on the boundary between organisations and customers, like flight bookings or tax returns. What they all have in common is that they need to interact with different people (end users, tax officers, customers, etc.) and they use information from multiple sources (input forms, databases, sensors, etc.).

Tasks. We call interactive units of work based on information sources *tasks*. Tasks model real world collaboration between users, are driven by work users do, and are assigned to some user. Users could be people out in the field or sitting behind their desks, as well as machines doing calculations or fetching data.

Task-oriented programming. Task-oriented programming (TOP) is a programming paradigm which targets the sweet spot between faithful modelling workflows and rapid prototyping of multi-user web applications supporting these workflows [23]. TOP focusses on modelling collaboration patterns. This gives rise to a user's need to interact and share information. Next to that, TOP automatically provides solutions to common development jobs like designing GUIs, connecting to databases, and servers-client communication.

Therefore, a language that supports TOP should choose the right level of abstraction to support two things. Firstly, it should provide primitive building blocks that are useful for high-level descriptions of how users collaborate, with each other and with machines. These building blocks are: *editors*, *composition*, and *shared data*. Secondly, it should be able to generate applications, including graphical user interfaces, from workflows modelled with said building blocks.

Users can work together in a number of ways, and this is reflected in TOP by task compositions. There is *sequential* composition, *parallel* composition, and *choice*. Users need to communicate in order to engage in these forms of collaboration. This is reflected in TOP by three kinds of communication mechanisms. There is data flow *alongside* control flow, where the result of a task is passed onto the next. There is data flow *across* control flow, where information is shared between multiple tasks. Finally, there is communication with the *outside* world, where information is entered into the system via input events. The end points where the outside world interacts with TOP applications are called editors. In generated applications, editors can take many forms, like input fields, selection boxes, or map widgets.

Utilisation. Currently, we know of two frameworks implementing TOP: iTasks and mTasks. iTasks is an implementation of TOP, in the form of a shallowly embedded domain-specific language in the lazy functional programming language Clean. It is a library that provides editors, monadic combinators, and shared data sources. iTasks uses the generic programming facilities of Clean to derive rich client and server applications from a single source. It has been used to model an incident management tool for the Dutch coast guard [15]. Also it has been used numerous times to prototype ideas for Command and Control [12, 24], and in a case study for the Dutch tax authority [25].

mTasks is a subset of iTasks, focusing on IoT devices and deployment on micro controllers. It has been used to control home thermostats and other home automation applications [13]. Both

implementations currently lack formal semantics which are suited to prove properties about tasks.

Challenges. Both iTasks and mTasks have been designed for developing real-world applications. They are constantly being extended and improved with this goal in mind. The different variations of task combinators and the details that come with real-world requirements, make it hard to see what the essence of TOP is. Also, the tight integration of both frameworks with Clean, makes it difficult to see where the boundaries are. This makes formal reasoning about TOP programs impossible.

In this paper, we want to take a step back and look at the spirit of TOP . We do this both formally and informally. Informally in the sense that we give an intuitive description of the features that define task-oriented programming. Formally in the sense that we develop a language which formalises these features as language constructs, and we give them a semantics in the style that is common in programming language research. We separate the task layer and the underlying host language, both syntactically and semantically. Thus making explicit which properties of TOP come from the task layer, and which come from functional programming. Our challenge, therefore, is to model the properties of TOP into a language and pave the way for formal treatment of TOP programs. We give this formal language the name $\widehat{\text{TOP}}$ (TopHat).

Contributions. Our contributions to workflow modelling, functional programming language design, and rapid application development are as follows.

(i) We describe the essential concepts of task-oriented programming. (ii) We present a formal language for modelling declarative workflows, called $\widehat{\text{TOP}}$, embedded in a simply typed λ -calculus. It is based on the aforementioned essential TOP concepts. (iii) We develop a layered operational semantics for $\widehat{\text{TOP}}$ that is driven by user input. The semantics of the task language is clearly separated from the semantics of the underlying host language. (iv) Along with this semantics, we present the following semantic observations on tasks: the current value, whether a term is stuck, the current user interface, and the accepted inputs. (v) We prove progress and type preservation for $\widehat{\text{TOP}}$. (vi) Using both the essential concepts and the formal language, we compare TOP with related work in areas ranging from business process modelling, to process algebras and reactive programming. (vii) We implemented the whole semantic system in the functional programming language Haskell [16]. (viii) To create executable applications, we implemented the task layer of $\widehat{\text{TOP}}$ in iTasks. This also demonstrates that the former is a subset of the latter.

Structure. In Section 2 we demonstrate the functionality of $\widehat{\text{TOP}}$ by means of an example, Section 3 gives an overview of the essential concepts of TOP . Section 4 introduces the $\widehat{\text{TOP}}$ language syntax and Section 5 the semantics. Then in Section 6 we show that certain properties hold for the language. We take a look at related work in Section 7 and conclude in Section 8.

2 EXAMPLE

In this section we develop an example program to demonstrate the capabilities of $\widehat{\text{TOP}}$. The example is a small flight booking system. It demonstrates communication on all three levels: with the

environment, across control flow, and alongside it. Also, it shows synchronisation and input validation.

The requirements of the application are as follows. (i) A user has to enter a list of passengers for which to book tickets. (ii) At least one of these passengers has to be an adult. (iii) After a valid list of passengers has been entered, the user has to pick seats. (iv) Only free seats may be picked. (v) Every passenger must have exactly one seat. (vi) Multiple users should be able to book tickets at the same time.

For this example we assume that the host language has lists and four functions on them: all, any, intersect, and difference. The functions all and any check if all or any elements in a list satisfy a given predicate. The functions intersect and difference compute the set-intersection and set-difference of two lists. We also make use of string equality ($\equiv$), dereferencing (!), reference assignment ($:=$), and expression sequencing (;). For brevity, we omit the type annotations of variable bindings.

Example 2.1 (Flight booking). We start by defining some type aliases. A passenger is a pair with name and age. A seat is a pair with a row number and a seat letter.

```
type PASSENGER = STRING  $\times$  INT
```

```
type SEAT = INT  $\times$  STRING
```

Choosing seats requires reading and updating shared information. The list of free seats is stored in a reference.

```
let freeSeats = ref [(1,"A"), (1,"B"), (1,"C"), ...]
```

Now we develop our workflow in a top-down manner. Our flight booking starts with an interactive task $\boxtimes(\text{LIST PASSENGER})$, where users can enter a list of passengers. A task $\boxtimes \tau$ is an empty editor that asks for a value of the given type τ . Passengers are valid if their name is not empty and their age is at least 0. Lists of passengers are valid if each passenger is valid, and at least one of the passengers is an adult. When the user has entered a valid list of passengers, the step after $\triangleright$ becomes enabled, and the user can proceed to picking seats. In case of an invalid list of passengers, the step is guarded by the failing task ζ .

```
let valid =  $\lambda p$ . not (fst  $p \equiv ""$ )  $\wedge$  snd  $p \geq 0$  in
```

```
let adult =  $\lambda p$ . snd  $p \geq 18$  in
```

Figure 1: Running web application of the flight booking example using a translation to iTasks. It shows three users booking a flight simultaneously. The first user entered name and age and continued picking seats. The second is entering details of two passengers. The ages are not filled in, therefore the Continue button is disabled. The message bubble shows that the age field only accepts integer values. The third user finished a booking, therefore, the first user can not pick seats 1b and 1c any more.

```

let allValid =  $\lambda ps.$  all valid  $ps \wedge$  any adult  $ps$  in
let bookFlight =  $\boxtimes$ (LIST PASSENGER)  $\triangleright \lambda ps.$ 
if allValid  $ps$  then chooseSeats  $ps$  else  $\perp$ 

```

A selection of seats is correct if every entered seat is free.

```

let correct =  $\lambda ss.$  intersect  $ss$  !freeSeats  $\equiv ss$  in
let chooseSeats =  $\lambda ps.$   $\boxtimes$ (LIST SEAT)  $\triangleright \lambda ss.$ 
if correct  $ss \wedge$  length  $ps \equiv$  length  $ss$ 
then confirmBooking  $ps$   $ss$  else  $\perp$ 

```

The function confirmBooking removes the selected seats from the shared list of free seats, and displays the end result using an editor, denoted by $\square$.

```

let confirmBooking =  $\lambda ps.$   $\lambda ss.$ 
  freeSeats := difference !freeSeats  $ss; \square\langle ps, ss \rangle$ 

```

The main task starts three bookFlight tasks, which could be performed by three different users in parallel.

```

bookFlight  $\bowtie$  bookFlight  $\bowtie$  bookFlight

```

A screenshot of the running application is shown in Fig. 1.

All instances of the bookFlight task have access to the shared list of free seats. Rewriting the example in a language without side effects would not only be cumbersome, obfuscating the code with explicit threading of state, but it would be impossible to model the parallel execution of three bookFlight tasks. It is not known upfront which task will finish first, and thus it is not possible to thread the free seat list between the parallel tasks.

3 INTUITION

This section gives an overview of the core concepts of task-oriented programming.

Tasks model collaboration. The central objective of TOP is to *coordinate collaboration*. The basic building blocks of TOP for expressing collaboration are task combinators. They express ways in which people can work together. Tasks can be executed after each other, at the same time, or conditionally. This motivates the combinators step, parallel, and choice.

Example 3.1 (Breakfast). The following program shows the different collaboration operators in the setting of making breakfast. Users have a choice ($\diamond$) whether they want tea or coffee. They always get an egg. The drink and the food are prepared in parallel ($\bowtie$). When both the drink and the food are prepared, users can step ($\triangleright$) to eating the result.

```

let mkBrkfst : TASK Drink  $\rightarrow$  TASK Food  $\rightarrow$  TASK  $\langle$ Drink,Food $\rangle$ 
  =  $\lambda mkDrink. \lambda mkFood. mkDrink \bowtie mkFood$  in
  mkBrkfst (mkTea  $\diamond$  mkCoffee) mkEgg  $\triangleright$  enjoyMorning

```

The way the combinators are defined matches real life closely. When we want to have breakfast, we have to complete several other tasks first before we can do so. We decide what we want to have and then prepare it. We can prepare the different items we have for breakfast in parallel, but not at the same time. For example, it is impossible to scramble eggs, and put on the kettle for tea simultaneously. Instead, what is meant by parallel is that the order in which we do tasks and the smaller tasks that they are composed of, does not matter. Then finally, only when every item we want to have for breakfast is ready, can we sit down and enjoy it.

Tasks are reusable. There are three ways in which tasks are modular. First, larger tasks are composed of smaller ones. Second, tasks are first-class, they can be arguments and results of functions. Third, tasks can be values of other tasks. These aspects make it possible for programmers to model custom collaboration patterns. Example 3.1 demonstrates how tasks can be parameterised by other tasks: mkBrkfst is a collaboration pattern that always works the same way, regardless of which food and drink are being prepared.

Tasks are driven by user input. Input events drive evaluation of tasks. When the system receives a valid event, it applies this event to the current task, which results in a new task. In this way the system communicates with the environment. Inputs are synchronous, which means the order of execution is completely determined by the order of the inputs.

In TOP, editors are the basic method of communication with the environment. Editors are modelled after input widgets from graphical user interfaces. There are different editors, denoted by different box symbols. Take for example an editor holding the integer seven: $\square 7$. Such an editor reacts to two kinds of events. The first are change events, for example the values 42 or 37, which are of the same type. The second is a clear event E, which empties the editor. This results in an empty editor $\boxtimes \text{INT}$, annotated to only hold integers.

The sole purpose of editors is to interact with users by remembering the last value that has been sent to them. There are no output events. As values of editors can be observed, for example by a user interface, editors facilitate both input and output. An empty editor ($\boxtimes$) stands for a prompt to input data, while a filled editor ($\square$) can be seen either as outputting a value, or as an input that comes with a default value.

Example 3.2 (Vending machine). This example demonstrates external communication and choice. It is a vending machine that dispenses a biscuit for one coin and a chocolate bar for two coins.

```

let vend : TASK SNACK =  $\boxtimes \text{INT} \triangleright \lambda n.$  if  $n \equiv 1$  then  $\square \text{Biscuit}$ 
else if  $n \equiv 2$  then  $\square \text{ChocolateBar}$  else  $\perp$ 

```

The editor $\boxtimes \text{INT}$ asks the user to enter an amount of money. This editor stands for a coin slot in a real machine that freely accepts and returns coins. There is a continue button that is initially disabled, due to the fact that the left hand side of the step combinator has no value. When the user has inserted exactly 1 or 2 coins, the continue button becomes enabled. When the user presses the continue button, the machine dispenses either a biscuit or a chocolate bar, depending on the amount of money. Snacks are modelled using a custom type.

Tasks can be observed. Several observations can be made on tasks. One of those is determining the value of a task. Not all tasks have a value, the empty editor for instance, which makes value observation partial. I.e., the value of $\square 7$ is 7, but the value of $\boxtimes \text{INT}$ is $\perp$.

Another observation is the set of input events a task can react to. For example, the task $\square 7$ can react to value events and empty events, as discussed before.

In order to render a task, we need to observe a task's user interface. This is done compositionally. User interfaces of combined tasks are composed of the user interfaces of the components. For example, of two tasks combined with a step combinator, only the left hand side is rendered. Two parallel tasks are rendered next to each other. Combining this information with the task's value

and possible inputs, we can display the current state of the task, together with buttons that show the actions a user can engage in.

The final observation is to determine whether a task results in a failure, denoted by $\perp$. The step combinator $\triangleright$ and the choice combinator $\diamond$ use this to prevent users from picking a failing task.

Tasks are never done. Tasks never terminate, they always keep reacting to events. Editors can always be changed or cleared, and step combinators move on to new tasks.

In a step $t \triangleright e$, the decision to move on from a task t to its continuation e is taken by $\triangleright$, not by t . The decision is based on a speculative evaluation of e . The step combinator $t \triangleright e$ passes the value v of t to the continuation e . Steps act like t as long as the step is guarded. A step is guarded if either the left task has no value, or the speculative evaluation of e applied to v yields the failure task $\perp$. Once it becomes unguarded, the step continues as the result of $e v$. Speculative evaluation is designed so that possible side effects are undone. The task $t \triangleright e$ additionally requires a continue event C to proceed.

Step combinators give rise to a form of internal communication. They represent data flow that *follows* control flow.

Tasks can share information. The step combinator is one form of internal communication, where task values are passed to continuations. Another form of internal communication is shared data. Shared data enables data flow *across* control flow, in particular between parallel tasks. Shared data sources are assignable references whose changes are immediately visible to all tasks interested in them. Users can not directly interact with shared data, a shared editor is required for that. If x is a reference of type τ , then $\blacksquare x$ is an editor whose value is that of x .

The semantics of $\widehat{\text{TOP}}$ requires all updates to shared data and all enabled internal steps to be processed before any further communication with the environment can take place.

Example 3.3 (Cigarette smokers). The cigarette smokers problem [7] is a surprisingly tricky synchronisation problem. We study it here because it demonstrates the capabilities of guarded steps. The problem is stated as follows. In order to smoke a cigarette, three ingredients are required: tobacco, paper, and a match. There are three smokers, each having one of the ingredients and requiring the other two. There is an agent that randomly provides two of those. The difficulty lies in the requirement that only the smoker may proceed whose missing ingredients are present.

Downey models availability of the ingredients with a semaphore for each ingredient. The agent randomly signals two of the three. The solution proposed by Downey involves an additional mutex, three additional semaphores, three additional threads called *pushers*, and three regular Boolean variables. The job of the pushers is to record availability of their ingredient in their Boolean variable, and check availability of other resources, waking the correct smoker when appropriate.

What is important is that the implementation of what is essentially deadlock-free waiting for two semaphores requires a substantial amount of additional synchronisation, together with non-trivial conditional statements. $\widehat{\text{TOP}}$ allows a simple solution to this problem, using guarded steps. Steps can be guarded with arbitrary expressions. The parallel combinator can be used to watch two

shared editors at the same time. Let match, paper, and tobacco be references to Booleans. The smokers are defined as follows.

```
let continue =  $\lambda \langle x, y \rangle . \text{if } x \wedge y \text{ then smoke else } \perp$  in
let tobaccoSmoker = ( $\blacksquare$  match  $\bowtie$   $\blacksquare$  paper)  $\triangleright$  continue in
let paperSmoker = ( $\blacksquare$  tobacco  $\bowtie$   $\blacksquare$  match)  $\triangleright$  continue in
let matchSmoker = ( $\blacksquare$  tobacco  $\bowtie$   $\blacksquare$  paper)  $\triangleright$  continue in
tobaccoSmoker  $\bowtie$  paperSmoker  $\bowtie$  matchSmoker
```

When the agent supplies two of the ingredients by setting the respective shares to True, only the step of the smoker that waits for those becomes enabled.

Tasks are predictable. Let t_1 and t_2 be tasks. The parallel combination $t_1 \bowtie t_2$ stands for two independent tasks carried out at the same time. This operator introduces interleaving concurrency. For the system it does not matter if the tasks are executed by two people actually in parallel, or by one person who switches between the tasks. The inputs sent to the component tasks are interleaved into a serial stream, which is sent to the parallel combinator. We assume that such a serialization is always possible. The tasks are truly independent of each other, all interleavings are permitted. The environment must prefix events to t_1 and t_2 respectively by F (first) and S (second). This unambiguously renames the inputs, removing any source of nondeterminism.

With concurrency comes the need for synchronisation, in situations where only some but not all interleavings are desired. The basic method for synchronisation in $\widehat{\text{TOP}}$ is built into the step combinator. The task $t \triangleright e$ can only continue execution when two conditions are met: Task t must have a value v , and $e v$ must not evaluate to $\perp$. Programmers can encode arbitrary conditions in $e v$, which are evaluated atomically between interaction steps. This allows a variety of synchronisation problems to be solved in an intuitive and straight-forward manner.

Hoare [9] states that nondeterminism is only ever useful for *specifying* systems, never for implementing them. $\widehat{\text{TOP}}$ is meant solely for implementation and does not have any form of nondeterminism. Input events for parallel tasks are disambiguated, internal steps ($\triangleright$) have a well-defined evaluation order, and internal choice ($\diamond$) is left-biased.

Recap. Collaboration in the real world consists of three aspects: communication, concurrency, and synchronisation. These aspects are reflected in TOP on a high level of abstraction, hiding the details of communication. For example, the cigarette smokers communicate with each other, but the programs do not explicitly mention sending or receiving events.

By focusing on collaboration instead of communication, TOP leads to specifications closer to real-world workflows which, at the same time, can be used to generate multi-user applications to support such workflows.

4 LANGUAGE

In this section, we present the constructs of $\widehat{\text{TOP}}$, our modular interactive workflow language. We define the host and task language, the types, and the static semantics. Then we describe the workings of each construct using examples. These constructs are formalised in Section 5.

4.1 Expressions

The host language is a simply typed λ -calculus, extended with some basic types and ML-style references. We use references to represent shared data sources. The grammar in Fig. 2 defines the syntax of the host language. It has abstractions, applications, variables, and constants for booleans, integers and strings. The symbol $\star$ stands for binary operators. For the result of parallel tasks we need pairs. Conditionals come in handy for defining guards. References will be used to implement shared editors. Our treatment of references closely follows the one by Pierce [22]. Creating a reference using the keyword **ref** yields a location l . While x denotes program variables, l denotes store locations. Locations are not intended to be directly manipulated by the programmer. The symbols $!$ and $:=$ stand for dereferencing and assignment. The unit value will be used as the result of assignments.

$e ::=$	Expressions
$\lambda x : \tau. e \mid e_1 e_2$	– abstraction, application
$x \mid c \mid e_1 \star e_2$	– variable, constant, operation
if e_1 then e_2 else $e_3 \mid \langle \rangle$	– branch, unit
$\langle e_1, e_2 \rangle \mid \text{fst } e \mid \text{snd } e$	– pair, projections
ref $e \mid !e \mid e_1 := e_2 \mid l$	– references, location
p	– pretask
$c ::=$	Constants
$B \mid I \mid S$	– boolean, integer, string

Figure 2: Language grammar

We use double quotation marks to denote strings. Integers are denoted by their decimal representation, and booleans are written True and False. We freely make use of the logic operators not, $\wedge$, and $\vee$, arithmetic operators $+$, $-$, $\times$, and the string append operator $\mathbin{\&\#}$. Furthermore, we use standard comparison operations $<$, $\leq$, $\equiv$, $\neq$, $\geq$, and $>$. The symbol $\star$ stands for any of those. The notation $e_1; e_2$ is an abbreviation for $(\lambda x : \text{UNIT}. e_2) e_1$, where x is a fresh variable. The notation **let** $x : \tau = e_1$ **in** e_2 is an abbreviation for $(\lambda x : \tau. e_2) e_1$.

The grammar in Fig. 3 specifies the syntactic category of *pretasks*. Pretasks are tasks that have unevaluated subexpressions with respect to the host language. How expressions are evaluated will be discussed in Section 5.1. Each pretask will be discussed in more detail in the following subsections. We use open symbols ($\square, \boxtimes, \triangleright, \diamond$) for tasks that require user input, and closed symbols ($\blacksquare, \blacktriangleright, \blacklozenge$) for tasks that can be evaluated without user input.

$p ::=$	Pretasks
$\square e \mid \boxtimes \tau \mid \blacksquare e$	– editors: valued, unvalued, shared
$e_1 \blacktriangleright e_2 \mid e_1 \triangleright e_2$	– steps: internal, external
$\frac{1}{2} \mid e_1 \boxtimes e_2$	– fail, combination
$e_1 \blacklozenge e_2 \mid e_1 \diamond e_2$	– choice: internal, external

Figure 3: Task grammar

Typing. Figure 4 shows the grammar of types used by $\widehat{\text{TOP}}$. It has functions, pairs, basic types, unit, references, and tasks.

Typing rules are of the form $\Gamma, \Sigma \vdash e : \tau$, which should be read as “in environment Γ and store typing Σ , expression e has type τ ”. Typing rules for expressions in the host language are presented in the appendix. The typing rules for pretasks are given in Fig. 5. Most typing rules lift the type of their subexpressions into the Task-type. The typing rules for steps make sure the continuations e_2 are

$\tau ::=$	Types
$\tau_1 \rightarrow \tau_2 \mid \tau_1 \times \tau_2 \mid \beta$	– function, product, basic
UNIT $\mid \text{REF } \tau \mid \text{TASK } \tau$	– unit, reference, task
$\beta ::=$	Basic types
BOOL $\mid \text{INT} \mid \text{STRING}$	– boolean, integer, string

Figure 4: Type grammar

functions which accept a well-typed value from the left hand side (T-THEN, T-NEXT). References, and therefore shared editors, can only be of a basic type so they do not introduce implicit recursion (T-UPDATE).

$\Gamma, \Sigma \vdash e : \tau$		
T-EDIT $\frac{\Gamma, \Sigma \vdash e : \tau}{\Gamma, \Sigma \vdash \square e : \text{TASK } \tau}$	T-ENTER $\frac{\Gamma, \Sigma \vdash \tau : \text{TASK } \tau}{\Gamma, \Sigma \vdash \boxtimes \tau : \text{TASK } \tau}$	T-UPDATE $\frac{\Gamma, \Sigma \vdash e : \text{REF } \beta}{\Gamma, \Sigma \vdash \blacksquare e : \text{TASK } \beta}$
T-THEN $\frac{\Gamma, \Sigma \vdash e_1 : \text{TASK } \tau_1 \quad \Gamma, \Sigma \vdash e_2 : \tau_1 \rightarrow \text{TASK } \tau_2}{\Gamma, \Sigma \vdash e_1 \blacktriangleright e_2 : \text{TASK } \tau_2}$	T-NEXT $\frac{\Gamma, \Sigma \vdash e_1 : \text{TASK } \tau_1 \quad \Gamma, \Sigma \vdash e_2 : \tau_1 \rightarrow \text{TASK } \tau_2}{\Gamma, \Sigma \vdash e_1 \triangleright e_2 : \text{TASK } \tau_2}$	
T-FAIL $\frac{}{\Gamma, \Sigma \vdash \frac{1}{2} : \text{TASK } \tau}$	T-AND $\frac{\Gamma, \Sigma \vdash e_1 : \text{TASK } \tau_1 \quad \Gamma, \Sigma \vdash e_2 : \text{TASK } \tau_2}{\Gamma, \Sigma \vdash e_1 \boxtimes e_2 : \text{TASK } (\tau_1 \times \tau_2)}$	
T-OR $\frac{\Gamma, \Sigma \vdash e_1 : \text{TASK } \tau \quad \Gamma, \Sigma \vdash e_2 : \text{TASK } \tau}{\Gamma, \Sigma \vdash e_1 \blacklozenge e_2 : \text{TASK } \tau}$	T-XOR $\frac{\Gamma, \Sigma \vdash e_1 : \text{TASK } \tau \quad \Gamma, \Sigma \vdash e_2 : \text{TASK } \tau}{\Gamma, \Sigma \vdash e_1 \diamond e_2 : \text{TASK } \tau}$	

Figure 5: Typing rules

4.2 Editors

Programs in $\widehat{\text{TOP}}$ model interactive workflows. Interaction means communication with end users. End users should be able to enter information into the system, change it, clear it, reenter it, and so on. To do this, we introduce the concept of *editors*. Editors are typed containers that either hold a value or are empty. Editors that have a value can be *changed* or *emptied*. Empty editors can be *filled*.

Editors stand for various forms of input and output, for example widgets in a GUI, form fields on a webpage, sensors, or network connections. Consider the editor for a person’s age from Example 2.1. Users can change the value until they are satisfied with it. Editors are meant to capture this constantly changing nature of user input. The user interface of an editor depends on its type. This could be an input field for strings, a toggle switch for booleans, or even a map with a pin for locations. It could also be a parser that tries to parse a line of text to match the type of the editor.

Valued and unvalued editors ($\square e, \boxtimes \tau$). Editors that hold an expression $e : \tau$ have type $\text{TASK } \tau$. Empty editors are annotated with a type in order to ensure type safety. Without type annotations, it would be possible to change the type of an editor, by clearing it and filling it with a value of different type.

Shared editors ($\blacksquare e$). Shared editors watch references, lifting their value into the task domain. If e is a reference $\text{REF } \tau$, then $\blacksquare e$ is of type $\text{TASK } \tau$. Shared editors cannot be cleared, only changed.

Changes to a shared editor are immediately visible to all shared editors watching the same reference. Imagine two users, Marco and

Christopher, both watching shared editors of the same coordinates. The editors are visualised as a pin on a map. When Marco moves his pin, he updates the value of the shared editor, thereby changing the value of the reference. This change is immediately reflected on Christopher's screen: The pin changes its position on his map. This way Marco and Christopher can work together to edit the same information.

Two other important use cases for shared editors are sensors and time. Sensors can be represented as external entities that periodically update a shared editor with their current sensor value. Similarly, the current time can be stored in a shared editor ($\blacksquare$ time) which is periodically updated by a clock. The actual sensor and the clock are not modelled in $\widehat{\text{TOP}}$. We assume that they exist as external users that send update events to the system. This allows programmers to write tasks that react to sensor values or timeouts.

4.3 Steps

Editors represent atomic units of work. In this section we look at ways to compose smaller tasks into bigger ones. Composing tasks can be done in two ways, sequential and parallel. Parallel composition comes in two variants: combining two tasks (*and*-parallel) and choosing between two tasks (*or*-parallel). We study sequential composition first, and after that combining and choosing.

Internal and external step ($t \blacktriangleright e, t \triangleright e$). Sequential composition has a task t on the left and a continuation e on the right. External steps ($\triangleright$) must be triggered by the user, while internal steps ($\blacktriangleright$) are taken automatically. The accompanying typing rules are T-THEN and T-NEXT. According to these rules, the left hand side must be a task $t : \text{TASK } \tau_1$, and the right hand side $e : \tau_1 \rightarrow \text{TASK } \tau_2$ must be a function that, given the task value of t , calculates the task with which to continue.

Steps are guarded, which means that the step combinators can only proceed when the following conditions are met. The left hand side must have a value, only then can the right hand side calculate the successor task. The successor task must not be ζ , introduced below. This is enforced on the semantic level, as described in the next section. The internal step can proceed immediately when these conditions are met. The external step must additionally receive a continue event C .

Example 4.1 (Conditional stepping). Consider the following:

$\boxtimes \text{INT} \blacktriangleright \lambda n. \text{if } n \equiv 42 \text{ then } \square \text{"Good"} \text{ else } \square \text{"Bad"}$

Initially, the step is guarded because the editor does not have a value. When users enter an integer, the program continues immediately with either $\square \text{"Good"}$ or $\square \text{"Bad"}$, depending on the input.

Fail (ζ). Fail is a task that never has a value and never accepts input. The typing rule T-FAIL states that it has type $\text{TASK } \tau$ for any type τ . Programmers can use ζ to tell steps that no sensible successor task can be determined.

Example 4.2 (Guarded stepping). Consider this slight variation on Example 4.1:

$\boxtimes \text{INT} \blacktriangleright \lambda n. \text{if } n \equiv 42 \text{ then } \square \text{"Good"} \text{ else } \zeta$

The user is asked to enter an integer. As long as the right hand side of $\blacktriangleright$ evaluates to ζ , the step cannot proceed, and the user can keep editing the integer. As soon as the value of the left hand side is 42,

the right hand side evaluates to something other than ζ , and the step proceeds to $\square \text{"Good"}$.

Example 4.3 (Waiting). With the language constructs seen so far it is possible to create a task that waits for a specified amount of time. To do this, we make use of a shared editor holding the current time (see Section 4.2), and a guarded internal step.

$\text{let wait} : \text{INT} \rightarrow \text{TASK UNIT} = \lambda \text{amount} : \text{INT}.$

$\blacksquare \text{time} \blacktriangleright \lambda \text{start} : \text{INT}.$

$\blacksquare \text{time} \blacktriangleright \lambda \text{now} : \text{INT}.$

$\text{if } \text{now} > \text{start} + \text{amount} \text{ then } \square \langle \rangle \text{ else } \zeta$

The first step is immediately taken, resulting in start to be the time at the moment wait is executed. The second step is guarded until the current time is greater to the start time plus the requested *amount*.

4.4 Parallel

A common pattern in workflow design is splitting up work into multiple tasks that can be executed simultaneously. In $\widehat{\text{TOP}}$, all parallel branches can progress independently, driven by input events. This requires inputs to be tagged in order to reach the intended task.

There are two ways to proceed after a parallel composition. One way is to wait for all tasks to produce results and combine those, the other to pick the first available result. Both ways introduce explicit forks and implicit joins in $\widehat{\text{TOP}}$.

Combination ($e_1 \bowtie e_2$). A combination of two tasks is a parallel *and*. It has a value only if both branches have a value. This is reflected in the typing rule T-AND, It shows that if the first task has type τ_1 , and the second has type τ_2 , their combination has the pair type $\tau_1 \times \tau_2$.

Example 4.4 (Combining). The task

$\boxtimes \text{INT} \bowtie \square \text{"Batman"} \blacktriangleright \lambda \langle n, s \rangle . \square (\text{replicate } n \text{"Na"} + s)$

can only step when both editors have values. When it steps, the continuation uses the pair to calculate the result.

Internal and external choice ($e_1 \blacklozenge e_2, e_1 \lozenge e_2$). Internal choice ($\blacklozenge$) is a parallel *or*. It picks the leftmost branch that has a value. Its typing rule T-OR states that both branches must have the same type $\text{TASK } \tau$. For example $\boxtimes \text{INT} \blacklozenge \square 37$ normalizes to $\square 37$, because $\boxtimes \text{INT}$ doesn't have a value. Users can work on both branches of an internal choice simultaneously.

External choice ($\lozenge$) is different in this regard. An external choice requires users to pick a branch before continuing with it. This means users cannot work on the branches of an external choice before picking one.

Example 4.5 (Delay). We illustrate the use of internal and external choice by means of an example that asks users to proceed with a given task or to cancel. If the user does not make a choice within a given time frame, the program proceeds automatically. The example makes use of the task wait from Example 4.3.

$\text{let cancel} : \text{TASK UNIT} = \square \langle \rangle \text{ in}$

$\text{let delay} : \text{INT} \rightarrow \text{TASK UNIT} \rightarrow \text{TASK UNIT} = \lambda n. \lambda \text{proceed}.$

$(\text{proceed} \lozenge \text{cancel}) \blacklozenge (\text{wait } n \blacktriangleright \lambda u : \text{UNIT}. \text{proceed})$

Note that delay is higher-order. It is a task which takes another task as parameter.

4.5 Annotations

Tasks can be annotated with additional information. The system can use this information in various ways. Possible use cases are labels for the user interface, resource consumption information for static resource analysis, or messages for automatic end-user feedback. Annotations are not covered in this paper. Our Haskell implementation of TOP supports annotating tasks with user IDs, so that individual tasks in a large workflow can be assigned to different users. These annotations are used to filter the user interfaces for each user so that they can only see their part of the workflow.

5 SEMANTICS

In this section we formalise the semantics of the language constructs described in Section 4. We organise this by following the structure of the language. Firstly, the task language is embedded in a simply typed λ -calculus. This requires a specification of the *evaluation* of terms in the host language, and how it handles the task language. Secondly, there are two ways to drive evaluation of task expressions, internally by the system itself, and externally by the user. This is done in two additional semantics, one for the internal *normalisation* of tasks, and another for the *interaction* with the end user.

The three main layers of semantics are thus evaluation, normalisation, and interaction. The semantics, together with *observations*, will be discussed in the following subsections. Figure 6 shows the relation between all semantics arrows. It also shows that there are two helper semantics, *handle* and *stride*. We use the convention that downward arrows are big-step semantics, and rightward arrows are small-step semantics.

Figure 6: Semantic functions defined in this report and their relation.

One of our explicit goals is to keep the semantics for evaluation and normalisation separate, to not mix general purpose programming notions with workflow specific semantics. This is achieved by letting tasks be values in the host language.

5.1 Evaluating expressions

The host language evaluates expressions using a big-step semantics. Evaluating an expression e in state σ into a value v in state σ' is denoted by $e, \sigma \Downarrow v, \sigma'$. To ease reasoning about references, we choose a call-by-value evaluation strategy.

Figure 7 shows values with respect to the evaluation semantics. Tasks are values, and the operands of task constructors are evaluated eagerly. Exceptions to this are steps and external choice, where some or all of the operands are not evaluated.

The rules to evaluate expressions e do not differ from standard work, except for the task constructs. The evaluation rules for tasks can be deduced from the value grammar. They are given in the appendix. Most task constructors are strict in their arguments. Only steps keep their right hand side unevaluated to delay side effects till the moment the step is taken. The same holds for both branches of the external choice.

$v ::=$		Values	
$\lambda x : \tau. e$	$\langle v_1, v_2 \rangle$	$\langle \rangle$	– abstraction, pair, unit
c	l	t	– constant, location, task
$t ::=$		Tasks	
$\square v$	$\boxtimes \tau$	$\blacksquare l$	– editors
$t_1 \blacktriangleright e_2$	$t_1 \triangleright e_2$		– steps
$\frac{1}{2}$	$t_1 \bowtie t_2$		– fail, combination
$t_1 \blacklozenge t_2$	$e_1 \lozenge e_2$		– choices

Figure 7: Value grammar

5.2 Task observations

The normalisation and interaction semantics make use of observations on tasks. Observations are semantic functions on the syntax tree of tasks. There are four semantic functions: $\mathcal{V}$ for the current task value, $\mathcal{F}$ to determine if a task fails, $\mathcal{I}$ for the currently accepted input events, and a function for generating user interfaces. The semantics make use of $\mathcal{V}$ and $\mathcal{F}$, while $\mathcal{I}$ is used for proving safety. The function for user interfaces is not used by the semantics, but by our implementation. It is only described in passing here.

Observable values ($\mathcal{V}$). Task values are used by steps to calculate the successor task. Filled editors are tasks which contain values, as are shared editors. Unvalued editors do not contain values, neither does the fail task. These facts propagate through all other task constructors. The partial function $\mathcal{V}$ associates a value v to tasks t where possible. Its definition is given in Fig. 8.

$\mathcal{V} : \text{Tasks} \times \text{States} \rightarrow \text{Values}$

$\mathcal{V}(\square v, \sigma)$	$= v$
$\mathcal{V}(\boxtimes \tau, \sigma)$	$= \perp$
$\mathcal{V}(\blacksquare l, \sigma)$	$= s(l)$
$\mathcal{V}(\frac{1}{2}, \sigma)$	$= \perp$
$\mathcal{V}(t_1 \blacktriangleright e_2, \sigma)$	$= \perp$
$\mathcal{V}(t_1 \triangleright e_2, \sigma)$	$= \perp$
$\mathcal{V}(t_1 \bowtie t_2, \sigma)$	$= \begin{cases} \langle v_1, v_2 \rangle & \text{when } \mathcal{V}(t_1, \sigma) = v_1 \wedge \mathcal{V}(t_2, \sigma) = v_2 \\ \perp & \text{otherwise} \end{cases}$
$\mathcal{V}(t_1 \blacklozenge t_2, \sigma)$	$= \begin{cases} v_1 & \text{when } \mathcal{V}(t_1, \sigma) = v_1 \\ v_2 & \text{when } \mathcal{V}(t_1, \sigma) = \perp \wedge \mathcal{V}(t_2, \sigma) = v_2 \\ \perp & \text{otherwise} \end{cases}$
$\mathcal{V}(t_1 \lozenge t_2, \sigma)$	$= \perp$

Figure 8: Values

Internal and external steps do not have an observable value, because calculating the value would require evaluation of the continuation. Parallel composition only has a value when both branches have values, in which case these values are paired. Internal choice has a value when one of the branches has a value. When both branches have a value, it takes the value of the left branch. External choice does not have a value because it waits for user input.

Failing ($\mathcal{F}$). In Section 4.3 we introduced $\frac{1}{2}$ to stand for an impossible task. Combinations of tasks can also be impossible. Take for example the parallel composition of two fails ($\frac{1}{2} \bowtie \frac{1}{2}$). This expression is equivalent to $\frac{1}{2}$, because it can not handle input and can not be further normalised. This intuition is formalised by the function $\mathcal{F}$ in Fig. 9. It determines whether a task is impossible. Such tasks are called *failing*.

Steps whose left hand sides are failing can never proceed because of the lack of an observable value. Therefore they are itself failing. The internal choice of two failing tasks is failing. External choices

$\mathcal{F} : \text{Tasks} \times \text{States} \rightarrow \text{Booleans}$
 $\mathcal{F}(\Box v, \sigma) = \text{False}$
 $\mathcal{F}(\Box \tau, \sigma) = \text{False}$
 $\mathcal{F}(\blacksquare l, \sigma) = \text{False}$
 $\mathcal{F}(\frac{1}{2}, \sigma) = \text{True}$
 $\mathcal{F}(t_1 \blacktriangleright e_2, \sigma) = \mathcal{F}(t_1, \sigma)$
 $\mathcal{F}(t_1 \triangleright e_2, \sigma) = \mathcal{F}(t_1, \sigma)$
 $\mathcal{F}(t_1 \bowtie t_2, \sigma) = \mathcal{F}(t_1, \sigma) \wedge \mathcal{F}(t_2, \sigma)$
 $\mathcal{F}(t_1 \blacklozenge t_2, \sigma) = \mathcal{F}(t_1, \sigma) \wedge \mathcal{F}(t_2, \sigma)$
 $\mathcal{F}(e_1 \blacklozenge e_2, \sigma) = \mathcal{F}(t_1, \sigma'_1) \wedge \mathcal{F}(t_2, \sigma'_2)$
 where $e_1, \sigma \Downarrow t_1, \sigma'_1$ and $e_2, \sigma \Downarrow t_2, \sigma'_2$

Figure 9: Failing

let the user pick a side and only then evaluate the corresponding subexpression. To determine if an external choice is failing, it needs to peek into the future to check if both subexpressions are failing.

User interface. $\widehat{\text{TOP}}$ is designed such that a user interface can be generated from a task's syntax tree. A possible graphical user interface is shown in Fig. 1, where tasks are rendered as HTML pages. Editors are rendered as input fields, external choices are represented by two buttons, and parallel tasks are rendered side by side. Steps only show the interface of their left hand side. In case of an external step they are accompanied by a button. When the guard condition of a step is not fulfilled, the button is disabled.

5.3 Normalising tasks

The normalisation semantics is responsible for reducing expressions of type Task until they are ready to handle input. It is a big-step semantics, and makes use of evaluation of the host language. We write $e, \sigma \Downarrow t, \sigma'$ to describe that an expression e in state σ normalises to task t in state σ' .

Normalisation rules are given in Fig. 10. Both rules ensure that expressions are first evaluated by the host language ($\Downarrow$), and then by the stride semantics ($\rightsquigarrow$). These two actions are repeated until the resulting state and task stabilise.

$$\boxed{e, \sigma \Downarrow t, \sigma'}$$

$$\text{N-DONE} \quad \frac{e, \sigma \Downarrow t, \sigma' \quad t, \sigma' \rightsquigarrow t', \sigma''}{e, \sigma \Downarrow t', \sigma''} \quad \sigma' = \sigma'' \wedge t = t'$$

$$\text{N-REPEAT} \quad \frac{e, \sigma \Downarrow t, \sigma' \quad t, \sigma' \rightsquigarrow t', \sigma'' \quad t', \sigma'' \Downarrow t'', \sigma'''}{e, \sigma \Downarrow t'', \sigma'''} \quad \sigma' \neq \sigma'' \vee t \neq t'$$

Figure 10: Normalisation semantics

The split between striding and normalisation is due to mutable references. Consider the following example, where $\sigma = \{l \mapsto \text{False}\}$.

$(\blacksquare l \blacktriangleright \lambda x:\text{Bool}. \text{if } x \text{ then } e \text{ else } \frac{1}{2}) \bowtie (l := \text{True}; \Box \langle \rangle)$

S-AND reduces this expression in one step to

$(\blacksquare l \blacktriangleright \lambda x:\text{Bool}. \text{if } x \text{ then } e \text{ else } \frac{1}{2}) \bowtie (\Box \langle \rangle)$

with $\sigma' = \{l \mapsto \text{True}\}$. This expression is not normalised, because the left task can take a step. The issue here lies in the fact that the right task updates l . To overcome this problem, the N-DONE and N-REPEAT rules ensure that striding is applied until the state σ becomes stable and no further normalisation can take place.

The striding semantics is responsible for reducing internal steps and internal choices. A stride from task t in state σ to t' in state σ' is denoted by $t, \sigma \rightsquigarrow t', \sigma'$. The rules for striding are given in Fig. 11. Tasks like editors, fail and external choice are not further reduced. For external choice and parallel there are congruence rules.

$$\boxed{t, \sigma \rightsquigarrow t', \sigma'}$$

$$\text{Step.}$$

$$\text{S-THENSTAY} \quad \frac{t_1, \sigma \rightsquigarrow t'_1, \sigma' \quad \mathcal{V}(t'_1, \sigma') = \perp}{t_1 \blacktriangleright e_2, \sigma \rightsquigarrow t'_1 \blacktriangleright e_2, \sigma'} \quad \mathcal{V}(t'_1, \sigma') = \perp$$

$$\text{S-THENFAIL} \quad \frac{t_1, \sigma \rightsquigarrow t'_1, \sigma' \quad e_2 \ v_1, \sigma' \Downarrow t_2, \sigma''}{t_1 \blacktriangleright e_2, \sigma \rightsquigarrow t'_1 \blacktriangleright e_2, \sigma'} \quad \mathcal{V}(t'_1, \sigma') = v_1 \wedge \mathcal{F}(t_2, \sigma'')$$

$$\text{S-THENCONT} \quad \frac{t_1, \sigma \rightsquigarrow t'_1, \sigma' \quad e_2 \ v_1, \sigma' \Downarrow t_2, \sigma''}{t_1 \blacktriangleright e_2, \sigma \rightsquigarrow t_2, \sigma''} \quad \mathcal{V}(t'_1, \sigma') = v_1 \wedge \neg \mathcal{F}(t_2, \sigma'')$$

$$\text{Choose.}$$

$$\text{S-ORLEFT} \quad \frac{t_1, \sigma \rightsquigarrow t'_1, \sigma' \quad t_1 \blacklozenge t_2, \sigma \rightsquigarrow t'_1, \sigma'}{t_1 \blacklozenge t_2, \sigma \rightsquigarrow t'_1, \sigma'} \quad \mathcal{V}(t'_1, \sigma') = v_1$$

$$\text{S-ORRIGHT} \quad \frac{t_1, \sigma \rightsquigarrow t'_1, \sigma' \quad t_2, \sigma' \rightsquigarrow t'_2, \sigma''}{t_1 \blacklozenge t_2, \sigma \rightsquigarrow t'_2, \sigma''} \quad \mathcal{V}(t'_1, \sigma') = \perp \wedge \mathcal{V}(t'_2, \sigma'') = v_2$$

$$\text{S-ORNONE} \quad \frac{t_1, \sigma \rightsquigarrow t'_1, \sigma' \quad t_2, \sigma' \rightsquigarrow t'_2, \sigma''}{t_1 \blacklozenge t_2, \sigma \rightsquigarrow t'_1 \blacklozenge t'_2, \sigma''} \quad \mathcal{V}(t'_1, \sigma') = \perp \wedge \mathcal{V}(t'_2, \sigma'') = \perp$$

$$\text{Ready.}$$

$$\text{S-EDIT} \quad \frac{\Box v, \sigma \rightsquigarrow \Box v, \sigma}{\Box v, \sigma \rightsquigarrow \Box v, \sigma} \quad \text{S-FILL} \quad \frac{\Box \tau, \sigma \rightsquigarrow \Box \tau, \sigma}{\Box \tau, \sigma \rightsquigarrow \Box \tau, \sigma} \quad \text{S-UPDATE} \quad \frac{\blacksquare l, \sigma \rightsquigarrow \blacksquare l, \sigma}{\blacksquare l, \sigma \rightsquigarrow \blacksquare l, \sigma}$$

$$\text{S-FAIL} \quad \frac{\frac{1}{2}, \sigma \rightsquigarrow \frac{1}{2}, \sigma}{\frac{1}{2}, \sigma \rightsquigarrow \frac{1}{2}, \sigma} \quad \text{S-XOR} \quad \frac{e_1 \blacklozenge e_2, \sigma \rightsquigarrow e_1 \blacklozenge e_2, \sigma}{e_1 \blacklozenge e_2, \sigma \rightsquigarrow e_1 \blacklozenge e_2, \sigma}$$

$$\text{Congruence.}$$

$$\text{S-NEXT} \quad \frac{t_1, \sigma \rightsquigarrow t'_1, \sigma' \quad t_1 \blacktriangleright e_2, \sigma \rightsquigarrow t'_1 \blacktriangleright e_2, \sigma'}{t_1 \blacktriangleright e_2, \sigma \rightsquigarrow t'_1 \blacktriangleright e_2, \sigma'}$$

$$\text{S-AND} \quad \frac{t_1, \sigma \rightsquigarrow t'_1, \sigma' \quad t_2, \sigma' \rightsquigarrow t'_2, \sigma''}{t_1 \bowtie t_2, \sigma \rightsquigarrow t'_1 \bowtie t'_2, \sigma''}$$

Figure 11: Striding semantics

Principles of stepping. Stepping away from task t_1 can only be performed when t_1 has a value: $\mathcal{V}(t_1) = v_1$. Only then can a new task t_2 be calculated from the expression e . On top of that, t_2 must not be failing: $\neg \mathcal{F}(t_2)$. These principles lead to the stepping rules in Fig. 11. S-THENSTAY does nothing, because the left side does not have a value. S-THENFAIL covers the case that the left side has a value but the calculated successor task is failing. This rule uses the semantics of the host language to evaluate the application $e_2 \ v_1$. When all required conditions are fulfilled, S-THENCONT allows stepping to the successor task.

Principles of choosing. Choosing between two tasks t_1 and t_2 can only be done when at least one of them has a value: $\mathcal{V}(t_1) = v_1 \vee \mathcal{V}(t_2) = v_2$. When both have a value, the left task is chosen. When none has a value, none can be chosen. These principles lead to the rules S-ORLEFT, S-ORRIGHT, and S-ORNONE, which encode that the choice operator picks the leftmost task that has a value.

5.4 Handling user inputs

The handling semantics is the outermost layer of the stack of semantics. It is responsible for performing external steps and choices, and for changing the values of editors. The rules of the interaction semantics are given in Fig. 12. The semantics is only applicable to normalised t . Sending an input event i to a task t first handles the event and then prepares the resulting task for the next input by normalising it.

$$\boxed{t, \sigma \xRightarrow{i} t', \sigma'}$$

$$\text{I-HANDLE} \quad \frac{t, \sigma \xrightarrow{i} t', \sigma' \quad t', \sigma' \Downarrow t'', \sigma''}{t, \sigma \xRightarrow{i} t'', \sigma''}$$

Figure 12: Interaction semantics

Inputs i are formed according to the grammar in Fig. 13. F and S in an input encode the path to the task for which the input is designated. There is a function $\mathcal{I}$ which calculates the possible input events a given task expects. It takes a normalised task and a state and returns a set of inputs that can be handled. The definition of this function is listed in Fig. 14.

$i ::=$	Inputs
$ \quad a \mid F i \mid S i$	– action, pass to first, pass to second
$a ::=$	Actions
$ \quad v \mid E \mid C$	– change, empty, continue
$ \quad L \mid R$	– go left, go right

Figure 13: Input grammar

$$\begin{aligned} \mathcal{I} : \text{Tasks} \times \text{States} &\rightarrow \mathcal{P}(\text{Inputs}) \\ \mathcal{I}(\Box v, \sigma) &= \{v' : \tau, E\} \quad \text{where } \Box v : \text{TASK } \tau \\ \mathcal{I}(\boxtimes \tau, \sigma) &= \{v' : \tau\} \\ \mathcal{I}(\blacksquare l, \sigma) &= \{v' : \tau\} \quad \text{where } \blacksquare l : \text{TASK } \tau \\ \mathcal{I}(\frac{1}{2}, \sigma) &= \emptyset \\ \mathcal{I}(t_1 \blacktriangleright e_2, \sigma) &= \mathcal{I}(t_1, \sigma) \\ \mathcal{I}(t_1 \triangleright e_2, \sigma) &= \mathcal{I}(t_1, \sigma) \cup \{C \mid \mathcal{V}(t_1, \sigma) = v_1 \wedge \\ &\quad e_2 v_1, \sigma \Downarrow t_2, \sigma' \wedge \neg \mathcal{F}(t_2, \sigma')\} \\ \mathcal{I}(t_1 \bowtie t_2, \sigma) &= \{F i \mid i \in \mathcal{I}(t_1, \sigma)\} \cup \{S i \mid i \in \mathcal{I}(t_2, \sigma)\} \\ \mathcal{I}(t_1 \blacklozenge t_2, \sigma) &= \{F i \mid i \in \mathcal{I}(t_1, \sigma)\} \cup \{S i \mid i \in \mathcal{I}(t_2, \sigma)\} \\ \mathcal{I}(e_1 \diamond e_2, \sigma) &= \{L \mid e_1, \sigma \Downarrow t_1, \sigma' \wedge \neg \mathcal{F}(t_1, \sigma')\} \cup \\ &\quad \{R \mid e_2, \sigma \Downarrow t_2, \sigma' \wedge \neg \mathcal{F}(t_2, \sigma')\} \end{aligned}$$

Figure 14: Inputs

Handling input is done by the *handling* semantics shown in Fig. 15. It is a small step semantics with labelled transitions. It takes a task t in a state σ and an input i , and yields a new task t' in a new state σ' .

H-CHANGE, H-EMPTY, H-FILL, H-UPDATE: Input events v and E are used to change the value of editors and empty them. Editors only accept values of the correct type.

H-NEXT: The C(ontinue) action triggers an external step. As with internal stepping, this is only possible if the left side has a value and the continuation is not failing.

H-PICKLEFT, H-PICKRIGHT: L and R are used to pick the left or right option of an external choice.

H-PASSTHEN, H-PASSNEXT: The step combinators pass all events other than C to the left side.

$$\boxed{t, \sigma \xrightarrow{i} t', \sigma'}$$

$$\text{Editing.}$$

$$\text{H-CHANGE} \quad \frac{}{\Box v, \sigma \xrightarrow{v'} \Box v', \sigma} \quad v, v' : \tau$$

$$\text{H-EMPTY} \quad \frac{}{\Box v, \sigma \xrightarrow{E} \boxtimes \tau, \sigma} \quad v : \tau$$

$$\text{H-FILL} \quad \frac{}{\boxtimes \tau, \sigma \xrightarrow{v'} \Box v', \sigma} \quad v' : \tau$$

$$\text{H-UPDATE} \quad \frac{}{\blacksquare l, \sigma \xrightarrow{v'} \blacksquare l, \sigma[l \mapsto v']} \quad \sigma(l), v' : \tau$$

$$\text{Continuing.}$$

$$\text{H-NEXT} \quad \frac{e_2 v_1, \sigma \Downarrow t_2, \sigma' \quad \mathcal{V}(t_1, \sigma) = v_1 \wedge \neg \mathcal{F}(t_2, \sigma')}{t_1 \triangleright e_2, \sigma \xrightarrow{C} t_2, \sigma'}$$

$$\text{H-PICKLEFT} \quad \frac{e_1, \sigma \Downarrow t_1, \sigma' \quad \neg \mathcal{F}(t_1, \sigma')}{e_1 \diamond e_2, \sigma \xrightarrow{L} t_1, \sigma'}$$

$$\text{H-PICKRIGHT} \quad \frac{e_2, \sigma \Downarrow t_2, \sigma' \quad \neg \mathcal{F}(t_2, \sigma')}{e_1 \diamond e_2, \sigma \xrightarrow{R} t_2, \sigma'}$$

$$\text{Passing.}$$

$$\text{H-PASSTHEN} \quad \frac{t_1, \sigma \xrightarrow{i} t'_1, \sigma'}{t_1 \blacktriangleright e_2, \sigma \xrightarrow{i} t'_1 \blacktriangleright e_2, \sigma'}$$

$$\text{H-PASSNEXT} \quad \frac{t_1, \sigma \xrightarrow{i} t'_1, \sigma'}{t_1 \triangleright e_2, \sigma \xrightarrow{i \neq C} t'_1 \triangleright e_2, \sigma'}$$

$$\text{H-FIRSTAND} \quad \frac{t_1, \sigma \xrightarrow{i} t'_1, \sigma'}{t_1 \bowtie t_2, \sigma \xrightarrow{F i} t'_1 \bowtie t_2, \sigma'}$$

$$\text{H-SECONDAND} \quad \frac{t_2, \sigma \xrightarrow{i} t'_2, \sigma'}{t_1 \bowtie t_2, \sigma \xrightarrow{S i} t_1 \bowtie t'_2, \sigma'}$$

$$\text{H-FIRSTOR} \quad \frac{t_1, \sigma \xrightarrow{i} t'_1, \sigma'}{t_1 \blacklozenge t_2, \sigma \xrightarrow{F i} t'_1 \blacklozenge t_2, \sigma'}$$

$$\text{H-SECONDOR} \quad \frac{t_2, \sigma \xrightarrow{i} t'_2, \sigma'}{t_1 \blacklozenge t_2, \sigma \xrightarrow{S i} t_1 \blacklozenge t'_2, \sigma'}$$

Figure 15: Handling semantics

H-FIRSTAND, H-SECONDAND, H-FIRSTOR, H-SECONDOR: Inputs F(irst) and S(econd) are used to direct inputs to the correct branch of parallel combinations.

5.5 Implementation

The semantics have been implemented in the Haskell programming language [16]. We use techniques presented by Jaskelioff et al. [10], Swierstra [27], and Peyton Jones [21]. The source code can be found on GitHub.¹ A command-line interface is part of this implementation. It prompts users to type input events, which get parsed and processed by the interaction semantics.

Also, we made an implementation of $\widehat{\text{TOP}}$ combinators on top of iTasks, so that $\widehat{\text{TOP}}$ specifications can be compiled to runnable applications. This shows that $\widehat{\text{TOP}}$ is a subset of iTasks.

6 PROPERTIES

In order to show our semantics is sane, we show that our evaluation, normalisation and handling semantics is type preserving. We additionally prove a progress theorem for our small-step handling semantics. We show that our failing function $\mathcal{F}$ indeed only indicates expressions that can not be normalised and that allow no further interaction. Finally, we prove that the function to compute all possible inputs $\mathcal{I}$ is sound and complete.

¹<https://github.com/timjs/tophat-haskell>

6.1 Type preservation

We show that the following three preservation Theorems hold.

THEOREM 6.1 (TYPE PRESERVATION UNDER EVALUATION). *For all expressions e and states σ such that $\Gamma, \Sigma \vdash e : \tau$ and $\Gamma, \Sigma \vdash \sigma$, if $e, \sigma \Downarrow e', \sigma'$, then $\Gamma, \Sigma \vdash e' : \tau$ and $\Sigma \vdash \sigma'$.*

Where $\Sigma \vdash \sigma$ means that for all $l \in \sigma$, it holds that $\Sigma(l) = \Gamma, \Sigma \vdash \sigma(l)$.

THEOREM 6.2 (TYPE PRESERVATION UNDER NORMALISATION). *For all expressions e and states σ such that $\Gamma, \Sigma \vdash e : \text{TASK } \tau$ and $\Sigma \vdash \sigma$, if $e, \sigma \Downarrow e', \sigma'$, then $\Gamma, \Sigma \vdash e' : \text{TASK } \tau$ and $\Sigma \vdash \sigma'$.*

THEOREM 6.3 (TYPE PRESERVATION UNDER HANDLING). *For all expressions e , states σ and inputs i such that $\Gamma, \Sigma \vdash e : \text{TASK } \tau$ and $\Sigma \vdash \sigma$, if $e, \sigma \xrightarrow{i} e', \sigma'$, then $\Gamma, \Sigma \vdash e' : \text{TASK } \tau$ and $\Sigma \vdash \sigma'$.*

All three Theorems are proven to be correct by induction over e . The full proofs are listed in the appendix. From Theorem 6.3 and Theorem 6.2 we directly obtain that the driving semantics also preserves types.

6.2 Progress

A well-typed term of a task type is guaranteed to progress after normalisation, unless it is failing.

We define what we mean with progress in Theorem 6.4.

THEOREM 6.4 (PROGRESS UNDER HANDLING). *For all well typed expressions e and states σ , if $e, \sigma \Downarrow e', \sigma'$, then either $\mathcal{F}(e', \sigma')$ or there exist e'', σ'' , and i such that $e', \sigma' \xrightarrow{i} e'', \sigma''$.*

Where a well typed expression e means that $\Gamma, \Sigma \vdash e : \tau$ for some type τ , and a well typed state means that $\Sigma \vdash \sigma$.

If an expression e and state σ are well-typed, then after normalisation, the pair e', σ' either fails, or there exists some input i that can be handled by it under the handling semantics. In order to prove this Theorem, we need to show that the failing function $\mathcal{F}$ behaves as desired.

THEOREM 6.5 (FAILING MEANS NO INTERACTION POSSIBLE). *For all well typed expressions e and states σ , and $e, \sigma \Downarrow e', \sigma'$, we have that $\mathcal{F}(e', \sigma') = \text{True}$, if and only if there is no input i such that $e', \sigma' \xrightarrow{i} e'', \sigma''$ for some e'' and σ'' .*

The Theorem above states that an expression e and state σ are failing, if, after normalisation, there exists no input that can be handled by it. We prove the theorem to be true by induction on e' .

We now have the ingredients to prove Theorem 6.4.

PROOF. Given $\Gamma, \Sigma \vdash e : \text{TASK } \tau$ and $\Sigma \vdash \sigma$ and after normalisation $e, \sigma \Downarrow e', \sigma'$, we find ourselves in either one of the following situations:

There exists an i such that $e', \sigma' \xrightarrow{i} e'', \sigma''$.

There does not exist an i such that $e', \sigma' \xrightarrow{i} e'', \sigma''$. In this case, we know that $\mathcal{F}(e', \sigma')$ must be true, by Theorem 6.5. $\square$

6.3 Soundness and Completeness of Inputs

In order to validate the function that calculates all possible inputs $\mathcal{I}$, we want to show that the set of possible inputs it produces is both sound and complete with respect to the handle semantics. By

sound we mean that all inputs in the set of possible inputs can actually be handled by the handle semantics, and by complete we mean that the set of possible inputs contains all inputs that can be handled by the handle semantics. Theorem 6.6 expresses exactly this property.

THEOREM 6.6 (INPUTS FUNCTION IS SOUND AND COMPLETE). *For all well typed expressions e , states σ , and inputs i , we have that $i \in \mathcal{I}(e, \sigma)$ if and only if $e, \sigma \xrightarrow{i} e', \sigma'$.*

We prove the above theorem by induction over e . The proof is listed in the appendix.

6.4 Outlook

At this point we have specified a formal language for task-oriented programming, given its semantics, and proved its safety. The main motive to formalise this paradigm, is to be able to reason about tasks. In future work, we plan on utilising the formalisation to do so. Firstly, we would like to express properties of tasks and prove them. For example, one would like to prove that, no matter what, in Example 3.1 breakfast is always being served. Secondly, we would like to explore what it means for two tasks to be equal. One could have noticed that some operators have a monadic or applicative feeling. The combination of $\bowtie$ and $\square$ could form a (lax) monoidal functor, $\blacklozenge$ is similar to applicative choice, and $\blacktriangleright$ looks like a bind operation. We need a correct understanding of equivalence of tasks, taking the interactive setting into account, to prove this. Thirdly, we do not know yet if the more complex combinators of iTasks are expressible in the basic combinators of $\widehat{\text{TOP}}$. We implemented $\widehat{\text{TOP}}$ on top of iTasks , so we know it is a subset, but we also know iTasks can do more. A more in depth description of future work can be found in Section 8.

7 RELATED WORK

The work presented in this paper lies on the boundary of many areas of study. People have looked at the problem of how to model and coordinate collaboration from many different perspectives. The following subsections give an overview of related work from the many different areas.

7.1 TOP implementations

iTasks. As mentioned earlier, iTasks is an implementation of TOP . iTasks has many features, and its basic combinators are versatile and powerful. Simpler combinators are implemented by restricting the powerful ones. This is useful for everyday programming, where having lots of functionality at one's fingertips is convenient and efficient. $\widehat{\text{TOP}}$ on the other hand does not include the many different variations of the step- and parallel combinators of iTasks . To name two examples, the combinators $(\gg|)$ and $(||-)$ are variations of step and parallel that ignore the value of the left task.

There have been two previous papers that describe semantics of iTasks , by Koopman et al. [14] and Plasmeijer et al. [23]. Both give a different semantics in the form of minimal implementations of a subset of the interface of iTasks . These semantics however do not make an explicit distinction between the host language and task language and they do not provide an explicit formal semantics. Therefore, they do not lend itself well for formal reasoning.

mTasks. The *mTasks* framework [13] is an implementation of $\widehat{\text{TOP}}$ geared towards IoT devices. As $\widehat{\text{TOP}}$ its basic combinators are a subset of *iTasks*. They are similar to those of $\widehat{\text{TOP}}$. However, on IoT devices it is useful to continue running tasks endlessly, which is done in *mTasks* using a *forever* combinator. This is currently not possible in $\widehat{\text{TOP}}$.

As for *iTasks*, there is currently no formal semantics for *mTasks*.

7.2 Workflow modelling

Much research has been done into workflow modelling. This work focusses on describing the collaboration between subsystems, rather than the communication between them. The systems described in the literature follow a *boxes and arrows* model of specifying workflows. Control flow, represented by arrows, usually can go unrestricted from anywhere to anywhere else in a workflow. We see $\widehat{\text{TOP}}$ as the functional programming of workflows, as opposed to this *goto*-style.

Workflow patterns. Workflow patterns are regarded as special kind of the design patterns in software engineering. They identify recurring patterns in workflow systems, much like the combinators defined by $\widehat{\text{TOP}}$. Work by van der Aalst et al. [30] defines a comprehensive list of these patterns, and examines their availability in industry workflow software. Workflow patterns are usually described in terms of control flow graphs, and no formal specification is given, which makes comparison and formal reasoning more difficult.

Workflow Nets & YAWL. Workflow Nets (WFN) [28] allow for the modelling and analysis of business processes. They are graphical in nature, and clearly display how every component is related to each other. A downside of WFN is that they do not facilitate higher order constructs. Also, they are often not directly executable.

A language based on WFN that is actually directly executable is YAWL by van der Aalst and ter Hofstede [29]. It facilitates modelling and execution of dynamic workflows, with support for *and*, *or* and *xor* workflow patterns. As mentioned, YAWL programs consist of WFN, and are therefore programmed visually.

BPEL. BPEL [20] is another popular business process calculus. The standardised language allows for the specification of actions within business processes, using an XML format. The language is mainly used for coordinating web services. Two workflow patterns are supported; execution of services can be done sequential or in parallel. On top of that, processes can be guarded by conditionals. There is no support for higher order processes however. Processes described in BPEL can be regarded as activity graphs, and they can also be rendered as such. The specified processes in BPEL are directly executable, just like YAWL.

7.3 Process algebras

Differences. There are two main differences between $\widehat{\text{TOP}}$ and process algebras. The first is a difference in scope. Process algebras focus on modelling the input/output behaviour of processes, by explicitly stating which actions are sent and received at certain points in the program. The goal of process algebras is formal reasoning about the interaction between processes. Typically, one wishes to

prove properties such as deadlock-freedom, liveness, or adherence to a protocol specification.

The focus of $\widehat{\text{TOP}}$ on the other hand is to model collaboration patterns, with the explicit goal of not having to specify how exactly subtasks communicate. The declarative specification of data dependencies between subtasks enables $\widehat{\text{TOP}}$ to hide such details.

The second difference concerns internal communication. There are two forms of communication between tasks: Passing values to continuations and sharing data. This is different from communication in process algebras, which is based on message-passing.

Similarities. There are some aspects that are similar in $\widehat{\text{TOP}}$ and process algebras. Internal communication in Hoare's CSP [9] is introduced with the concealment operator. The semantics of CSP requires that all concealed actions are handled to exhaustion before any action with the environment can take place. This is somewhat similar to $\widehat{\text{TOP}}$, where all enabled internal steps must be taken until the system can react to input events again. Contrast this with Milner's CCS [18], where concealed actions are visible to the outside as τ -actions, and can be interleaved with external communication.

Another similarity between $\widehat{\text{TOP}}$ and process algebras, or any system with concurrency for that matter, is the need for synchronisation. Broadly speaking, concurrency means that different parts of a program can interact with the environment independently, in an interleaved manner. Synchronisation means that only some, but not all, of the possible interleavings are desirable. The semantics of the step combinators in $\widehat{\text{TOP}}$, together with the fact that internal communication happens atomically, allows for concise and intuitive synchronisation code.

7.4 Reactive programming

HipHop & Esterel. HipHop [3, 4] is a programming language tailored to the development of synchronous reactive web systems. From a single source, both server and client applications can be generated. Programs are written in the Hop language, a Scheme dialect. Communication is based on a reactive layer embedded in Hop. The set of HipHop reactive statements is based on those of the Esterel language [2, 5]. Each reactive component starts by specifying possible input and output events. The component then proceeds as a state machine.

Input events are sent to such a machine programmatically using Hop, or are explicitly wired to events from the client. They are optionally associated with a Hop value. As Hop is a dynamic language, and HipHop uses strings to identify events, events and their possible associated values are not statically checked. Events are aggregated until the moment the machine is asked to react. The machine is executed and reacts by building a multi-set of output events. The execution of a HipHop machine is atomic. The set of inputs is not influenced by the current computations.

As with $\widehat{\text{TOP}}$, HipHop is a DSL embedded in a general purpose programming language. Another similarity is that both specifications lead to executable server and client applications from a single source. However, both HipHop and Esterel are more low level regarding their specification. Where $\widehat{\text{TOP}}$ takes tasks and collaboration as a starting point, HipHop focusses on synchronous communication and atomic execution of reactive machines.

This difference in focus shows in the way both systems define events. In HipHop programmers can define and use their own events. Inputs in $\widehat{\text{TOP}}$ are not extensible and not visible to the developer. They are a completely separate entity living on the semantic level.

Another important difference is the way in which both systems handle events. In HipHop the programmer decides when a machine should process its events. This could be just one event, or a multi-set of events that are processed simultaneously. $\widehat{\text{TOP}}$ always processes an input the moment it occurs and only handles a single event in one instance.

Functional reactive programming. Functional Reactive Programming (FRP) is a paradigm to describe dynamic changes of values in a declarative way. This is done by specifying networks of values, called behaviours, that can depend on each other and on external events. Behaviours can change over time, or triggered by events. When a behaviour changes, all other behaviours that depend on it are updated automatically. The underlying implementation that takes care of the updating usually can tie input devices, like mouse and keyboard, to event streams and behaviours to output facilities, like text fields. This allows for declarative specifications of applications with user interfaces.

The idea of FRP was pioneered by Elliott and Hudak [8]. In the meantime there are many variants and implementations, where reactive-banana [1], FrTime [6], and Flapjax [17] belong to the most well-known.

FRP and TOP are different systems that have different goals in mind. Whereas FRP expresses automatically updating data dependencies, TOP expresses collaboration patterns. TOP has no notion of time. Tasks cannot change spontaneous over time, while behaviours can. Only input events can change task values. The biggest conceptual difference between a workflow in TOP and a data network in FRP is that an event to a task only causes updates up until the next step, while an event in FRP propagates through the whole network.

That being said, there are some concepts that are similar in TOP and FRP. The *stepper* behaviour, for example, is associated with an event and yields the value of the most recent event. This is similar to editors in TOP. Furthermore, both systems can be used to declaratively program user interfaces, albeit in FRP the programmer has to construct the GUI elements manually, and connect inputs and outputs to the correct events and behaviours. In TOP graphical user interfaces are automatically derived.

7.5 Session types

Session types are a type discipline that can be used to check whether communicating programs conform to a certain protocol. Session types are expressions in some process calculus that describe the input/output behaviour of such programs. Session types are useful for programming languages where modules communicate with each other via messages, like CSP, π -calculus, or Go, to name a few. The only form of messages in TOP are input events which drive execution, but modules do not communicate using messages. Therefore, session types are not applicable to TOP in the sense used in the literature.

Formal reasoning about TOP programs is one of our future goals for $\widehat{\text{TOP}}$. The ideas and techniques of session types could be useful

for specifying that a list of inputs of a certain form leads to desired task values. The details are a topic for future work.

8 CONCLUSION

In this paper we have identified and intuitively described the essence of task-oriented programming. We then formalised this essence by developing a domain-specific language for declarative interactive workflows, called $\widehat{\text{TOP}}$. The task language and the host language are clearly separated, to make explicit where the boundaries are. The semantics of the task layer is driven by user input. We have compared $\widehat{\text{TOP}}$ with workflow modelling languages, process algebras, functional reactive programming and session types to point out differences and similarities. Finally, we have proven type safety and progress for our language.

Future work. There are a couple of ways in which we would like to continue this line of work.

One of the main motivations to formalise task-oriented programming is to be able to reason about programs. In this paper we reason about the language itself, but it would be nice to prove properties about individual programs. To this end, we are very interested to see if it is possible to develop an axiomatic semantics for $\widehat{\text{TOP}}$ that allows us to do so. There are certain properties of our language that make this particularly complex: We have to deal with parallelism, user interaction, and references.

We would also like to prove whether certain programs are equivalent, for example to show that the monad laws hold for our step combinator. This requires a notion of equality, which in the presence of side effects most certainly needs some form of coalgebraic input-output conformance. We have implemented the reduction semantics of our language in Haskell, whose type system could aid in the formalisation of such proofs.

Another form of reasoning about programs is static analysis. Klinik et al. [11] have developed a cost analysis for tasks that require resources in order to be executed. This analysis was developed for a simpler task language, and could be brought over to the one developed here.

Naus and Jeuring [19] have looked at building a generic feedback system for rule-based problems. A workflow system typically is rule based, as outlined in their work. It would be interesting to fit the generic feedback system to $\widehat{\text{TOP}}$ in order to support end-users working in applications developed in this language.

Additionally, we would like to develop visualisations for $\widehat{\text{TOP}}$ language constructs. An assistive development environment integrating these visualisations and the presented textual language would aid domain experts to model workflows in a more accessible manner. A system that visualises iTask programs has been developed in the past [26].

REFERENCES

- [1] Heinrich Apfelmus. 2019. reactive-banana. <https://wiki.haskell.org/Reactive-banana>. [Accessed 13-February-2019].
- [2] Gérard Berry and Georges Gonthier. 1992. The Esterel Synchronous Programming Language: Design, Semantics, Implementation. *Sci. Comput. Program.* 19, 2 (1992), 87–152. [https://doi.org/10.1016/0167-6423\(92\)90005-V](https://doi.org/10.1016/0167-6423(92)90005-V)
- [3] Gérard Berry, Cyprien Nicolas, and Manuel Serrano. 2011. HipHop: A synchronous reactive extension for Hop. In *Proceedings of the 1st ACM SIGPLAN international workshop on Programming language and systems technologies for internet clients*. ACM, 49–56.

- [4] Gérard Berry and Manuel Serrano. 2013. Hop and HipHop : Multitier Web Orchestration. *CoRR* abs/1312.0078 (2013). arXiv:1312.0078 <http://arxiv.org/abs/1312.0078>
- [5] Frédéric Boussinot and Robert De Simone. 1991. The Esterel language. *Proc. IEEE* 79, 9 (1991), 1293–1304.
- [6] Gregory Cooper and Shriram Krishnamurthi. 2004. *FrTime: Functional Reactive Programming in PLT Scheme*. Technical Report CS-03-20. Department of Computer Science, Brown University, Rhode Island.
- [7] Allen B. Downey. 2008. *The Little Book of Semaphores*. Green Tea Press.
- [8] Conal Elliott and Paul Hudak. 1997. Functional Reactive Animation. In *Proceedings of the 1997 ACM SIGPLAN International Conference on Functional Programming (ICFP '97), Amsterdam, The Netherlands, June 9-11, 1997*. 263–273.
- [9] C. A. R. Hoare. 1985. *Communicating Sequential Processes*. Prentice Hall International.
- [10] Mauro Jaskieloff, Neil Ghani, and Graham Hutton. 2011. Modularity and Implementation of Mathematical Operational Semantics. *Electr. Notes Theor. Comput. Sci.* 229, 5 (2011), 75–95. <https://doi.org/10.1016/j.entcs.2011.02.017>
- [11] Markus Klinik, Jan Martin Jansen, and Rinus Plasmeijer. 2017. The Sky is the Limit: Analysing Resource Consumption Over Time Using Skylines. In *Proceedings of the 29th Symposium on Implementation and Application of Functional Programming Languages, IFL 2017*. ACM.
- [12] Bram Kool. 2017. Integrated Mission Management voor C2-ondersteuning. Bachelor's Thesis. Dutch Defence Academy, Den Helder, The Netherlands.
- [13] Pieter Koopman, Mart Lubbers, and Rinus Plasmeijer. 2018. A task-based dsl for microcomputers. In *Proceedings of the Real World Domain Specific Languages Workshop*. 4.
- [14] Pieter W. M. Koopman, Rinus Plasmeijer, and Peter Achten. 2008. An Executable and Testable Semantics for iTasks. In *Proceedings of the 20th Symposium on Implementation and Application of Functional Programming Languages, IFL 2008*.
- [15] Bas Lijnse, Jan Martin Jansen, and Rinus Plasmeijer. 2012. Incidone: A Task-Oriented Incident Coordination Tool. In *Proceedings of ISCRAM*.
- [16] Simon Marlow et al. 2010. *Haskell 2010 language report*.
- [17] Leo A. Meyerovich, Arjun Guha, Jacob P. Baskin, Gregory H. Cooper, Michael Greenberg, Aleks Bromfield, and Shriram Krishnamurthi. 2009. Flapjax: a programming language for Ajax applications. In *Proceedings of the 24th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2009, October 25-29, 2009, Orlando, Florida, USA*. 1–20.
- [18] Robin Milner. 1989. *Communication and concurrency*. Prentice Hall.
- [19] Nico Naus and Johan Jeuring. 2016. Building a Generic Feedback System for Rule-Based Problems. In *Trends in Functional Programming - 17th International Conference, TFP 2016, College Park, MD, USA, June 8-10, 2016, Revised Selected Papers (Lecture Notes in Computer Science)*, David Van Horn and John Hughes (Eds.), Vol. 10447. Springer, 172–191. https://doi.org/10.1007/978-3-030-14805-8_10
- [20] OASIS. 2019. Web Services Business Process Execution Language. https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wsbpel. [Accessed 12-February-2019].
- [21] Simon Peyton Jones. 2001. Tackling the awkward squad: monadic input/output, concurrency, exceptions, and foreign-language calls in Haskell. In *Engineering theories of software construction, Marktoberdorf Summer School 2000, Marktoberdorf, Germany*.
- [22] Benjamin C. Pierce. 2002. *Types and programming languages*. MIT Press.
- [23] Rinus Plasmeijer, Bas Lijnse, Steffen Michels, Peter Achten, and Pieter W. M. Koopman. 2012. Task-oriented programming in a pure functional language. In *Principles and Practice of Declarative Programming, PPDP'12, Leuven, Belgium, 2012*.
- [24] Jurriën Stutterheim. 2017. *A Cocktail of Tools*. Ph.D. Dissertation. Radboud University, Nijmegen, The Netherlands.
- [25] Jurriën Stutterheim, Peter Achten, and Rinus Plasmeijer. 2017. Maintaining Separation of Concerns Through Task Oriented Software Development. In *Trends in Functional Programming - 18th International Symposium, TFP 2017, Canterbury, UK*.
- [26] Jurriën Stutterheim, Rinus Plasmeijer, and Peter Achten. 2014. Tonic: An Infrastructure to Graphically Represent the Definition and Behaviour of Tasks. In *Trends in Functional Programming - 15th International Symposium, TFP 2014, Soesterberg, The Netherlands, May 26-28, 2014, Revised Selected Papers (Lecture Notes in Computer Science)*, Jurriaan Hage and Jay McCarthy (Eds.), Vol. 8843. Springer, 122–141. https://doi.org/10.1007/978-3-319-14675-1_8
- [27] Wouter Swierstra. 2008. Data types à la carte. *J. Funct. Program.* 18, 4 (2008), 423–436. <https://doi.org/10.1017/S0956796808006758>
- [28] Wil M. P. van der Aalst. 1998. The Application of Petri Nets to Workflow Management. *Journal of Circuits, Systems, and Computers* 8, 1 (1998), 21–66. <https://doi.org/10.1142/S0218126698000043>
- [29] Wil M. P. van der Aalst and Arthur H. M. ter Hofstede. 2005. YAWL: yet another workflow language. *Inf. Syst.* 30, 4 (2005), 245–275. <https://doi.org/10.1016/j.is.2004.02.002>
- [30] Wil M. P. van der Aalst, Arthur H. M. ter Hofstede, Bartek Kiepuszewski, and Alistair P. Barros. 2003. Workflow Patterns. *Distributed and Parallel Databases* 14, 1 (2003), 5–51. <https://doi.org/10.1023/A:1022883727209>

A ADDITIONAL RULES

A.1 Evaluation rules

$$\boxed{e, \sigma \downarrow v, \sigma'}$$

$$\text{E-APP} \frac{e_1, \sigma \downarrow \lambda x : \tau. e'_1, \sigma' \quad e_2, \sigma' \downarrow v_2, \sigma'' \quad e'_1[x \mapsto v_2], \sigma'' \downarrow v_1, \sigma'''}{e_1 e_2, \sigma \downarrow v_1, \sigma'''}$$

$$\text{E-IFTRUE} \frac{e_1, \sigma \downarrow \text{True}, \sigma' \quad e_2, \sigma' \downarrow v, \sigma''}{\text{if } e_1 \text{ then } e_2 \text{ else } e_3, \sigma \downarrow v, \sigma''} \quad \text{E-REF} \frac{e, \sigma \downarrow v, \sigma' \quad l \notin \text{Dom}(\sigma')}{\text{ref } e, \sigma \downarrow l, \sigma'[l \mapsto v]}$$

$$\text{E-IFFALSE} \frac{e_1, \sigma \downarrow \text{False}, \sigma' \quad e_3, \sigma' \downarrow v, \sigma''}{\text{if } e_1 \text{ then } e_2 \text{ else } e_3, \sigma \downarrow v, \sigma''} \quad \text{E-DEREF} \frac{e, \sigma \downarrow l, \sigma' \quad !e, \sigma \downarrow \sigma'(l), \sigma'}{!e, \sigma \downarrow \sigma'(l), \sigma'}$$

$$\text{E-VALUE} \frac{}{v, \sigma \downarrow v, \sigma} \quad \text{E-ASSIGN} \frac{e_1, \sigma \downarrow l, \sigma' \quad e_2, \sigma' \downarrow v_2, \sigma''}{e_1 := e_2, \sigma \downarrow \langle \rangle, \sigma''[l \mapsto v_2]}$$

$$\text{E-PAIR} \frac{e_1, \sigma \downarrow v_1, \sigma' \quad e_2, \sigma' \downarrow v_2, \sigma''}{\langle e_1, e_2 \rangle, \sigma \downarrow \langle v_1, v_2 \rangle, \sigma''}$$

$$\text{E-EDIT} \frac{e, \sigma \downarrow v, \sigma'}{\Box e, \sigma \downarrow \Box v, \sigma'} \quad \text{E-ENTER} \frac{}{\Box \tau, \sigma \downarrow \Box \tau, \sigma} \quad \text{E-UPDATE} \frac{e, \sigma \downarrow l, \sigma'}{\blacksquare e, \sigma \downarrow \blacksquare l, \sigma'}$$

$$\text{E-THEN} \frac{e_1, \sigma \downarrow t_1, \sigma'}{e_1 \blacktriangleright e_2, \sigma \downarrow t_1 \blacktriangleright e_2, \sigma'} \quad \text{E-NEXT} \frac{e_1, \sigma \downarrow t_1, \sigma'}{e_1 \triangleright e_2, \sigma \downarrow t_1 \triangleright e_2, \sigma'}$$

$$\text{E-AND} \frac{e_1, \sigma \downarrow t_1, \sigma' \quad e_2, \sigma' \downarrow t_2, \sigma''}{e_1 \bowtie e_2, \sigma \downarrow t_1 \bowtie t_2, \sigma''} \quad \text{E-FAIL} \frac{}{\frac{1}{2}, \sigma \downarrow \frac{1}{2}, \sigma}$$

$$\text{E-OR} \frac{e_1, \sigma \downarrow t_1, \sigma' \quad e_2, \sigma' \downarrow t_2, \sigma''}{e_1 \blacklozenge e_2, \sigma \downarrow t_1 \blacklozenge t_2, \sigma''} \quad \text{E-XOR} \frac{}{e_1 \diamond e_2, \sigma \downarrow e_1 \diamond e_2, \sigma}$$

A.2 Typing rules

$$\boxed{\Gamma, \Sigma \vdash e : \tau}$$

$$\text{T-VAR} \frac{x : \tau \in \Gamma}{\Gamma, \Sigma \vdash x : \tau} \quad \text{T-LOC} \frac{\Sigma(l) = \beta}{\Gamma, \Sigma \vdash l : \text{REF } \beta} \quad \text{T-PAIR} \frac{\Gamma, \Sigma \vdash e_1 : \tau_1 \quad \Gamma, \Sigma \vdash e_2 : \tau_2}{\Gamma, \Sigma \vdash \langle e_1, e_2 \rangle : \tau_1 \times \tau_2}$$

$$\text{T-ABS} \frac{\Gamma[x : \tau_1], \Sigma \vdash e : \tau_2}{\Gamma, \Sigma \vdash \lambda x : \tau_1. e : \tau_1 \rightarrow \tau_2} \quad \text{T-APP} \frac{\Gamma, \Sigma \vdash e_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma, \Sigma \vdash e_2 : \tau_1}{\Gamma, \Sigma \vdash e_1 e_2 : \tau_2}$$

$$\text{T-IF} \frac{\Gamma, \Sigma \vdash e_1 : \text{BOOL} \quad \Gamma, \Sigma \vdash e_2 : \tau \quad \Gamma, \Sigma \vdash e_3 : \tau}{\Gamma, \Sigma \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : \tau} \quad \text{T-REF} \frac{\Gamma, \Sigma \vdash e : \beta}{\Gamma, \Sigma \vdash \text{ref } e : \text{REF } \beta}$$

$$\text{T-DEREF} \frac{\Gamma, \Sigma \vdash e : \text{REF } \beta}{\Gamma, \Sigma \vdash !e : \beta} \quad \text{T-ASSIGN} \frac{\Gamma, \Sigma \vdash e_1 : \text{REF } \beta \quad \Gamma, \Sigma \vdash e_2 : \beta}{\Gamma, \Sigma \vdash e_1 := e_2 : \text{UNIT}}$$

B PROOFS

B.1 Theorem 6.1

PROOF. We prove Theorem 6.1 by induction on e :

Case $e = \lambda x : \tau. e_1 e_2, x, c, l, e_1 \star e_2, \text{if } e_1 \text{ then } e_2 \text{ else } e_3, \langle e_1, e_2 \rangle, \langle \rangle, \text{ref } e, !e, e_1 := e_2$

Preservation has been proven for these cases by Pierce [22].

$$\text{E-EDIT} \frac{e, \sigma \downarrow v, \sigma'}{\Box e, \sigma \downarrow \Box v, \sigma'}$$

Given that $\Gamma, \Sigma \vdash \Box e : \text{TASK } \tau$ and $\Sigma \vdash s, \text{T-EDIT}$ gives us that $\Gamma, \Sigma \vdash e : \tau$. The induction hypothesis gives us that $e, s \downarrow v, s'$ also preserves, and thus $\Gamma, \Sigma \vdash v : \tau$ and $\Sigma \vdash s'$. Therefore $\Gamma, \Sigma \vdash \Box v : \text{TASK } \tau$.

$$\text{E-ENTER} \frac{}{\Box \tau, \sigma \downarrow \Box \tau, \sigma}$$

Evaluation does not alter e and s , therefore this case holds trivially.

$$\text{E-UPDATE} \frac{e, \sigma \downarrow l, \sigma'}{\blacksquare e, \sigma \downarrow \blacksquare l, \sigma'}$$

Given that $\Gamma, \Sigma \vdash \Box e : \text{TASK } \tau$ and $\Sigma \vdash s, \text{T-UPDATE}$ gives us that $\Gamma, \Sigma \vdash e : \text{ref } \tau$. The induction hypothesis gives us that $e, s \downarrow l, s'$ also preserves, and thus $\Gamma, \Sigma \vdash l : \text{ref } \tau$ and $\Sigma \vdash s'$. Therefore $\Gamma, \Sigma \vdash \blacksquare l : \text{TASK } \tau$.

$$\text{E-FAIL} \frac{}{\frac{1}{2}, \sigma \downarrow \frac{1}{2}, \sigma}$$

Evaluation does not alter e and s , therefore this case holds trivially.

$$\text{E-THEN} \frac{e_1, \sigma \downarrow t_1, \sigma'}{e_1 \blacktriangleright e_2, \sigma \downarrow t_1 \blacktriangleright e_2, \sigma'}$$

Given that $\Gamma, \Sigma \vdash e_1 \blacktriangleright e_2 : \text{TASK } \tau$ and $\Sigma \vdash s, \text{T-THEN}$ gives us that $\Gamma, \Sigma \vdash e_1 : \text{TASK } \tau_1$ and $\Gamma, \Sigma \vdash e_2 : \tau_1 \rightarrow \text{TASK } \tau$. By the induction hypothesis, we know that $e_1, s \downarrow t_1, s'$ preserves and thus $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau_1$ and $\Sigma \vdash s'$. Therefore $\Gamma, \Sigma \vdash t_1 \blacktriangleright e_2 : \text{TASK } \tau$.

$$\text{E-NEXT} \frac{e_1, \sigma \downarrow t_1, \sigma'}{e_1 \triangleright e_2, \sigma \downarrow t_1 \triangleright e_2, \sigma'}$$

Given that $\Gamma, \Sigma \vdash e_1 \triangleright e_2 : \text{TASK } \tau$ and $\Sigma \vdash s, \text{T-NEXT}$ gives us that $\Gamma, \Sigma \vdash e_1 : \text{TASK } \tau_1$ and $\Gamma, \Sigma \vdash e_2 : \tau_1 \rightarrow \text{TASK } \tau$. By the induction hypothesis, we know that $e_1, s \downarrow t_1, s'$ preserves and thus $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau_1$ and $\Sigma \vdash s'$. Therefore $\Gamma, \Sigma \vdash t_1 \triangleright e_2 : \text{TASK } \tau$.

$$\text{E-AND} \frac{e_1, \sigma \downarrow t_1, \sigma' \quad e_2, \sigma' \downarrow t_2, \sigma''}{e_1 \bowtie e_2, \sigma \downarrow t_1 \bowtie t_2, \sigma''}$$

Given that $\Gamma, \Sigma \vdash e_1 \bowtie e_2 : \text{TASK}(\tau_1 \times \tau_2)$ and $\Sigma \vdash s, \text{T-AND}$ gives us that $\Gamma, \Sigma \vdash e_1 : \text{TASK } \tau_1$ and $\Gamma, \Sigma \vdash e_2 : \text{TASK } \tau_2$. By the induction hypothesis, we know that both $e_1, s \downarrow t_1, s'$ and $e_2, s' \downarrow t_2, s''$ preserve and thus $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau_1, \Sigma \vdash s', \Gamma, \Sigma \vdash t_2 : \text{TASK } \tau_2$ and $\Sigma \vdash s''$. Therefore $\Gamma, \Sigma \vdash t_1 \bowtie t_2 : \text{TASK}(\tau_1 \times \tau_2)$.

Case
$$\frac{\text{E-OR} \quad e_1, \sigma \downarrow t_1, \sigma' \quad e_2, \sigma' \downarrow t_2, \sigma''}{e_1 \blacklozenge e_2, \sigma \downarrow t_1 \blacklozenge t_2, \sigma''}$$

Given that $\Gamma, \Sigma \vdash e_1 \blacklozenge e_2 : \text{TASK } \tau$ and $\Sigma \vdash s$, T-OR gives us that $\Gamma, \Sigma \vdash e_1 : \text{TASK } \tau$ and $\Gamma, \Sigma \vdash e_2 : \text{TASK } \tau$. By the induction hypothesis, we have that both $e_1, s \downarrow t_1, s'$ and $e_2, s' \downarrow t_2, s''$ preserve and thus $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau, \Sigma \vdash s', \Gamma, \Sigma \vdash t_2 : \text{TASK } \tau$ and $\Sigma \vdash s''$. Therefore $\Gamma, \Sigma \vdash t_1 \blacklozenge t_2 : \text{TASK } \tau$.

Case
$$\frac{\text{E-XOR}}{e_1 \diamond e_2, \sigma \downarrow e_1 \diamond e_2, \sigma}$$

Evaluation does not alter e and s , therefore this case holds trivially. $\square$

B.2 Lemma B.1

LEMMA B.1 (TASK VALUE PRESERVES TYPES). *For all expressions e and states σ such that $\Gamma, \Sigma \vdash e : \text{TASK } \tau$ and $\Sigma \vdash \sigma$, if $\mathcal{V}(e, \sigma) = v$, then $v : \tau$.*

PROOF. We prove Lemma B.1 by induction over e .

Case $\mathcal{V}(\square v, s) = v$

By T-Edit, if $\Gamma, \Sigma \vdash \square v : \text{TASK } \tau$, then $\Gamma, \Sigma \vdash v : \tau$.

Case $\mathcal{V}(\boxtimes \tau, s) = \perp$

Since this case does not lead to a value, the lemma holds trivially.

Case $\mathcal{V}(\blacksquare l, s) = s(l)$

Given that $\Gamma, \Sigma \vdash \blacksquare l : \text{TASK } \tau$ and $\Sigma \vdash s$, we know that $\Gamma, \Sigma \vdash s(l) : \tau$ by definiton.

Case $\mathcal{V}(\frac{1}{2}, s) = \perp$

Since this case does not lead to a value, the lemma holds trivially.

Case $\mathcal{V}(t_1 \blacktriangleright e_2, s) = \perp$

Since this case does not lead to a value, the lemma holds trivially.

Case $\mathcal{V}(t_2 \triangleright e_2, s) = \perp$

Since this case does not lead to a value, the lemma holds trivially.

Case $\mathcal{V}(t_1 \bowtie t_2, s) = \langle v_1, v_2 \rangle$ given that $\mathcal{V}(t_1, s) = v_1 \wedge \mathcal{V}(t_2, s) = v_2$

By T-AND we have that $\Gamma, \Sigma \vdash t_1 \bowtie t_2 : \text{TASK}(\tau_1 \times \tau_2)$ and $\Gamma, \Sigma \vdash t_1 : \tau_1$ and $\Gamma, \Sigma \vdash t_2 : \tau_2$. By the induction hypothesis, $\mathcal{V}(t_1, s) = v_1$ and $\mathcal{V}(t_2, s) = v_2$ preserve, and thus $\Gamma, \Sigma \vdash v_1 : \tau_1$ and $\Gamma, \Sigma \vdash v_2 : \tau_2$. This gives us that $\Gamma, \Sigma \vdash \langle v_1, v_2 \rangle : \text{TASK}(\tau_1 \times \tau_2)$.

Case $\mathcal{V}(t_1 \bowtie t_2, s) = \perp$ given that $\neg(\mathcal{V}(t_1, s) = v_1 \wedge \mathcal{V}(t_2, s) = v_2)$

Since this case does not lead to a value, the lemma holds trivially.

Case $\mathcal{V}(t_1 \blacklozenge t_2, s) = v_1$ given that $\mathcal{V}(t_1, s) = v_1$

By T-OR we have that $\Gamma, \Sigma \vdash t_1 \blacklozenge t_2 : \text{TASK } \tau$, and $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau$ and $\Gamma, \Sigma \vdash t_2 : \text{TASK } \tau$. By the induction hypothesis, we have that $\Gamma, \Sigma \vdash v_1 : \tau$.

Case $\mathcal{V}(t_1 \blacklozenge t_2, s) = v_2$ given that $\mathcal{V}(t_1, s) = \perp \wedge \mathcal{V}(t_2, s) = v_2$

By T-OR we have that $\Gamma, \Sigma \vdash t_1 \blacklozenge t_2 : \text{TASK } \tau$, and $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau$ and $\Gamma, \Sigma \vdash t_2 : \text{TASK } \tau$. By the induction hypothesis, we have that $\Gamma, \Sigma \vdash v_2 : \tau$.

Case $\mathcal{V}(t_1 \blacklozenge t_2, s) = \perp$ given that $\mathcal{V}(t_1, s) = \perp \wedge \mathcal{V}(t_2, s) = \perp$

Since this case does not lead to a value, the lemma holds trivially.

Case $\mathcal{V}(t_1 \diamond t_2, s) = \perp$

Since this case does not lead to a value, the lemma holds trivially. $\square$

B.3 Lemma B.2

LEMMA B.2 (STRIDING PRESERVES TYPES). *For all expressions e and states σ such that $\Gamma, \Sigma \vdash e : \text{TASK } \tau$ and $\Sigma \vdash \sigma$, if $e, \sigma \rightsquigarrow e', \sigma'$, then $\Gamma, \Sigma \vdash e' : \text{TASK } \tau$ and $\Sigma \vdash \sigma'$.*

PROOF. We prove Lemma B.2 by induction on e :

Case
$$\frac{\text{S-FAIL}}{\frac{1}{2}, \sigma \rightsquigarrow \frac{1}{2}, \sigma}$$

Since this case does not alter the expression, the theorem holds trivially.

Case
$$\frac{\text{S-XOR}}{e_1 \diamond e_2, \sigma \rightsquigarrow e_1 \diamond e_2, \sigma}$$

Since this case does not alter the expression, the theorem holds trivially.

Case
$$\frac{\text{S-UPDATE}}{\blacksquare l, \sigma \rightsquigarrow \blacksquare l, \sigma}$$

Since this case does not alter the expression, the theorem holds trivially.

Case
$$\frac{\text{S-FILL}}{\boxtimes \tau, \sigma \rightsquigarrow \boxtimes \tau, \sigma}$$

Since this case does not alter the expression, the theorem holds trivially.

Case
$$\frac{\text{S-EDIT}}{\square v, \sigma \rightsquigarrow \square v, \sigma}$$

Since this case does not alter the expression, the theorem holds trivially.

Case
$$\frac{\text{S-AND}}{t_1, \sigma \rightsquigarrow t'_1, \sigma' \quad t_2, \sigma' \rightsquigarrow t'_2, \sigma''}{t_1 \bowtie t_2, \sigma \rightsquigarrow t'_1 \bowtie t'_2, \sigma''}$$

Given that $\Gamma, \Sigma \vdash t_1 \bowtie t_2 : \text{TASK}(\tau_1 \times \tau_2)$, by T-AND we have $\Gamma, \Sigma \vdash t_1 : \tau_1$ and $\Gamma, \Sigma \vdash t_2 : \tau_2$. By the induction hypothesis, we also have $\Gamma, \Sigma \vdash t'_1 : \tau_1$ and $\Gamma, \Sigma \vdash t'_2 : \tau_2$. This gives us that $\Gamma, \Sigma \vdash t'_1 \bowtie t'_2 : \text{TASK}(\tau_1 \times \tau_2)$.

S-NEXT
Case $\frac{t_1, \sigma \rightsquigarrow t'_1, \sigma'}{t_1 \triangleright e_2, \sigma \rightsquigarrow t'_1 \triangleright e_2, \sigma'}$
 Given that $\Gamma, \Sigma \vdash e_1 \triangleright e_2 : \text{TASK } \tau$, T-THEN gives us that $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau_1$ and $\Gamma, \Sigma \vdash e_2 : \tau_1 \rightarrow \text{TASK } \tau$. By the induction hypothesis, we know that $t_1 \rightsquigarrow t'_1$ preserves and thus $\Gamma, \Sigma \vdash t'_1 : \text{TASK } \tau_1$. Therefore $\Gamma, \Sigma \vdash t'_1 \triangleright e_2 : \text{TASK } \tau$.

S-ORLEFT
Case $\frac{t_1, \sigma \rightsquigarrow t'_1, \sigma'}{t_1 \blacklozenge t_2, \sigma \rightsquigarrow t'_1, \sigma'} \mathcal{V}(t'_1, \sigma') = v_1$
 Given that $\Gamma, \Sigma \vdash t_1 \blacklozenge t_2 : \text{TASK } \tau$, by T-OR we have $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau$. By the induction hypothesis, we know that $t_1 \rightsquigarrow t'_1$ preserves and thus $\Gamma, \Sigma \vdash t'_1 : \text{TASK } \tau$.

S-ORRIGHT
Case $\frac{t_1, \sigma \rightsquigarrow t'_1, \sigma' \quad t_2, \sigma' \rightsquigarrow t'_2, \sigma''}{t_1 \blacklozenge t_2, \sigma \rightsquigarrow t'_1, \sigma' \wedge t'_2, \sigma''} \mathcal{V}(t'_1, \sigma') = \perp \wedge \mathcal{V}(t'_2, \sigma'') = v_2$
 Given that $\Gamma, \Sigma \vdash t_1 \blacklozenge t_2 : \text{TASK } \tau$, by T-OR we have $\Gamma, \Sigma \vdash t_2 : \text{TASK } \tau$. By the induction hypothesis, we know that $t_2 \rightsquigarrow t'_2$ preserves and thus $\Gamma, \Sigma \vdash t'_2 : \text{TASK } \tau$.

S-ORNONE
Case $\frac{t_1, \sigma \rightsquigarrow t'_1, \sigma' \quad t_2, \sigma' \rightsquigarrow t'_2, \sigma''}{t_1 \blacklozenge t_2, \sigma \rightsquigarrow t'_1 \blacklozenge t'_2, \sigma''} \mathcal{V}(t'_1, \sigma') = \perp \wedge \mathcal{V}(t'_2, \sigma'') = \perp$
 Given that $\Gamma, \Sigma \vdash t_1 \blacklozenge t_2 : \text{TASK } \tau$, by T-OR we have $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau$ and $\Gamma, \Sigma \vdash t_2 : \text{TASK } \tau$. By the induction hypothesis, we know that $t_1 \rightsquigarrow t'_1$ and $t_2 \rightsquigarrow t'_2$ preserve, and thus $\Gamma, \Sigma \vdash t'_1 \blacklozenge t'_2 : \text{TASK } \tau$.

S-THENSTAY
Case $\frac{t_1, \sigma \rightsquigarrow t'_1, \sigma'}{t_1 \blacktriangleright e_2, \sigma \rightsquigarrow t'_1 \blacktriangleright e_2, \sigma'} \mathcal{V}(t'_1, \sigma') = \perp$
 Given that $\Gamma, \Sigma \vdash t_1 \blacktriangleright e_2 : \text{TASK } \tau$, by T-THEN we have $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau_1$ and $\Gamma, \Sigma \vdash e_2 : \tau_1 \rightarrow \text{TASK } \tau$. By the induction hypothesis, we know that $t_1 \rightsquigarrow t'_1$ preserves, and thus $\Gamma, \Sigma \vdash t'_1 \blacktriangleright e_2 : \text{TASK } \tau$.

S-THENFAIL
Case $\frac{t_1, \sigma \rightsquigarrow t'_1, \sigma' \quad e_2 v_1, \sigma' \downarrow t_2, \sigma''}{t_1 \blacktriangleright e_2, \sigma \rightsquigarrow t'_1 \blacktriangleright e_2, \sigma'} \mathcal{V}(t'_1, \sigma') = v_1 \wedge \mathcal{F}(t_2, \sigma'')$
 Given that $\Gamma, \Sigma \vdash t_1 \blacktriangleright e_2 : \text{TASK } \tau$, by T-THEN we have $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau_1$ and $e_2 : \tau_1 \rightarrow \text{TASK } \tau$. By the induction hypothesis, we know that $t_1 \rightsquigarrow t'_1$ preserves, and thus $\Gamma, \Sigma \vdash t'_1 \blacktriangleright e_2 : \text{TASK } \tau$.

S-THENCONT
Case $\frac{t_1, \sigma \rightsquigarrow t'_1, \sigma' \quad e_2 v_1, \sigma' \downarrow t_2, \sigma''}{t_1 \blacktriangleright e_2, \sigma \rightsquigarrow t'_1 \blacktriangleright e_2, \sigma''} \mathcal{V}(t'_1, \sigma') = v_1 \wedge \neg \mathcal{F}(t_2, \sigma'')$
 Given that $\Gamma, \Sigma \vdash t_1 \blacktriangleright e_2 : \text{TASK } \tau$, by T-THEN we have $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau_1$ and $\Gamma, \Sigma \vdash e_2 : \tau_1 \rightarrow \text{TASK } \tau$. By the induction hypothesis, we know that $t_1 \rightsquigarrow t'_1$ preserves. By Lemma B.1, we know that $\mathcal{V}(t'_1) = v_1$ preserves. By Theorem 6.1 we know that $e_2 v_1 \downarrow t_2$ preserves. And finally by the induction hypothesis, we know that $t_2 \rightsquigarrow t'_2$ preserves. Therefore $\Gamma, \Sigma \vdash t'_2 : \text{TASK } \tau$. $\square$

B.4 Theorem 6.2

PROOF. We prove Theorem 6.2 by induction on e :

N-DONE
Case $\frac{e, \sigma \downarrow t, \sigma' \quad t, \sigma' \rightsquigarrow t', \sigma''}{e, \sigma \Downarrow t', \sigma''} \sigma' = \sigma'' \wedge t = t'$

Given that $\Gamma, \Sigma \vdash e : \text{TASK } \tau$ and $\Sigma \vdash s$, we know that $\Gamma, \Sigma \vdash t : \text{TASK } \tau$ and $\Sigma \vdash s'$ by Theorem 6.1. Then by Lemma B.2, we have $\Gamma, \Sigma \vdash t' : \text{TASK } \tau$ and $\Sigma \vdash s''$.

N-REPEAT
Case $\frac{e, \sigma \downarrow t, \sigma' \quad t, \sigma' \rightsquigarrow t', \sigma'' \quad t', \sigma'' \Downarrow t'', \sigma'''}{e, \sigma \Downarrow t'', \sigma'''} \sigma' \neq \sigma'' \vee t \neq t''$

Given that $\Gamma, \Sigma \vdash e : \text{TASK } \tau$ and $\Sigma \vdash s$, we know that $\Gamma, \Sigma \vdash t : \text{TASK } \tau$ and $\Sigma \vdash s'$ by Theorem 6.1. Then by Lemma B.2, we have $\Gamma, \Sigma \vdash t' : \text{TASK } \tau$ and $\Sigma \vdash s''$. Then by the induction hypothesis, we finally obtain that $\Gamma, \Sigma \vdash t'' : \text{TASK } \tau$ and $\Sigma \vdash s'''$. $\square$

B.5 Theorem 6.3

We require the following Lemma for this proof.

LEMMA B.3. *Given that $\Sigma \vdash s$, $\Sigma(l) = \tau$ and $\Gamma, \Sigma \vdash v : \tau$, it holds that $\Sigma \vdash s[l \mapsto v]$*

This lemma follows immediately from definition.

PROOF. We prove Theorem 6.3 by induction on e :

H-CHANGE
Case $\frac{}{\Box v, \sigma \xrightarrow{v'} \Box v', \sigma} v, v' : \tau$

Given that $\Gamma, \Sigma \vdash \Box v : \text{TASK } \tau$ and $\Sigma \vdash s$, the H-CHANGE rule additionally gives us that $v, v' : \tau$. Therefore by T-EDIT we have that $\Gamma, \Sigma \vdash \Box v' : \text{TASK } \tau$.

H-EMPTY
Case $\frac{}{\Box v, \sigma \xrightarrow{E} \Box \tau, \sigma} v : \tau$

Given that $\Gamma, \Sigma \vdash \Box v : \text{TASK } \beta$ and $\Sigma \vdash s$, the H-EMPTY rule additionally gives us that $v : \tau$. Then by T-ENTER we have $\Gamma, \Sigma \vdash \Box \tau : \text{TASK } \tau$.

H-FILL
Case $\frac{}{\Box \tau, \sigma \xrightarrow{v'} \Box v', \sigma} v' : \tau$

Given that $\Gamma, \Sigma \vdash \Box \tau$ and $\Sigma \vdash s$, the H-FILL rule additionally gives us that $v' : \tau$. Then by T-ENTER we have $\Gamma, \Sigma \vdash \Box v' : \text{TASK } \tau$.

H-UPDATE
Case $\frac{}{\blacksquare l, \sigma \xrightarrow{v'} \blacksquare l, \sigma[l \mapsto v']} \sigma(l), v' : \tau$

Given that $\Gamma, \Sigma \vdash \blacksquare l : \text{TASK } \tau$ and $\Sigma \vdash s$. This gives us that $\Sigma(l) = \tau$, and we additionally obtain $s(l), v' : \tau$ by H-UPDATE. By application of Lemma B.3 this case holds.

Case $\frac{\text{H-PICKLEFT}}{e_1, \sigma \Downarrow t_1, \sigma'} \neg \mathcal{F}(t_1, \sigma')$
 $e_1 \diamond e_2, \sigma \xrightarrow{L} t_1, \sigma'$
 Given that $\Gamma, \Sigma \vdash t_1 \diamond t_2 : \text{TASK } \tau$ and $\Sigma \vdash s$, then by T-XOR we have $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau$.

Case $\frac{\text{H-PICKRIGHT}}{e_2, \sigma \Downarrow t_2, \sigma'} \neg \mathcal{F}(t_2, \sigma')$
 $e_1 \diamond e_2, \sigma \xrightarrow{R} t_2, \sigma'$
 Given that $\Gamma, \Sigma \vdash t_1 \diamond t_2 : \text{TASK } \tau$ and $\Sigma \vdash s$, then by T-XOR we have $\Gamma, \Sigma \vdash t_2 : \text{TASK } \tau$.

Case $\frac{\text{H-NEXT}}{e_2 \triangleright e_1, \sigma \Downarrow t_2, \sigma'} \mathcal{V}(t_1, \sigma) = v_1 \wedge \neg \mathcal{F}(t_2, \sigma')$
 $t_1 \triangleright e_2, \sigma \xrightarrow{C} t_2, \sigma'$
 Given that $\Gamma, \Sigma \vdash t_1 \triangleright e_2 : \text{TASK } \tau$ and $\Sigma \vdash s$. Then by T-NEXT, we have $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau_1$ and $\Gamma, \Sigma \vdash e_2 : \tau_1 \rightarrow \text{TASK } \tau$. Then by T-THEN we obtain $\Gamma, \Sigma \vdash t_1 \blacktriangleright e_2 : \text{TASK } \tau$.

Case $\frac{\text{H-PASSTHEN}}{t_1, \sigma \xrightarrow{i} t'_1, \sigma'} t_1 \blacktriangleright e_2, \sigma \xrightarrow{i} t'_1 \blacktriangleright e_2, \sigma'$
 Given that $\Gamma, \Sigma \vdash t_1 \blacktriangleright e_2 : \text{TASK } \tau$ and $\Sigma \vdash s$, T-THEN gives us that $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau_1$ and $\Gamma, \Sigma \vdash e_2 : \tau_1 \rightarrow \text{TASK } \tau$. By the induction hypothesis, we know that $t_1, s \xrightarrow{i} t'_1, s'$ also preserves and thus $\Gamma, \Sigma \vdash t'_1 : \text{TASK } \tau_1$ and $\Gamma, \Sigma \vdash s' : \tau_1 \rightarrow \text{TASK } \tau$. By T-THEN we now obtain that $\Gamma, \Sigma \vdash t'_1 \blacktriangleright e_2 : \text{TASK } \tau$.

Case $\frac{\text{H-PASSNEXT}}{t_1, \sigma \xrightarrow{i} t'_1, \sigma'} t_1 \triangleright e_2, \sigma \xrightarrow{i \neq C} t'_1 \triangleright e_2, \sigma'$
 Given that $\Gamma, \Sigma \vdash t_1 \triangleright e_2 : \text{TASK } \tau$ and $\Sigma \vdash s$, T-NEXT gives us that $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau_1$ and $\Gamma, \Sigma \vdash e_2 : \tau_1 \rightarrow \text{TASK } \tau$. By the induction hypothesis, we know that $t_1, s \xrightarrow{i} t'_1, s'$ also preserves and thus $\Gamma, \Sigma \vdash t'_1 : \text{TASK } \tau_1$ and $\Gamma, \Sigma \vdash s' : \tau_1 \rightarrow \text{TASK } \tau$. By T-NEXT we now obtain that $\Gamma, \Sigma \vdash t'_1 \triangleright e_2 : \text{TASK } \tau$.

Case $\frac{\text{H-FIRSTAND}}{t_1, \sigma \xrightarrow{i} t'_1, \sigma'} t_1 \bowtie t_2, \sigma \xrightarrow{Fi} t'_1 \bowtie t_2, \sigma'$
 Given that $\Gamma, \Sigma \vdash t_1 \bowtie t_2 : \text{TASK}(\tau_1 \times \tau_2)$ and $\Sigma \vdash s$, T-AND gives us that $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau_1$ and $\Gamma, \Sigma \vdash t_2 : \text{TASK } \tau_2$. By the induction hypothesis, we know that $t_1, s \xrightarrow{i} t'_1, s'$ also preserves and thus $\Gamma, \Sigma \vdash t'_1 : \text{TASK } \tau_1$ and $\Sigma \vdash s'$. Therefore by T-NEXT we obtain $\Gamma, \Sigma \vdash t'_1 \bowtie t_2 : \text{TASK}(\tau_1 \times \tau_2)$.

Case $\frac{\text{H-SECONDAND}}{t_2, \sigma \xrightarrow{i} t'_2, \sigma'} t_1 \bowtie t_2, \sigma \xrightarrow{Si} t_1 \bowtie t'_2, \sigma'$
 Given that $\Gamma, \Sigma \vdash t_1 \bowtie t_2 : \text{TASK}(\tau_1 \times \tau_2)$ and $\Sigma \vdash s$, T-AND gives us that $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau_1$ and $\Gamma, \Sigma \vdash t_2 : \text{TASK } \tau_2$. By the induction hypothesis, we know that $t_2, s \xrightarrow{i} t'_2, s'$ also preserves and thus

$\Gamma, \Sigma \vdash t'_2 : \text{TASK } \tau_2$ and $\Sigma \vdash s'$. Therefore by T-NEXT we obtain $\Gamma, \Sigma \vdash t_1 \bowtie t'_2 : \text{TASK}(\tau_1 \times \tau_2)$.

Case $\frac{\text{H-FIRSTOR}}{t_1, \sigma \xrightarrow{i} t'_1, \sigma'} t_1 \blacklozenge t_2, \sigma \xrightarrow{Fi} t'_1 \blacklozenge t_2, \sigma'$
 Given that $\Gamma, \Sigma \vdash t_1 \blacklozenge t_2 : \text{TASK } \tau$ and $\Sigma \vdash s$, T-OR gives us that $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau$ and $\Gamma, \Sigma \vdash t_2 : \text{TASK } \tau$. By the induction hypothesis we know that $t_1, s \xrightarrow{i} t'_1, s'$ also preserves, and therefore $\Gamma, \Sigma \vdash t'_1 : \text{TASK } \tau$ and $\Sigma \vdash s'$. By T-OR we now obtain $\Gamma, \Sigma \vdash t'_1 \blacklozenge t_2 : \text{TASK } \tau$.

Case $\frac{\text{H-SECONDOR}}{t_2, \sigma \xrightarrow{i} t'_2, \sigma'} t_1 \blacklozenge t_2, \sigma \xrightarrow{Si} t_1 \blacklozenge t'_2, \sigma'$
 Given that $\Gamma, \Sigma \vdash t_1 \blacklozenge t_2 : \text{TASK } \tau$ and $\Sigma \vdash s$, T-OR gives us that $\Gamma, \Sigma \vdash t_1 : \text{TASK } \tau$ and $\Gamma, \Sigma \vdash t_2 : \text{TASK } \tau$. By the induction hypothesis we know that $t_2, s \xrightarrow{i} t'_2, s'$ also preserves, and therefore $\Gamma, \Sigma \vdash t'_2 : \text{TASK } \tau$ and $\Sigma \vdash s'$. By T-OR we now obtain $\Gamma, \Sigma \vdash t_1 \blacklozenge t'_2 : \text{TASK } \tau$. $\square$

B.6 Theorem 6.5

PROOF. We prove Theorem 6.5 by induction on e' .

Case $e = \perp$
 $\mathcal{F}(\perp, s) = \text{True}$, and there is no handling rule that applies to fail.

Case $e = \square v$
 $\Gamma, \Sigma \vdash \square v : \text{TASK } \tau$, $\mathcal{F}(\square v, s) = \text{False}$, and there exists an input i , namely $v' : \tau$.

Case $e = \boxtimes \tau$
 $\mathcal{F}(\boxtimes \tau) = \text{False}$, and there exists an input i , namely $v : \tau$.

Case $e = \blacksquare l$
 Given that $\Gamma, \Sigma \vdash \blacksquare l : \text{TASK } \tau$, $\mathcal{F}(\blacksquare l, s) = \text{False}$, and there exists an input i , namely $v : \tau$.

Case $e = t_1 \blacktriangleright e_2$
 $\mathcal{F}(t_1 \blacktriangleright e_2, s) = \mathcal{F}(t_1, s)$. If there exists an i for t_1 , then this i also applies to $t_1 \blacktriangleright e_2$. This case therefore holds by the induction hypothesis.

Case $e = t_1 \triangleright e_2$
 $\mathcal{F}(t_1 \triangleright e_2, s) = \mathcal{F}(t_1, s)$. If there exists an i for t_1 , then this i also applies to $t_1 \triangleright e_2$. This case therefore holds by the induction hypothesis.

Case $e = e_1 \diamond e_2$
 We normalise the two expressions first, $e_1, s \rightsquigarrow t_1, s'$, $e_2, s \rightsquigarrow t_2, s'$ and we can then be in two situations. One, we can have that $\mathcal{F}(t_1, s')$ and $\mathcal{F}(t_2, s')$ are both true. If that is so, then by definition, we have both $\mathcal{F}(e_1 \diamond e_2, s)$ and no rule of the handling

semantics applies, and therefore there exists no input for this case.

Or we are in the situation where one or both of the two sub expressions does not fail. In that case, we know that $\mathcal{F}(e_1 \diamond e_2, s)$ does not hold, and that at least one of the handling rules applies. Therefore, there must be an input i , namely L, R or both.

Case $e = t_1 \bowtie t_2$

We can again find ourselves in one of two situations. In the first case, both sub expressions fail, $\mathcal{F}(t_1, s)$ and $\mathcal{F}(t_2, s)$. In that case, we know that $\mathcal{F}(t_1 \bowtie t_2, s)$ also fails by definition. By the induction hypothesis, we know that for both t_1 and t_2 there is no input that can be handled. Since the only two rules for $\bowtie$ that handle input just pass this input on to one of the two expressions, we know that indeed no i applies.

In the case that one or both sub expressions do not fail, then by definition $t_1 \bowtie t_2$ not failing under s . Again by induction hypothesis, we know that for one or both of the expressions, there exists an i that can be handled. Then by H-FirstAnd and H-SecondAnd, we know that we can pass this i , by prefixing it with either F or S.

Case $e = t_1 \blacklozenge t_2$

We can again find ourselves in one of two situations. In the first case, both sub expressions fail, $\mathcal{F}(t_1, s)$ and $\mathcal{F}(t_2, s)$. In that case, we know that $\mathcal{F}(t_1 \blacklozenge t_2, s)$ also fails by definition. By the induction hypothesis, we know that for both t_1 and t_2 there is no input that can be handled. Since the only two rules for $\blacklozenge$ that handle input just pass this input on to one of the two expressions, we know that indeed no i applies.

In the case that one or both sub expressions do not fail, then by definition $t_1 \blacklozenge t_2$ not failing under s . Again by induction hypothesis, we know that for one or both of the expressions, there exists an i that can be handled. Then by H-FirstOr and H-SecondOr, we know that we can pass this i , by prefixing it with either F or S. $\square$

B.7 Theorem 6.6

PROOF. Case $e = \square v : \text{TASK } \tau, i = v' : \tau$

Given that $\frac{\text{H-CHANGE}}{\square v, \sigma \xrightarrow{v'} \square v', \sigma}$, we have by definition that $I(\square v : \text{TASK } \tau, s) = \{v' : \tau, E\}$, which includes $v' : \tau$.

Case $e = \square v : \text{TASK } \tau, i = E$

Given that $\frac{\text{H-EMPTY}}{\square v, \sigma \xrightarrow{E} \boxtimes \tau, \sigma}$, we have by definition that $I(\square v : \text{TASK } \tau, s) = \{v' : \tau, E\}$, which includes E.

Case $e = \boxtimes \tau, i = v' : \tau$

Given that $\frac{\text{H-FILL}}{\boxtimes \tau, \sigma \xrightarrow{v'} \square v', \sigma}$, we have by definition that $I(\boxtimes \tau, s) = \{v' : \tau\}$, which includes $v' : \tau$.

Case $e = \blacksquare l : \text{TASK } \tau, i = v' : \tau$

Given that $\frac{\text{H-UPDATE}}{\blacksquare l, \sigma \xrightarrow{v'} \blacksquare l, \sigma[l \mapsto v']}$, we have by definition that $I(\blacksquare l : \text{TASK } \tau, s) = \{v' : \tau\}$, which includes $v' : \tau$.

Case $e = t_1 \diamond t_2, i = L$

Given that $\frac{\text{H-PICKLEFT}}{e_1, \sigma \Downarrow t_1, \sigma' \neg \mathcal{F}(t_1, \sigma')}$, we have by definition that $I(t_1 \diamond t_2, s) = \{L, R\}$, which includes L.

Case $e = t_1 \diamond t_2, i = R$

Given that $\frac{\text{H-PICKRIGHT}}{e_2, \sigma \Downarrow t_2, \sigma' \neg \mathcal{F}(t_2, \sigma')}$, we have by definition that $I(t_1 \diamond t_2, s) = \{L, R\}$, which includes R.

Case $e = t_1 \triangleright e_2, i = C$

Given that $\frac{\text{H-NEXT}}{e_2 v_1, \sigma \Downarrow t_2, \sigma' \mathcal{V}(t_1, \sigma) = v_1 \wedge \neg \mathcal{F}(t_2, \sigma')}$, we have by definition that $I(t_1 \triangleright e_2, s) = I(t_1, s) \cup \{C \mid \mathcal{V}(t_1, s) = v_1 \wedge \neg \mathcal{F}(e_2 v_1, s \rightsquigarrow)\}$. If the H-Next rule applies, this means that the conditions $\mathcal{V}(t_1, s) = v_1 \wedge \neg \mathcal{F}(e_2 v_1, s \rightsquigarrow)$ are fulfilled, and therefore C is contained.

Case $e = t_1 \triangleright e_2, i \neq C$

Given that $\frac{\text{H-PASSNEXT}}{t_1, \sigma \xrightarrow{i} t'_1, \sigma'}$, we have by definition that $I(t_1 \triangleright e_2, s) = I(t_1, s) \cup \{C \mid \mathcal{V}(t_1, s) = v_1 \wedge \neg \mathcal{F}(e_2 v_1, s \rightsquigarrow)\}$. By the induction hypothesis, we have that $i \in I(t_1, s)$, and by definition of I , i is therefore also included in this case.

Case $e = t_1 \blacktriangleright e_2, i$

Given that $\frac{\text{H-PASSTHEN}}{t_1, \sigma \xrightarrow{i} t'_1, \sigma'}$, we have by definition that $I(t_1 \blacktriangleright e_2, s) = I(t_1, s)$. By the induction hypothesis, we have that $i \in I(t_1, s)$, and by definition of I , i is therefore also included in this case.

Case $e = t_1 \bowtie t_2, i = F i$

Given that $\frac{\text{H-FIRSTAND}}{t_1, \sigma \xrightarrow{i} t'_1, \sigma'}$ we have by definition that $I(t_1 \bowtie t_2, s) = \{F i \mid i \in I(t_1, s)\} \cup \{S i \mid i \in I(t_2, s)\}$. By the induction hypothesis, we have that $i \in I(t_1, s)$, and by definition of I , $F i$ is therefore also included in this case.

Case $e = t_1 \bowtie t_2, i = S\ i$

H-SECONDAND

Given that $\frac{t_2, \sigma \xrightarrow{i} t'_2, \sigma'}{\quad}$ we have by definition

$$\frac{t_1 \bowtie t_2, \sigma \xrightarrow{S\ i} t_1 \bowtie t'_2, \sigma'}{\quad}$$

that $\mathcal{I}(t_1 \bowtie t_2) = \{F\ i \mid i \in \mathcal{I}(t_1, s)\} \cup \{S\ i \mid i \in \mathcal{I}(t_2, s)\}$. By the induction hypothesis, we have that $i \in \mathcal{I}(t_2, s)$, and by definition of $\mathcal{I}$, $S\ i$ is therefore also included in this case.

Case $e = t_1 \blacklozenge t_2, i = F\ i$

H-FIRSTOR

Given that $\frac{t_1, \sigma \xrightarrow{i} t'_1, \sigma'}{\quad}$ we have by definition that

$$\frac{t_1 \blacklozenge t_2, \sigma \xrightarrow{F\ i} t'_1 \blacklozenge t_2, \sigma'}{\quad}$$

$\mathcal{I}(t_1 \blacklozenge t_2, s) = \{F\ i \mid i \in \mathcal{I}(t_1, s)\} \cup \{S\ i \mid i \in \mathcal{I}(t_2, s)\}$. By the induction hypothesis, we have that $i \in \mathcal{I}(t_1, s)$, and by definition of $\mathcal{I}$, $F\ i$ is therefore also included in this case.

Case $e = t_1 \blacklozenge t_2, i = S\ i$

H-FIRSTOR

Given that $\frac{t_1, \sigma \xrightarrow{i} t'_1, \sigma'}{\quad}$ we have by definition that

$$\frac{t_1 \blacklozenge t_2, \sigma \xrightarrow{F\ i} t'_1 \blacklozenge t_2, \sigma'}{\quad}$$

$\mathcal{I}(t_1 \blacklozenge t_2, s) = \{F\ i \mid i \in \mathcal{I}(t_1, s)\} \cup \{S\ i \mid i \in \mathcal{I}(t_2, s)\}$. By the induction hypothesis, we have that $i \in \mathcal{I}(t_2, s)$, and by definition of $\mathcal{I}$, $S\ i$ is therefore also included in this case.

□