

# PURELY FUNCTIONAL DISTRIBUTED SYSTEMS PROGRAMMING

SEYED H. HAERI      PETER VAN ROY

**ABSTRACT.** Code written for distributed systems requires to handle communication between nodes in the system. Handling that communication involves mutating the state of its medium – or, at least that is the common impression. Hence, distributed systems programming is construed as effectful (aka impure) ipso facto. This work refutes that construal. We show that the above effectfulness (aka impurity) is accidental complexity.

Our  $\lambda(\text{refut})$  is a pure functional model for distributed systems programming.  $\lambda(\text{refut})$  ships *remote futures*: a built-in facility for node communication. To use remote futures, the  $\lambda(\text{refut})$  programmer does not mutate variables. Hence, purity of the programming. The payback is that  $\lambda(\text{refut})$ 's term rewriting is impure. But, so is also the semantics of the famous purely functional languages.

## 1. INTRODUCTION

Reasoning about distributed programs is difficult. That is a well-known fact. That difficulty is partly because of the relatively little programming languages (PL) research on the topic. But, it also is partly due to the unnecessary complexity conveyed by the traditional conception of distributed systems (DSs).

Reasoning about sequential programs written for a single machine, in contrast, is the bread and butter of the PL research. Many models of functional programming, for example, have received considerable and successful attention from the PL community. The mathematically rich nature of those models is the main motivation.

In particular, pure functional programming, i.e., programming with no side-effects, has been a major source of attraction to the PL community. In that setting, many interesting properties are exhibited elegantly, e.g., equational reasoning, referential transparency, and idempotence. In short, equational reasoning is the process of interpreting code by substituting expressions for their equivalents; referential transparency is correctness of replacing an expression by its value (or any other expression having the same value); and, idempotence is when a function is invariant with respect to multiple calls. Every single one of those nice properties can boost code efficiency considerably.

The traditional conception about DSs programming, however, lacks those nice properties. In that conception, side-effects are inherent in distributed programming. The argument is: No sensible distributed program can run without communication between nodes; and, such communications **always** alter the state of the communication medium; hence, the effectfulness.

In earlier work [3, 4], we raised the point that effectfulness can be relative. What a node might find effectful is not necessarily effectful to other nodes too. In other words, side-effects might not be observable to every node in the DS.

This work takes that relativism and observability one step further. Side-effects observable to the PL's compiler (or, interpreter, evaluator, etc.) need not to be observable to the nodes in the DS. This idea, on its own, is not our invention. However, to the best of our knowledge, we are the first to notice that one can leverage that idea for a pure functional model for DSs programming.

In a simple two level world, a PL can have an object level as well as a meta-level. The object level is the one at which the programmer codes. The meta-level is the one that processes the object level. Let us set those PLs that supply metaprogramming aside for a moment. A pure

PL does not offer means for state manipulation at the object level. Nevertheless, it is perfectly fine for a pure code to incur change of state at the meta-level, and, hence, be effectful to that level.

Consider the common **let**-expressions of functional PLs (FPLs). In a pure PL, the programmer cannot rebind a variable. Memory allocation is, however, inevitable for the variable. In addition, reductions might update the runtime environment’s bookkeeping of the variable. The Haskell **let**-bindings, in particular, are that way [6, 1]. Despite that, Haskell is today’s most famous pure FPL because that effectfulness to memory is only at the meta-level.

In  $\lambda(\text{refut})$  too, for node communication, the programmer does not manipulate a communication medium at the object level. Instead, communication is by a node launching a remote task at the other node, whilst having a single-assignment handle to it. That is similar to dataflow futures of Niehren et al. [7], but,  $\lambda(\text{refut})$  provides a **let** notation for that. Thus, the set of remotely bound variables varies during runtime. Furthermore, a remote task may reduce over time too – as if the variable bound to it is being reassigned. Nonetheless,  $\lambda(\text{refut})$  programming remains pure. Because, at  $\lambda(\text{refut})$ ’s object level, the set of remotely bound variables is not observable. That set is only for  $\lambda(\text{refut})$ ’s meta-level bookkeeping. In other words, the impurity of message passing does not disappear in  $\lambda(\text{refut})$ ; it only is not observable at the object level.

$\lambda(\text{refut})$  is an important improvement to the Distributed  $\lambda$ -Calculus [2] because, contrary to  $\lambda(\text{refut})$ , the latter system cannot handle user input. In addition,  $\lambda(\text{refut})$  is an improvement over the classical model of Kahn networks [5] because it can handle non-deterministic ordering of inputs from other nodes in the DS. (Cf. Examples 3.1 and 3.3, respectively.)

An important shortcoming of  $\lambda(\text{refut})$  is that it cannot model distributed programs such as client-server, in which the **order** of messages is significant. The server is expected to be FIFO: earlier requests should be handled before the later ones. Of course, clients might have different priorities. Yet, the FIFO property is expected of the server for clients of the same priority.

Here is how this paper is organised: In § 2, we present the syntax and semantics of  $\lambda(\text{refut})$ . In § 3, we provide three case-studies for  $\lambda(\text{refut})$ : Example 3.1, 3.3, and 3.6 describe a master-slave architecture, a sensor network, and a replicated key-value store, respectively.

The study of how to extend  $\lambda(\text{refut})$  to handle significance in the order of messages is future work. Ports are our first candidate for that future work. Ports are a standard medium for message passing in DSs. Each port has an associated queue to guarantee the FIFO property required, say, for a client-server scenario.

## 2. FORMALISM

**Definition 2.1.** *The  $\lambda(\text{refut})$  syntax follows.*

$p ::= g \text{ where } d$	programs	$g ::= t \mid x \Rightarrow t \mid g \parallel g$	configurations
$e ::= x \mid c \mid \lambda x.e \mid e_1 e_2 \mid f e$	expressions	$t ::= e^a$	tasks
$\mid \text{let } \{b\} \text{ in } e \mid e; e$	expressions	$b ::= x = e \mid x = t \mid b, b$	bindings
$d ::= f(\bar{x}) = \{e\} \mid d, d$	definitions		

□

A program  $p$  in  $\lambda(\text{refut})$  is a configuration  $g$  of tasks, given some definitions  $d$ . A configuration is a set of tasks (with or without a handle to) that run concurrently, denoted by “ $\parallel$ ” in the syntax. A task  $e^a$  is an expression  $e$  to run on a node  $a$ . Remote futures have handles to a task as in  $x \Rightarrow e^a$ . What is new in a remote future  $x \Rightarrow e^a$  in comparison with dataflow futures of Niehren et al [7] is  $e$  running remotely on the node  $a$ .

The syntax for expressions is routine: variables ( $x$ ), constants ( $c$ ),  $\lambda$ -abstractions ( $\lambda x.e$ ), function applications ( $e_1 e_2$ ), application of named functions on expressions ( $f e$ ), **let**-expressions (**let**  $\{b\}$  **in**  $e$ ), and expression sequencing ( $e_1; e_2$ ). Bindings of a **let**-expression can be of two kind: local bindings ( $x = e$ ) and remote bindings ( $x = e^a$ ). The latter kind gives the task of

evaluating  $e$  to the node  $a$  and keeps a handle  $x$  to it so  $e$ 's value can be used via  $x$  when ready, if at all. When  $a$  in  $x = e^a$  is the current node, the binding is local and the node annotation will be disregarded. The notation  $f(\bar{x}) = \{e\}$  is list comprehension for  $f(x_1, \dots, x_n) = \{e\}$ , for a given  $n$ . We may drop the curly braces around  $e$  in  $f(x_1, \dots, x_n) = \{e\}$  when no confusion. So may we also do for the separator commas between definitions and between bindings of a **let**-expression.

We assume a set of values ranged over by  $v, v', \dots, v_1, v_2, \dots$ . Subscripts and priming do not change the syntactic category. For example,  $x_1, x_2, \dots, x', x'', \dots$  are all variables.

We take the order of functions introduced after the **where** clause to be irrelevant. Such we also take the order of bindings in a **let** expression and that of the constituents of a configuration. Moreover, we take **let**  $\{b_1, b_2\}$  in  $e$  to be a shorthand for **let**  $\{b_1\}$  in (**let**  $\{b_2\}$  in  $e$ ). The above dismissals of ordering is an informal account of our structural congruence below.

$$\begin{aligned} g \text{ where } d_1, d_2 &\equiv g \text{ where } d_2, d_1 & \text{let } \{b_1, b_2\} \text{ in } e &\equiv \text{let } \{b_2, b_1\} \text{ in } e \\ g_1 \parallel g_2 &\equiv g_2 \parallel g_1 & \text{let } \{b_1, b_2\} \text{ in } e &\equiv \text{let } \{b_1\} \text{ in let } \{b_2\} \text{ in } e \end{aligned}$$

**Definition 2.2.** *Evaluation contexts  $C$  of  $\lambda(\text{refut})$  are of three sorts  $E, G, P$  (corresponding to  $e, g$ , and  $p$  of the  $\lambda(\text{refut})$  syntax, respectively):*

$$\begin{aligned} C &::= (E, \delta) \mid (G, \delta) \mid (P, \delta) \\ E &::= \square \mid E \ e \mid (\lambda x. e) \ E \mid f \ E \mid E; e \mid \text{let } \{b\} \text{ in } E \mid \text{let } \{x = E\} \text{ in } e \\ G &::= E^a \mid x \Rightarrow E^a \mid G \parallel g \\ P &::= G \text{ where } d \end{aligned}$$

where  $\delta$  is a mapping from variables to expressions. □

Given that  $\lambda(\text{refut})$  programs contain **where** clauses,  $\lambda(\text{refut})$  contexts carry  $\delta$  around to recall the functions defined in the **where** clause. The context  $E$  is standard call-by-value with sequencing and **let**-expressions. Notice that  $E$  has no representative for remote futures. That is because a remote future is to only reduce at its respective remote node.

We are now prepared for the  $\lambda(\text{refut})$  semantics.

**Definition 2.3.** *The reduction semantics of  $\lambda(\text{refut})$  is of the form  $C[e] \rightarrow C[e']$ , where we will leave the  $\delta$  out when invariant upon the reduction. The reduction rules follow:*

$$\begin{aligned} (P, \delta)[g \text{ where } d_1, \dots, d_n] &\rightarrow (G, \delta[f_i \mapsto \lambda \bar{x}_i. e_i]_{i=1}^n)[g] & \text{where } d_i \text{ is } f_i(\bar{x}_i) = e_i & \text{ (PRG)} \\ & & \text{and } f_i \notin \text{domain}(\delta) & \\ E[(\lambda x. e) \ v] &\rightarrow E[e[v/x]] & \text{ (AP-E)} \\ E[\text{let } \{x = v\} \text{ in } e] &\rightarrow E[e[v/x]] & \text{ (LT-E)} \\ E[v; e] &\rightarrow E[e] & \text{ (SQ)} \\ E[f] &\rightarrow E[\lambda \bar{x}. e] & \text{where } \delta(f) = \lambda \bar{x}. e & \text{ (LD-F)} \\ G[\text{let } \{x = t\} \text{ in } e] &\rightarrow G[e[y_a/x]] \parallel y_a \Rightarrow t & \text{where } y \text{ fresh and node} = a & \text{ (LT-T)} \\ G[x] \parallel x \Rightarrow v^a &\rightarrow G[v] \parallel x \Rightarrow v^a & \text{ (DF)} \end{aligned}$$

Technically, the following rule is not required. Yet, we adopt it as a shorthand for an application of (LD-F) followed by a long-enough sequence of (AP-E)s. Suppose that  $\delta(f) = \lambda \bar{x}. e$ . Then,

$$E[f \ \bar{v}] \rightarrow E[e[\bar{v}/\bar{x}]]. \quad \text{ (AP-F)} \quad \square$$

(PRG) simply dumps the functions defined at the **where** clause into  $\delta$ . The function application rule (AP-E) is routine call-by-value. The named function application rule (AP-F) is not much more complicated. The main difference lies in the comprehension notation:  $f \ \bar{v}$  abbreviates  $f \ v_1 \ \dots \ v_n$ , for some given  $n$ ; and,  $e'[\bar{v}/\bar{x}]$  abbreviates  $e'[v_1/x_1, \dots, v_n/x_n]$ , again, for some given  $n$ . According to (LT-E), a **let**-bound variable can only be substituted when it

is bound to a value (possibly after reduction). (SQ) is routine for sequencing. (LD-F) uses  $\delta$  to substitute the body of a named function for its name.

(LT-T) is less trivial.  $\lambda(\text{refut})$  maintains a metavariable **node** to represent the node in which the reduction is taking place. The notation **node** =  $a$  in (LT-T) is for “let us assume that the node at which this let-expression is running is  $a$ .” In such a case, we launch a new concurrent task  $t$  at the right node and give it a fresh handle. (That is at a node  $b$  when  $t = e^b$ .) The subscript  $a$  in  $y_a$  is to prevent variable shadowing whilst evaluating the same let-expression at different nodes. Furthermore, we choose  $y$  to be fresh to prevent shadowing upon subsequent evaluations of the let-expression. Example 3.3 makes use of both above freshness tricks. Of course, due to the refreshing of  $x$  to  $y_a$ , occurrences of the latter need to be renamed in  $e$  accordingly.

(DF) describes the situation when a dataflow variable  $x$  has a handle to a value (possibly the result of a formerly more complicated expression). In such a situation, according to (DF), the value  $x$  is bound to can be substituted for its occurrences.

Here is an important note in relation to (DF) and (LT-T). One might have come to ask whether  $\lambda(\text{refut})$  advises for too much concurrency. After all, according to (LT-T), every remote let-binding spawns a new concurrent evaluation. We find that alright because if one modifies (DF) to

$$G[x] \parallel x \Rightarrow v^a \rightarrow G[v] \parallel x \Rightarrow v^a \quad \text{when } \text{node} = b \text{ and } a \neq b$$

one can help the scheduler with putting threads of  $b$  that have references to  $x$  to sleep until the above rule is applied to them. Furthermore, the scheduler can always perform garbage collection on a handled task  $x \Rightarrow v^a$  when there no longer is a reference to  $x$  in the configuration. We drop those in Definition 2.3 for simplicity.

### 3. USING $\lambda(\text{refut})$

**Example 3.1.** Consider a master-slave scenario with a master node  $m$  and a sufficiently large set of slave nodes  $S = \{s_1, s_2, \dots\}$ . Let us assume the availability of an input function that inputs a single item from the user. Let us also assume the any metafunction defined below

$$\frac{x_i \Rightarrow v^{a_j} \text{ for } i, j \in \{1, \dots, n\}}{\text{any}(x_1, x_2, \dots, x_n) \rightarrow v}.$$

The following  $\lambda(\text{refut})$  code gives the task of computing a function  $f$  to as many slaves as requested from  $m$ . The code is done as soon as any of those slaves is done with the computation of  $f$ , if ever.

```

1 (let {k = input} in distribute f k)m
2 where
3   distribute(g, n) = let
4     {xi = (g ())si}i=1n
5   in
6     any(x1, ..., xn)
7   f() = {...}
```

Let us call the above  $\lambda(\text{refut})$  program  $p$ . Let us assume that the user enters  $k_0$  for  $k$ . Then, dropping the **where** clauses in the presentation

$p \rightarrow^*$  (PRG), (AP-E)  
 (let {k =  $k_0$ } in distribute f k)<sup>m</sup>  $\rightarrow$  (LT-E)  
 (distribute f  $k_0$ )<sup>m</sup>  $\rightarrow$  (AP-F)  
 (let {x<sub>i</sub> = (f ())<sup>s<sub>i</sub></sup>}<sub>i=1</sub><sup>k<sub>0</sub></sup> in any(x<sub>1</sub>, ..., x<sub>k<sub>0</sub></sub>))<sup>m</sup>  $\rightarrow^*$  (LT-T)  
 (any(x<sub>1</sub>, ..., x<sub>k<sub>0</sub></sub>))<sup>m</sup> || x<sub>1</sub> | => (f ())<sup>s<sub>1</sub></sup> || ... || x<sub>k<sub>0</sub></sub> | => (f ())<sup>s<sub>k<sub>0</sub></sub></sup>  $\square$

Example 3.1 demonstrates an important step forward after the Distributed  $\lambda$ -Calculus [2]:  $\lambda(\text{refut})$  can also handle user input.

**Remark 3.2.** One may argue that `input` is not pure and so is not  $p$  in Example 3.1. We would like to draw the reader's attention to the common monadic treatments of I/O, for example, in order to give that a pure interface on a single node. This paper takes such (or similar) treatments for granted to focus on purity of distribution.  $\square$

**Example 3.3.** Consider a control unit  $c$  receiving updates from sensors  $s_1, \dots, s_n$ . Initially, the sensors are configured to take samples every 1 second. Based on their updates, the duration between two sampling may change, as instructed by the control unit in response to the update. The order in which the sensors send their update is unknown. So is also the order in which they are handled. The  $\lambda(\text{refut})$  code below describes that scenario.

```

1 (sensor 1)s1 || ... || (sensor 1)sn
2 where
3   sensor(k) = let
4     s = sample k
5     u = (update s)c
6   in
7     sensor u
8   update(s) = {...}, sample(k) = {...}

```

To demonstrate a sample run, we take  $n = 2$  and call the above code  $p$ . Again, in the demonstration, we will drop the `where` clause.

$$\begin{aligned}
 p &\rightarrow^* && (\text{PRG}), (\text{AP-F}) \\
 (\text{let } s = \text{sample } 1, u = (\text{update } s)^c \text{ in sensor } u)^{s_1} || (\text{sensor } 1)^{s_2} &\rightarrow^* && (\text{AP-F}), (\text{LT-E}) \\
 (\text{let } u = (\text{update } k_1^1)^c \text{ in sensor } u)^{s_1} || (\text{sensor } 1)^{s_2} &\rightarrow && (\text{LT-T}) \\
 (\text{sensor } u_1^1)^{s_1} || (\text{sensor } 1)^{s_2} || u_1^1 &\Rightarrow (\text{update } k_1^1)^c \rightarrow && (\text{AP-F}) \\
 (\text{sensor } u_1^1)^{s_1} || (\text{sensor } 1)^{s_2} || u_1^1 &\Rightarrow (v_1^1)^c \rightarrow && (\text{DF}) \\
 (\text{sensor } v_1^1)^{s_1} || (\text{sensor } 1)^{s_2} || u_1^1 &\Rightarrow (v_1^1)^c \rightarrow && (\text{AP-F}) \\
 (\text{sensor } v_1^1)^{s_1} || (\text{let } s = \text{sample } 1, u = (\text{update } s)^c \text{ in sensor } u)^{s_2} &&& \\
 || u_1^1 &\Rightarrow (v_1^1)^c \rightarrow^* && (\text{AP-F}), (\text{LT-E}), (\text{LT-T}), (\text{AP-F}), (\text{DF}) \\
 (\text{sensor } v_1^1)^{s_1} || (\text{sensor } v_1^2)^{s_2} || u_1^1 &\Rightarrow (v_1^1)^c || u_1^2 &\Rightarrow (v_1^2)^c \rightarrow \dots
 \end{aligned}$$

Of course, the scheduler might choose a very different interleaving.  $\square$

**Remark 3.4.** In Example 3.3, according to the wiring of the (LT-T) rule, each call to `sensor` will place a new task in the configuration. The handle to that task is different from those of the previous calls and those of the other nodes. (In the sample run, we write  $u_j^i$  for the  $j^{\text{th}}$  handle of  $s_i$ .) Note also that the application `sensor u` in line 7 needs not to be  $(\text{sensor } u)^{s_i}$ , where `node` =  $s_i$ . Thanks to  $G$  in (LT-T), the node information is invariant upon the reduction. (In fact, according to Definition 2.1, replacing `sensor u` by  $(\text{sensor } u)^{s_i}$  would be ungrammatical.)  $\square$

**Remark 3.5.** Notice how, in Example 3.3, each message sent from a sensor to the control unit places a new remote task on the configuration. That is, however, not really a problem since each such task can be garbage-collected as soon as the respective (DF) is instantiated. It is easy to get the scheduler trigger that garbage collection.  $\square$

Example 3.3 is important because it shows how  $\lambda(\text{refut})$  relaxes a stringent restriction of the Kahn networks [5], whilst still remaining pure. Putting Kahn networks into the context of distribute systems, the restriction would be that each node knows which node will provide its next input. That knowledge is not provided to the node  $c$  in Example 3.3, for instance.

Example 3.3 also develops over Example 3.1 by showing how, in  $\lambda(\text{refut})$  nodes can pass messages back and forth. Each sensor sends its  $s$  to the control unit and receives  $u$  back as the updated  $k$ . All that message passing is purely functional.

**Example 3.6.** Consider a CRDT<sup>1</sup> key-value store with replicas  $r_1, \dots, r_n$  and a front-end node  $f$ . Assume that the value type here is a semilattice. The  $\lambda(\text{refut})$  code below provides a possible run of such a system along with the functions required.

```

1  (write k1 v1; write k1 v2; write k2 v2; write k1 v3; read k1)f
2  where
3    write(k, v) = let
4      i = replica-no ()
5      ret = (local-write k v)ri
6    in
7      ret
8
9    read(k) = let
10     {rvi = (local-read k)ri}i=1n
11   in
12     rv1  $\sqcap$  ...  $\sqcap$  rvn
13
14   local-write(k, v) = {...}, local-read(k) = {...}, replica-no() = {...}

```

Given that the value type is a semilattice, it suffices for correctness to only write to one of the replicas (*local-write*). The function *write* above does that after choosing the replica, say based on the proximity to the user or traffic load. (We assume that *local-write* returns a Boolean flagging its success.) A crucial factor in the correctness of this implementation is a propagation policy enforcing eventual consistency. It is only under that condition that the ordering of writes does not matter. Reading is different: It requires reading all the replicas first (*local-read*) and then performing a join operation ( $\sqcap$ ) on the results.

For a sample run, we again take  $n = 2$  and call the  $\lambda(\text{refut})$  code above p. Furthermore, let  $e_1 = \text{write } k_1 \ v_1$ ,  $e_2 = \text{write } k_1 \ v_2$ ,  $e_3 = \text{write } k_2 \ v_2$ ,  $e_4 = \text{write } k_1 \ v_3$ , and  $e_5 = \text{read } k_1$ . Let also  $v_2 \sqsubseteq v_3 \sqsubseteq v_1$ . Dropping the *where* clauses like the two previous examples, one gets

$p \rightarrow^*$  (PRG),(AP-F)

$((\text{let } i = \text{replica-no } () , \text{ ret} = (\dots)^{r_i} \text{ in ret}); e_2; \dots; e_5)^f \rightarrow$  (AP-F)

$((\text{let } i = 2 , \text{ ret} = (\dots)^{r_i} \text{ in ret}); e_2; \dots; e_5)^f \rightarrow$  (LT-E)

$(\text{let ret} = (\text{local-write } k_1 \ v_1)^{r_2}; e_2; \dots; e_5)^f \rightarrow$  (LT-T)

$(\text{ret}_1^2; e_2; \dots; e_5)^f \parallel \text{ret}_1^2 \mid \Rightarrow ( \text{local-write } k_1 \ v_1 )^{r_2} \rightarrow$  (AP-F)

$(\text{ret}_1^2; e_2; \dots; e_5)^f \parallel \text{ret}_1^2 \mid \Rightarrow ( \text{false} )^{r_2} \rightarrow$  (DF)

$( \text{false} ; e_2 \dots; e_5 )^f \parallel \text{ret}_1^2 \mid \Rightarrow (\text{false})^{r_2} \rightarrow$  (SQ)

$( \text{write } k_1 \ v_2 ; \dots; e_5 )^f \parallel \text{ret}_1^2 \mid \Rightarrow (\text{false})^{r_2} \rightarrow^*$  (...)

$(\text{ret}_2^2; \dots; e_5)^f \parallel \text{ret}_1^2 \mid \Rightarrow (\text{false})^{r_2} \parallel \text{ret}_2^2 \mid \Rightarrow ( \text{true} )^{r_2} \rightarrow^*$  (...)

$( \text{write } k_2 \ v_2 ; \dots; e_5 )^f \parallel \text{ret}_1^2 \mid \Rightarrow (\text{false})^{r_2} \parallel \text{ret}_2^2 \mid \Rightarrow (\text{true})^{r_2} \rightarrow^*$  (...)

$( \text{read } k_1 )^f \parallel \text{ret}_1^2 \mid \Rightarrow (\text{false})^{r_2} \parallel \text{ret}_2^2 \mid \Rightarrow (\text{true})^{r_2}$

$\parallel \text{ret}_1^1 \mid \Rightarrow (\text{true})^{r_2} \parallel \text{ret}_2^2 \mid \Rightarrow (\text{true})^{r_2} \rightarrow$  (AP-F)

$(\text{let } \{rv_i = (\text{local-read } k_1)^{r_i}\}_{i=1}^2 \text{ in } rv_1 \sqcap rv_2)^f \parallel \text{ret}_1^2 \mid \Rightarrow (\text{false})^{r_2} \parallel \text{ret}_2^2 \mid \Rightarrow (\text{true})^{r_2}$

$\parallel \text{ret}_1^1 \mid \Rightarrow (\text{true})^{r_2} \parallel \text{ret}_2^2 \mid \Rightarrow (\text{true})^{r_2} \rightarrow^*$  (LT-T)

$(rv_1 \sqcap rv_2)^f \parallel rv_1 \mid \Rightarrow (\text{local-read } k_1)^{r_1} \parallel rv_2 \mid \Rightarrow ( \text{local-read } k_1 )^{r_2} \parallel \text{ret}_1^2 \mid \Rightarrow (\text{false})^{r_2}$

$\parallel \text{ret}_2^2 \mid \Rightarrow (\text{true})^{r_2} \parallel \text{ret}_1^1 \mid \Rightarrow (\text{true})^{r_2} \parallel \text{ret}_2^2 \mid \Rightarrow (\text{true})^{r_2} \rightarrow$  (AP-F)

$(rv_1 \sqcap rv_2)^f \parallel rv_1 \mid \Rightarrow (\text{local-read } k_1)^{r_1} \parallel rv_2 \mid \Rightarrow ( v_3 )^{r_2} \parallel \text{ret}_1^2 \mid \Rightarrow (\text{false})^{r_2}$

$\parallel \text{ret}_2^2 \mid \Rightarrow (\text{true})^{r_2} \parallel \text{ret}_1^1 \mid \Rightarrow (\text{true})^{r_2} \parallel \text{ret}_2^2 \mid \Rightarrow (\text{true})^{r_2} \rightarrow$  (DF)

$(rv_1 \sqcap v_3)^f \parallel rv_1 \mid \Rightarrow ( \text{local-read } k_1 )^{r_1} \parallel rv_2 \mid \Rightarrow (v_3)^{r_2} \parallel \text{ret}_1^2 \mid \Rightarrow (\text{false})^{r_2}$

$\parallel \text{ret}_2^2 \mid \Rightarrow (\text{true})^{r_2} \parallel \text{ret}_1^1 \mid \Rightarrow (\text{true})^{r_2} \parallel \text{ret}_2^2 \mid \Rightarrow (\text{true})^{r_2} \rightarrow$  (AP-F)

$(rv_1 \sqcap v_3)^f \parallel rv_1 \mid \Rightarrow ( v_2 )^{r_1} \parallel rv_2 \mid \Rightarrow (v_3)^{r_2} \parallel \text{ret}_1^2 \mid \Rightarrow (\text{false})^{r_2}$

$\parallel \text{ret}_2^2 \mid \Rightarrow (\text{true})^{r_2} \parallel \text{ret}_1^1 \mid \Rightarrow (\text{true})^{r_2} \parallel \text{ret}_2^2 \mid \Rightarrow (\text{true})^{r_2} \rightarrow$  (DF)

$( v_2 \sqcap v_3 )^f \parallel rv_1 \mid \Rightarrow (v_2)^{r_1} \parallel rv_2 \mid \Rightarrow (v_3)^{r_2} \parallel \text{ret}_1^2 \mid \Rightarrow (\text{false})^{r_2}$

<sup>1</sup>Conflict-free Replicated Datatype [8]



$$\begin{aligned}
& || \text{ret}_2^2 \mid => (\text{true})^{r_2} || \text{ret}_1^1 \mid => (\text{true})^{r_2} || \text{ret}_2^1 \mid => (\text{true})^{r_2} \rightarrow \\
& (v_3)^f || \text{rv}_1 \mid => (v_2)^{r_1} || \text{rv}_2 \mid => (v_3)^{r_2} || \text{ret}_1^2 \mid => (\text{false})^{r_2} \\
& || \text{ret}_2^2 \mid => (\text{true})^{r_2} || \text{ret}_1^1 \mid => (\text{true})^{r_2} || \text{ret}_2^1 \mid => (\text{true})^{r_2} \quad \square
\end{aligned}
\tag{AP-F}$$

**Remark 3.7.** Pure functional programming of a single node database is not the focus of this article. Example 3.6 assumed that writing to a database is done in a pure fashion, namely, *local-write* is pure. For more on that, see the recent developments of Rob Norris<sup>2</sup>. Our focus in Example 3.6 is on showing pure message passing between nodes of a replicated key-value store.  $\square$

## REFERENCES

- [1] S. H. Haeri. Observational Equivalence and a New Operational Semantics for Lazy Evaluation with Selective Strictness. In Z. Majkic, S.-Y. Hsieh, J. Ma, I. M. M. El Emary, and K. S. Husain, editors, *TMFCS*, pages 143–150. ISRST, July 2010.
- [2] S. H. Haeri and P. Van Roy. Distributed  $\lambda$ -Calculus. Technical report, ICTEAM, UCLouvain, Belgium, October 2019. Available online at: <http://hdl.handle.net/2078.1/228131>.
- [3] S. H. Haeri and P. Van Roy. A Family of  $\lambda$ -Calculi with Ports. In 21<sup>st</sup> *TFP*. LNCS, February 2020. Accepted, Available online at: <http://hdl.handle.net/2078.1/228133>.
- [4] S. H. Haeri and P. Van Roy. Piecewise Relative Observational Purity. In 4<sup>th</sup> *ProWeb*. ACM, March 2020. Accepted.
- [5] G. Khan. The Semantics of a Simple Language for Parallel Programming. *Inf. Proc.*, 74:471–475, 1974.
- [6] J. Launchbury. A Natural Semantics for Lazy Evaluation. In 20<sup>th</sup> *POPL*, pages 144–154. ACM, 1993.
- [7] J. Niehren, J. Schwinghammer, and G. Smolka. A Concurrent Lambda Calculus with Futures. *TCS*, 364(3):338–356, 2006.
- [8] M. Shapiro, N. M. Preguiça, C. Baquero, and M. Zawirski. Conflict-Free Replicated Data Types. In X. Défago, F. Petit, and V. Villain, editors, 13<sup>th</sup> *SSS*, volume 6976 of *LNCS*, pages 386–400. Springer, October 2011.

<sup>2</sup><https://www.youtube.com/watch?v=NJrgj1vQeAI>