

Treo: Textual Syntax for Reo Connectors

Kasper Dokter Farhad Arbab

Formal methods
Centrum Wiskunde & Informatica
Amsterdam, Netherlands

{K.P.C.Dokter, Farhad.Arbab}@cwi.nl

Reo is an interaction-centric model of concurrency for compositional specification of communication and coordination protocols. Formal verification tools exist to ensure correctness and compliance of protocols specified in Reo, which can readily be (re)used in different applications, or composed into more complex protocols. Recent benchmarks show that compiling such high-level Reo specifications produces executable code that can compete with or even beat the performance of hand-crafted programs written in languages such as C or Java using conventional concurrency constructs.

The original declarative graphical syntax of Reo does not support intuitive constructs for parameter passing, iteration, recursion, or conditional specification. This shortcoming hinders Reo's uptake in large-scale practical applications. Although a number of Reo-inspired syntax alternatives have appeared in the past, none of them follows the primary design principles of Reo: a) declarative specification; b) all channel types and their sorts are user-defined; and c) channels compose via shared nodes. In this paper, we offer a textual syntax for Reo that respects these principles and supports flexible parameter passing, iteration, recursion, and conditional specification. In on-going work, we use this textual syntax to compile Reo into target languages such as Java, Promela, and Maude.

1 Introduction

The advent of multicore processors has intensified the significance of coordination in concurrent applications. A programmer tackles the coordination concern of an application by specifying a (usually implicit) protocol that defines all possible permissible interactions among different active components of the application. Depending on the language used, programmers define their protocols at different levels of abstraction. A threading library, for instance, generally offers only basic synchronization primitives, such as locks and semaphores, that can be inserted into imperative code to ensure execution follows an implicitly defined protocol. Exogenous coordination languages offer syntax to programmers to explicitly define their interaction protocols at a high level of abstraction.

Reo [1, 2] is an example of such a coordination language that defines an interaction protocol as a *connector*: a graph-like structure that enables (a)synchronous data flow along its edges (cf., Figure 1). Each edge, called a *channel*, has a user-defined type and two channel ends. The type determines the behavior of the channel, specified as a constraint on the flows of data at its two ends. This constraint is expressed in a user-defined semantic sort, such as timed data streams, constraint automata, or coloring semantics [13]. A channel end is either a *source end* through which the channel accepts data, or a *sink end* through which the channel offers data. Multiple channel ends coincident at a vertex of the connector together form a *node*. Nodes have predefined ‘merge-replicate’ behavior: a node repeatedly accepts a datum from one of its coincident sink ends, chosen non-deterministically, and offers a copy of that datum through every one of its coincident source ends.

Tools for Reo have been implemented as a collection of Eclipse plugins called the ECT [8]. The main plugin in this tool set consists of a graphical editor that allows a user to draw a connector on a canvas.

The graphical editor has an intuitive interface with a flat learning curve. However, it does not provide constructs to express parameter passing, iteration, recursion, or conditional construction of connector graphs. Such language constructs are more easily offered by familiar programming language constructs in a textual representation of connectors.

In the context of Vereofy (a model checker for Reo), Baier, Blechmann, Klein, and Klüppelholz developed the Reo Scripting Language (RSL) and its companion language, the Constraint Automata Reactive Module Language (CARML) [3, 19]. RSL is the first textual language for Reo that includes a construct for iteration, and a limited form of parameter passing. Primitive channels and nodes are defined in CARML, a guarded command language for specification of constraint automata. Programmers then combine CARML specified constraint automata as primitives in RSL to construct complex connectors and/or complete systems. In contrast to the declarative nature of the graphical syntax of Reo, RSL is imperative.

Jongmans developed the First-Order Constraint Automata with Memory Language (FOCAML) [12], a textual declarative language that enables compositional construction of connectors from a (pre-defined set of) primitive components. As a textual representation for Reo, however, FOCAML has poor support for its primary design principle: Reo channels are user-defined, not tied to any specific formalism to express its semantics, and compose via shared nodes with predefined merge-replicate behavior. Although FOCAML components are user-defined, FOCAML requires them to be of the same predefined semantic sort (i.e., constraint automata with memory [4]). The primary concept of Reo nodes does not exist in FOCAML, which forces explicit construction of their ‘merge-replicate’ behavior in FOCAML specifications.

Jongmans et al. have shown by benchmarks that compiling Reo specifications can produce executable code whose performance competes with or even beats that of hand-crafted programs written in languages such as C or Java using conventional concurrency constructs [18, 14, 16, 15, 17]. A textual syntax for Reo that preserves its declarative, compositional nature, allows user-defined primitives, and faithfully complies with the semantics of its nodes can significantly facilitate the uptake of Reo for specification of protocols in large-scale practical applications.

In this paper, we introduce Treo, a declarative textual language for component-based specification of Reo connectors with user-defined semantic sorts and predefined node behavior. We recall the basics of Reo (Section 2). We describe the structure of a Treo file by means of an abstract syntax (Section 3). In Listing 1, we provide a concrete syntax of Treo as an ANTLR4 grammar [22]. In on-going work, we currently use Treo to compile Reo into target languages such as Java, Promela, and Maude [23]. The construction of the Treo compiler is based on the theory of stream constraints [7].

In order to preserve the agnosticism of Reo regarding the concrete semantics of its primitives, Treo uses the notion of user-defined *semantic sorts*. A user-defined semantic sort consists of a set of component instances together with a composition operator \wedge , a substitution operator $[/]$, and a trivial component \top (Section 4). The composition operator defines the behavior of composite components as a composition of its operands. The substitution operator binds nodes in the interface or passes values to parameters.

For a given semantic sort, we define the meaning of abstract Treo programs (Section 5). Treo is very liberal with respect to parameter values. A component definition not only accepts the usual (structured) data as actual parameters, but also other component instances and other component definitions. Among other benefits, this flexible parameter passing supports *component sharing*, which is useful to preserve component encapsulation [5, Figure 2].

A given semantics sort may possibly distinguish between inputs and outputs. Thus, not all combinations of components may result in a valid composite component. For example, the composition may not be defined, if two components share an output. In Treo, however, it is safe to compose compo-

nents on their outputs, because, complying with the semantics of Reo, the compiler inserts special *node components* to ensure well-formed compositions (Section 6).

We conclude by discussing related work (Section 7), and pointing out future work (Section 8).

2 Reo

We briefly recall the basics of the Reo language and refer to [1] and [2] for further details. Reo is a language for specification of interaction protocols, originally proposed with a graphical syntax. A Reo program, called a *connector*, is a graph-like structure whose edges consist of *channels* that enable synchronous and asynchronous data flow and whose vertices consist of *nodes* that synchronously route data among multiple channels. Each channel has a type and two channel ends. Each channel end is either a *source end*, through which the channel accepts data, or a *sink end*, through which the channel offers data. The type of a channel completely defines the behavior of the channel in some user-defined semantic sort. Reo is agnostic regarding the semantic sort that expresses the behavior of its channel types, so long as the semantic sort preserves Reo’s compositional construction principle (cf., Definition 4.1). Table 1 shows some frequently used channels and an example node together with an informal description of their behavior.

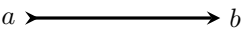

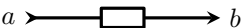
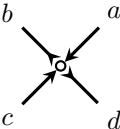
	A Sync channel accepts datum from its source end a , when its simultaneous offer of this datum at its sink end b succeeds.
	A SyncDrain channel simultaneously accepts data from both its source ends a and b and loses the data.
	An empty FIFO ₁ accepts data from its source end a and becomes a full FIFO ₁ . A full FIFO ₁ offers its stored data at its sink end b and, when its offer succeeds, it becomes an empty FIFO ₁ again.
	A Reo node accepts a datum from one of its coincident sink ends (a or c), when its simultaneous offer to dispense a copy of this datum through every one of its coincident source ends (b and d) succeeds.

Table 1: Informal description of the behavior of nodes and of some channels in Reo.

The key concept in Reo is composition, which allows a programmer to build complex connectors out of simpler ones. For example, using the channels in Table 1, we can construct the Alternator _{k} connector, for $k \geq 2$, as shown in Figure 1. For $k = 2$, the Alternator₂ consists of four nodes (a_1, a_2, b_1 , and b_2) and four channels, namely a SyncDrain channel (between a_1 and a_2), two Sync channels (from a_1 to b_1 , and from a_2 to b_2), and a FIFO₁ channel (from b_2 to b_1).




Figure 1: Construction of the Alternator _{k} Reo connector, for $k \geq 2$.

The behavior of the Alternator₂ connector is as follows. Suppose that the environment is ready to offer a datum at each of the nodes a_1 and a_2 , and ready to accept a datum from node b_1 . According to Table 1, nodes a_1 and a_2 both offer a copy of their received datum to the SyncDrain channel. The SyncDrain channel ensures that nodes a_1 and a_2 accept data from the environment only simultaneously. The Sync channel from a_1 to b_1 ensures that node b_1 simultaneously obtains a copy of the datum offered at a_1 . By definition, node b_1 either accepts a datum from the connected Sync channel or it accepts a datum from the FIFO₁ channel (but not from both simultaneously), and offers this datum immediately to the environment. Because the FIFO₁ is initially empty, b_1 has no choice but to accept and dispense the datum from a_1 . Simultaneously, the Sync channel from a_2 to b_2 ensures that the value offered at a_2 is stored in the FIFO₁ buffer. In the next step, the environment at node b_1 has no choice but to retrieve the datum in the buffer, after which the behavior repeats.

3 Treo syntax

We now present a textual representation for the graphical Reo connectors in Section 2. Table 2 shows the abstract syntax of Treo.

$$\begin{array}{ll}
 K ::= I \mid KND & D ::= V \mid \langle U_0 \rangle (U_1) \{ C \} \\
 L ::= \varepsilon \mid L, T \mid L, T_0..T_1 & C ::= V \mid A \mid C_0 C_1 \mid \{ C \mid P \} \mid D(L)(U) \\
 U ::= \varepsilon \mid U, V & T ::= V \mid C \mid D \mid [L] \mid T_0 : T_1 \mid T[L] \mid F(L) \\
 V ::= N \mid V[L] & P ::= V \in T \mid R(L) \mid \neg P \mid P_0 \wedge P_1 \mid P_0 \vee P_1 \mid (P)
 \end{array}$$

Table 2: Abstract syntax of Treo, with start symbol K (a source file), and terminal symbols for imports (I), primitive components (A), functions (F), relations (R), names (N), and the empty list (ε). The bold vertical bar in $\{C \mid P\}$ is just text.

We introduce the symbols in the abstract syntax by identifying them in some concrete examples. These concrete examples are Treo programs that can be parsed using the concrete Treo syntax shown in Listing 1.

Listing 1: Concrete ANTLR4 syntax of Treo (Treo.g4).

```

1 grammar Treo;
2 file      : sec? imp* assg* EOF;
3 sec      : 'section' name ';' ;
4 imp      : 'import' name ';' ;
5 assg     : ID defn ;
6 defn     : var | params? nodes comp ;
7 comp     : defn vals? args | var | '{' atom+ '}' | '{' comp* ('|' pred)? '}'
8          | 'for' '(' ID 'in' list ')' comp
9          | 'if' '(' pred ')' comp ('else' '(' pred ')' comp)* ('else' comp)? ;
10 atom    : STRING ; /* Example syntax for primitive components */
11 pred     : 'true' | 'false' | '(' pred ')' | var 'in' list
12          | term op=('<=' | '<' | '>=' | '>' | '=' | '!=') term
13          | var | 'forall' ID 'in' list ':' pred | 'exists' ID 'in' list ':' pred
14          | 'not' pred | pred ('and'|',' ) pred | pred 'or' pred | pred 'implies' pred ;
15 term    : var | NAT | BOOL | STRING | DEC | comp | defn | list | 'len(' term ')'

```

```

16      | '(' term ')' | <assoc=right> term list | <assoc=right> term '^' term
17      | '-' term | term op=('*' | '/' | '%' | '+' | '-') term ;
18 vals  : '<' '>' | '<' term (',' term)* '>' ;
19 list  : '[' ']' | '[' item (',' item)* ']' ;
20 item  : term | term '..' term | term ':' term ;
21 args  : '(' ')' | '(' var (',' var)* ')' ;
22 params : '<' '>' | '<' var (',' var)* '>' ;
23 nodes : '(' ')' | '(' node (',' node)* ')' ;
24 node  : var (io=('?' | '!' | ':') ID?)? ;
25 var   : name list* ;
26 name  : (ID '..')* ID ;
27 NAT   : ('0' | [1-9][0-9]*) ;
28 DEC   : ('0' | [1-9][0-9]*) '.' [0-9]+ ;
29 BOOL  : 'true' | 'false' ;
30 ID    : [a-zA-Z_][a-zA-Z0-9_]* ;
31 STRING : '\"' .*? '\"' ;
32 SPACES : [ \t\r\n]+ -> skip ;
33 SL_COMM : '//' .*? ('\n'|EOF) -> skip ;
34 ML_COMM : '/*' .*? '*/' -> skip ;

```

Consider the following Treo file (K in Table 2) representing the Alternator₂:

```

import syncdrain; import sync; import fifo1;
alternator2(a1,a2,b1) { sync(a1,b1) syncdrain(a1,a2) sync(a2,b2) fifo1(b2,b1) }

```

On the first line, we import (I) three different *component definitions*. On the second line, we define the `alternator2` component (ND). Its definition (D) has no parameters ($\langle U_0 \rangle$), and three nodes, `a1`, `a2`, and `b1`, in its interface ($\langle U_1 \rangle$). The body ($\{C\}$) of this definition consists of a set of *component instances* that interact via shared nodes. The first component instance `sync(a1,b1)` is an instantiation ($D\langle L \rangle(U)$) of the imported `sync` definition (D) with nodes `a1` and `b1` ($\langle U \rangle$) and without any parameters ($\langle L \rangle$).

All nodes that occur in the body, but not in the interface, are hidden. Hiding renames a node to a fresh inaccessible name, which prevents it from being shared with other components. In the case of `alternator2`, node `b2` is not part of the interface, and hence hidden.

Constructed from existing components, `alternator2` is a *composite* component (C_0C_1). However, not every component is constructed from existing components, and we call such components *primitive* (A). The following Treo code shows a possible (primitive) definition of the `fifo1` component.

```

fifo1(a?,b!) { empty -{a},true-> full; full -{b},true-> empty; }

```

The definition of the `fifo1` differs from the definition of the `alternator2` in two ways.

The first difference is that the `fifo1` component is (in this case) defined directly as a *constraint automaton* [4]. Constraint automata constitute a popular semantic sort for specification of Reo component types, and forms the basis of the Lykos compiler [12]. However, constraint automata are not the *de facto* standard: the literature offers more than thirty different semantic sorts for specification of Reo components [13], such as the coloring semantics and timed data stream semantics. To accommodate the generality that disparate semantics allow, Treo features *user-defined semantic sorts*, which means that the syntax for primitive components is user-defined. For example, this means that we may also define the `fifo1` component by referring to a Java file via `fifo1(a?,b!){ "MyFIF01.java" }`.

The second difference is that the nodes `a` and `b` in the interface are *directed*. That is, each of its interface nodes is either of type input or output, designated by the markers `?` and `!`, respectively. In

Reo, it is safe to join two channels on a shared sink node (e.g., node b_1 in Figure 1). However, the composition operators in most Reo semantics do not automatically produce the correct behavior for such nodes (e.g., see [4, Section 4.3] for further details). Therefore, most Reo semantics require *well-formed* compositions, wherein each node has at most one input channel end and at most one output channel end.

The restriction of well-formed compositions can be very inconvenient in practice. To ensure well-formed compositions, a programmer must implement every Reo node with more than one input or output channel end as a *node component*. The interface of this node component is determined by its *degree*, which is a pair (i,o) giving the numbers of its coincident source and sink ends. Such explicit node components make component constructions verbose and hard to maintain. For convenience, the Treo compiler uses the above input/output annotations to compute the degree of each node in a composition, and subsequently inserts the correct node components in the construction. We may view the input/output annotations as syntactic sugar that ensures well-formed compositions. This feature allows programmers to remain oblivious to these annotations and well-formed composition.

The ellipses in Figure 1 signify the parametrized construction of the Alternator_k connector, for $k > 2$. This notation is informal and not supported in the graphical Reo editor [8], which offers no support for parametrized constructions. In Treo, however, we can define the Alternator_k connector as:

```
alternator<k>(a[1:k],b[1]) { sync(a[1],b[1])
  { syncdrain(a[i-1],a[i]) sync(a[i],b[i]) fifo1(b[i],b[i-1]) | i in [2..k] } }
```

The definition of the `alternator` depends on a parameter k . Since Treo is a strongly typed language with type-inferencing, there is no need to specify a type for the (integer) parameter k . The interface consists of an array of nodes $a[1:k]$ and the single node $b[1]$. Here, $[1:k]$ is an abbreviation for the list $[1..k]$ that contains a single list of length k . The array $a[1:2]$ stands for the *slice* $[a[1],a[2]]$ of a , while the expression $a[1..2]$ stands for the element $a[1][2]$ in a (cf., Equation (2)). For iteration, we write $\{ \dots | i \text{ in } [2..k] \}$ using set-comprehension ($\{C | P\}$ in Table 2).

Instead of defining `alternator` iteratively, we may also provide a recursive definition as follows:

```
recursive_alternator(a[1:k],b[1],b[k]) { recursive_alternator(a[1:k-1],b[1],b[k-1])
  { syncdrain(a[k-1],a[k]) sync(a[k],b[k]) fifo1(b[k-1],b[k]) | k > 1 } }
```

Here, the value of k is defined by the size of $a[1:k]$, and we use set-comprehension $\{ \dots | k > 1 \}$ for conditional construction, as well. Indeed, the resulting set of component instances is non-empty, only if $k > 1$ holds. Although Treo syntax allows recursive definitions, the semantics presented in Section 5 does not yet support recursion, which we leave as future work.

We illustrate the practicality of Treo by providing code for a chess playing program [12, Figure 3.29]. In this program, two teams of chess engines compete in a game of chess. We define a chess team as the following Treo component:

```
import parse; /* and the other imports */
team<engine[1:n]>(inp,out) {
  for (i in [1..n]) {
    engine[i](inp,best[i]) parse(best[i],p[i])
    if (i > 1) concatenate(a[i-1],p[i],a[i]) }
  sync(best[1],a[1]) majority(a[n],b) syncdrain(b,c)
  fifo1(inp,c) move(b,d) concatenate(c,d,out) }
```

The for-loop `for (i in [1..n]) ...` and if-statement `if (i > 1) ...` are just syntactic sugar for set-comprehensions $\{ \dots | i \text{ in } [1..n] \}$ and $\{ \dots | i > 1 \}$, respectively. The team component depends on an array `engine[1:n]` of parameters. This array does not contain the usual data values,

but consists of Treo component definitions. In the body of the team component, these definitions are instantiated via `engine[i] (inp, best[i])`. In RSL [3, 19] and FOCAML [12], it is impossible to pass a component as a parameter, which makes these languages less expressive than Treo.

We may view the team component as an example of role-oriented programming [6]. Indeed, the team component encapsulates a list of chess engines in a component, so that they can collectively be used as a single participant in a chess match:

```
match() { fifo1full<"">(a,b) fifo1(c,d) team<[eng1, eng2]>(a,d) team<[eng3]>(b,c) }
```

Treo treats not only component definitions, but also component instances as values. By passing a single component instance as a parameter to multiple components, this feature allows *component (instance) sharing* (cf., [5, Figure 2]). Hence, it is straightforward to implement a chess match, wherein a single instance of a chess engine plays against itself.

4 Semantic sorts

As noted in Section 3, Reo channels can be defined in many different semantic formalisms [13], such as the constraint automaton semantics, the coloring semantics, or the timed data stream semantics. Although each sort of Reo semantics has its unique properties, each of them can be used to define a collection of composable components with parameters and nodes, which we call a *semantic sort*:

Definition 4.1 (Semantic sort). A semantic sort over a set of names \mathcal{N} with values from \mathcal{V} is a tuple $(\mathcal{C}, \wedge, [/\], \top)$ that consists of a set of components \mathcal{C} , a composition operator $\wedge : \mathcal{C} \times \mathcal{C} \rightarrow \mathcal{C}$, a substitution operator $[/\] : \mathcal{C} \times (\mathcal{N} \cup \mathcal{V}) \times \mathcal{N} \rightarrow \mathcal{C}$, and a trivial component $\top \in \mathcal{C}$.

We assume that the set of names and the set of values are disjoint, i.e., $\mathcal{N} \cap \mathcal{V} = \emptyset$. For convenience, we write $C \wedge C'$ for $\wedge(C, C')$, and $C[y/x]$ for $[/\](C, y, x)$. For any semantic sort T , we write \mathcal{C}_T for its set of components, \wedge_T for its composition operator, $[/\]_T$ for its substitution operator, and \top_T for its trivial component. The composition operator \wedge_T ensures that the behavior of finite non-empty compositions is well-defined. To empty compositions we assign the trivial component \top_T . The substitution operator $[/\]_T$ allows us to change the interface of a component via renaming or instantiation. Let $C \in \mathcal{C}_T$ be a component and $x \in \mathcal{N}$ a name. For a name $y \in \mathcal{N}$, the construct $C[y/x]_T$ renames every occurrence of name x in C to y . For a value $y \in \mathcal{V}$, the construct $C[y/x]_T$ instantiates (parameter) x in C to y . (See Example 4.3 for an example of the distinction between renaming and instantiations.)

A semantic sort T implicitly defines an interface for each component $C \in \mathcal{C}$ via the map $\text{supp} : \mathcal{C}_T \rightarrow 2^{\mathcal{N}}$ defined as $\text{supp}(C) = \{x \in \mathcal{N} \mid C[y/x]_T \neq C, \text{ for some name } y \in \mathcal{N}\}$. If name x does not ‘occur’ in C , substitution of x by any name y does not affect C , i.e., $C[y/x]_T = C$.

Example 4.1 (Systems of differential equations). The set ODE of systems of ordinary differential equations with variables from \mathcal{N} and values $\mathcal{V} = \{v : \mathbb{R} \rightarrow \mathbb{R}\}$ constitute a semantic sort. Composition is union, substitution is binding a name or value to a given name, and the trivial component is the empty system of equations. Using the ODE semantic sort, we can define continuous systems in Treo. \triangle

Example 4.2 (Process calculi). Consider the process calculus CSP, proposed by Hoare [10]. The set CSP of all such process algebraic terms comprises a semantic sort. Each process can participate in a number of events, which we can interpret as names from a given set \mathcal{N} . We model the composition of CSP processes P and Q by means of the interface parallel operator $P \parallel [X] Q$, where $X \subseteq \mathcal{N}$ is the set of event names shared by P and Q . We define substitution as simply (1) renaming the event, if a name is substituted for an event; or (2) hiding the event, if a values is substituted for an event. Since neither STOP nor SKIP shares any event with its environment, we may use either one to denote the trivial component. \triangle

Example 4.3 (I/O-components). Let T be a semantic sort over \mathcal{N} and \mathcal{V} . We define the I/O-component sort IO_T over T using the notion of a primitive I/O-component of sort T .

A *primitive I/O-component* P of sort T is a tuple (C, I, O) , where $C \in \mathcal{C}_T$ is a component, $I \subseteq \mathcal{N}$ is a set of input names, $O \subseteq \mathcal{N}$ is a set of output names. For $P \subseteq \mathcal{N}$ and $x \in \mathcal{N}$ and $y \in \mathcal{N} \cup \mathcal{V}$, define

$$P[y/x] = \begin{cases} (P - \{x\}) \cup \{y\} & \text{if } x \in P \text{ and } y \in \mathcal{N} \\ P - \{x\} & \text{if } x \in P \text{ and } y \in \mathcal{V} \\ P & \text{otherwise} \end{cases} \quad (1)$$

We define substitution on primitive I/O-components as $(C, I, O)[y/x] = (C[y/x], I[y/x], O[y/x])$, for all $x \in \mathcal{N}$ and $y \in \mathcal{N} \cup \mathcal{V}$. We denote the set of primitive I/O-components over T as \mathcal{P}_T .

An *I/O-component* of sort T is a sequence $P_1 \cdots P_n \in \mathcal{P}_T^*$, with $n \geq 0$, of primitive I/O-components of sort T . Composition of I/O-components is concatenation \cdot of sequences. The trivial I/O-component is the empty sequence ε . We define substitution of composite I/O-components as $(P_1 \cdots P_n)[y/x] = P_1[y/x] \cdots P_n[y/x]$, for all $x \in \mathcal{N}$ and $y \in \mathcal{N} \cup \mathcal{V}$. Hence, $\text{IO}_T = (\mathcal{P}_T^*, \cdot, [/], \varepsilon)$ is a semantic sort. \triangle

5 Denotational semantics

We define the denotational semantics of the Treo language over a fixed, but arbitrary, semantic sort T . The main purpose of this denotational semantics is to provide a clear abstract structure that guides the implementation of Treo parsers. The syntax to which this denotational semantics applies is the abstract syntax in Table 2. The general structure of our denotational semantics is quite standard, and adheres to Schmidt's notation [24].

Although Treo syntax allows recursive definitions, the semantics presented in this section does not support this feature. Since not all recursive definitions define finite compositions of components, extending the current semantics with recursion is not straightforward, and we leave it as future work.

Variables and terms in Treo are structured as non-rectangular arrays. The set of all (*ragged*) arrays over a set X is the smallest set X^\square such that both $X \subseteq X^\square$ and $[x_0, \dots, x_{n-1}] \in X^\square$, if $n \geq 0$ and $x_i \in X^\square$ for all $0 \leq i < n$. For example, the set \mathbb{N}^\square of ragged arrays over integers contains all natural numbers from \mathbb{N} as 'atomic' arrays, as well as the array $[37, [], [[2, [55], 3]]] \in \mathbb{N}^\square$. Every ragged array has a length, which can be computed via the map $\text{len} : X^\square \rightarrow \mathbb{N}$ defined inductively as $\text{len}(x) = 0$, if $x \in X$, and $\text{len}([x_0, \dots, x_{n-1}]) = n$, otherwise. If $x = [x_0, \dots, x_{n-1}] \in X^\square$ is a ragged array, we access its entries via the function application $x(i) = x_i$, for every $0 \leq i < n$. We extend the access map \mathbb{N}^\square by defining

$$x([i_0, \dots, i_n]) = \begin{cases} x(i_0)([i_1, \dots, i_n]) & \text{if } i_0 \in \mathbb{N} \\ [x(i_{00})([i_1, \dots, i_n]), \dots, x(i_{0m})([i_1, \dots, i_n])] & \text{if } i_0 = [i_{00}, \dots, i_{0m}] \end{cases}, \quad (2)$$

whenever the right-hand side is defined. Two ragged arrays $x \in X^\square$ and $y \in Y^\square$ have the same structure ($x \simeq y$) iff $x \in X$ and $y \in Y$, or $\text{len}(x) = \text{len}(y)$ and $x(i) \simeq y(i)$ for all $0 \leq i < \text{len}(x)$. We can flatten a ragged array from X^\square to a sequence over X via the map $\text{flatten} : X^\square \rightarrow X^*$ defined as $\text{flatten}(x) = x$, if $x \in X$, and $\text{flatten}([x_0, \dots, x_{n-1}]) = \text{flatten}(x_0) \cdots \text{flatten}(x_{n-1})$, otherwise.

Suppose that semantic sort T is defined over a set of names \mathcal{N} and a set of values \mathcal{V} , with $\mathcal{N} \cap \mathcal{V} = \emptyset$. For simplicity, we assume that, for every component $C \in \mathcal{C}_T$, its support $\text{supp}(C) \subseteq \mathcal{N}$ is finite. Since Treo views components as values, we assume the inclusion $\mathcal{C}_T \subseteq \mathcal{V}$.

We assume that the set of names \mathcal{N} is closed under taking subscripts from \mathbb{N} . That is, if $x \in \mathcal{N}$ is a name and $i \in \mathbb{N}$ is a natural number, then we can construct a fresh name $x_i \in \mathcal{N}$. To construct sequences

of data with variable lengths, we use a map $\text{lst} : \mathbb{N}^2 \rightarrow \mathbb{N}^\square$ that constructs from a pair $(i, j) \in \mathbb{N}^2$ of integers a finite ordered list $[i, i+1, \dots, j]$ in \mathbb{N}^\square .

Recall from Section 3 that a component accepts an arbitrary but finite number of parameters and nodes. Therefore, we define a component definition as a map $D : \mathcal{V}^\square \times \mathcal{N}^\square \rightarrow \mathcal{C}_T \cup \{\zeta\}$ that takes an array of parameter values from \mathcal{V}^\square and an array of nodes from \mathcal{N}^\square and returns a component or an error ζ . Let $\mathcal{D} = (\mathcal{C}_T \cup \{\zeta\})^{\mathcal{V}^\square \times \mathcal{N}^\square}$ be the set of all definitions. As mentioned earlier, Treo also allows definitions as values, which amounts to the inclusion $\mathcal{D} \subseteq \mathcal{V}$.¹

We evaluate every Treo construct in its *scope* $\sigma : N \rightarrow \mathcal{V}^\square$, with $N \subseteq \mathcal{N}$ finite, which assigns a value to a finite collection of locally defined names. We write $\Sigma = \{\sigma : N \rightarrow \mathcal{V}^\square \mid N \subseteq \mathcal{N} \text{ finite}\}$ for the set of scopes. For a name $x \in \mathcal{N}$ and a value $d \in \mathcal{V}^\square$, we have a scope $\{x \mapsto d\} : \{x\} \rightarrow \mathcal{V}^\square$ defined as $\{x \mapsto d\}(x) = d$. For any two scopes $\sigma, \sigma' \in \Sigma$, we have a composition $\sigma\sigma' \in \Sigma$ such that for every $x \in \text{dom}(\sigma) \cup \text{dom}(\sigma')$ we have $(\sigma\sigma')(x) = \sigma'(x)$, if $x \in \text{dom}(\sigma')$, and $(\sigma\sigma')(x) = \sigma(x)$, otherwise. The composite scope $\sigma\sigma'$ can be viewed as an extension of σ that includes definitions and updates from σ' .

Let Names be the set of parse trees with root N , and let $\mathbf{N}[-] : \text{Names} \rightarrow \mathcal{N}$ be the semantics of names. We define the semantics of variables as a map $\mathbf{V}[-] : \text{Variables} \rightarrow (\mathcal{N}^\square \cup \{\zeta\})^\Sigma$, where Variables is the set of parse trees with root V . For a scope $\sigma \in \Sigma$, we define $\mathbf{V}[-](\sigma)$ as follows:

1. $\mathbf{V}[N](\sigma) = \mathbf{N}[N]$;
2. $\mathbf{V}[V[L]](\sigma) = \begin{cases} x(k) & \text{if } \mathbf{V}[V](\sigma) = x \in \mathcal{N}^\square \text{ and } \mathbf{L}[L](\sigma) = k \in \mathbb{N}^\square \\ \zeta & \text{otherwise} \end{cases}$.

Since \mathcal{N} is closed under taking subscripts, we can define $n(i) = n_i$, for all $n \in \mathcal{N}$ and $i \in \mathbb{N}$, which ensures that $x(k) \in \mathcal{N}^\square$ is always defined.

The semantics of arguments is a map $\mathbf{U}[-] : \text{Arguments} \rightarrow (\mathcal{N}^\square \cup \{\zeta\})^\Sigma$, where Arguments is the set of all parse trees with root U . For a scope $\sigma \in \Sigma$, we define $\mathbf{U}[-](\sigma)$ as follows:

1. $\mathbf{U}[\varepsilon](\sigma) = []$;
2. $\mathbf{U}[U, V](\sigma) = \begin{cases} [x_1, \dots, x_{n+1}] & \text{if } \mathbf{U}[U](\sigma) = [x_1, \dots, x_n] \text{ and } \mathbf{V}[V](\sigma) = x_{n+1} \\ \zeta & \text{otherwise} \end{cases}$.

Let Functions be the set of parse trees with root F , and let $\mathbf{F}[-] : \text{Functions} \rightarrow \{\mathcal{V}^k \rightarrow \mathcal{V} \mid k \in \mathbb{N}\}$ be the semantics of functions. The semantics of terms is a map $\mathbf{T}[-] : \text{Terms} \rightarrow (\mathcal{V}^\square \cup \{\zeta\})^\Sigma$, where Terms is the set of parse trees with root T . For a scope $\sigma \in \Sigma$, we define $\mathbf{T}[-](\sigma)$ inductively as follows:

1. $\mathbf{T}[V](\sigma) = \begin{cases} \sigma(\mathbf{V}[V](\sigma)) & \text{if defined} \\ \zeta & \text{otherwise} \end{cases}$;
2. $\mathbf{T}[C](\sigma) = \mathbf{C}[C](\sigma)$, which is well-defined since $\mathcal{C}_T \subseteq \mathcal{V}$;
3. $\mathbf{T}[D](\sigma) = \mathbf{D}[D](\sigma)$, which is well-defined since $\mathcal{D} \subseteq \mathcal{V}$;
4. $\mathbf{T}[L](\sigma) = \mathbf{L}[L](\sigma)$;
5. $\mathbf{T}[T_0 : T_1](\sigma) = \begin{cases} \text{lst}(x_0, x_1 - 1) & \text{if } \mathbf{T}[T_i](\sigma) = x_i \in \mathbb{N} \text{ for } i \in \{0, 1\} \\ \zeta & \text{otherwise} \end{cases}$;
6. $\mathbf{T}[T[L]](\sigma) = \begin{cases} x(k) & \text{if } \mathbf{T}[T](\sigma) = x \in \mathcal{V}^\square \text{ and } \mathbf{L}[L](\sigma) = k \in \mathbb{N}^\square \\ \zeta & \text{otherwise} \end{cases}$;

¹ Such a set of values \mathcal{V} exists only if $\mathcal{V} \mapsto \mathcal{C}_T \cup (\mathcal{C}_T \cup \{\zeta\})^{\mathcal{V}^\square \times \mathcal{N}^\square}$ admits a pre-fixed point. In this work, we simply assume that such \mathcal{V} exists.

$$7. \mathbf{T}[\mathbf{F}(L)](\sigma) = \begin{cases} \mathbf{F}[\mathbf{F}](\mathbf{L}[L](\sigma)) & \text{if } \mathbf{F}[\mathbf{F}] : \mathcal{V}^k \rightarrow \mathcal{V} \text{ and } \text{len}(\mathbf{L}[L](\sigma)) = k \\ \downarrow & \text{otherwise} \end{cases}$$

The semantics of lists is a map $\mathbf{L}[-] : \text{Lists} \rightarrow (\mathcal{V}^\square \cup \{\downarrow\})^\Sigma$, where Lists is the set of parse trees with root L . For a given scope $\sigma \in \Sigma$, we define $\mathbf{S}[-](\sigma)$ inductively as follows:

1. $\mathbf{L}[\varepsilon](\sigma) = []$;
2. $\mathbf{L}[L, T](\sigma) = \begin{cases} [x_1, \dots, x_{n+1}] & \text{if } \mathbf{L}[L](\sigma) = [x_1, \dots, x_n] \in \mathcal{V}^\square \text{ and } \mathbf{T}[T](\sigma) = x_{n+1} \in \mathcal{V} \\ \downarrow & \text{otherwise} \end{cases}$;
3. $\mathbf{L}[L, T_0..T_1](\sigma) = \begin{cases} [x_1, \dots, x_{n+k}] & \text{if } \mathbf{L}[L](\sigma) = [x_1, \dots, x_n] \in \mathcal{V}^\square, \mathbf{T}[T_i](\sigma) = a_i \in \mathcal{V}, \\ & \text{for } i \in \{0, 1\}, \text{ and } \text{lst}(a_0, a_1) = [x_{n+1}, \dots, x_{n+k}] \\ \downarrow & \text{otherwise} \end{cases}$.

Since we use predicates in Treo for list comprehension, we define the semantics of predicates as a map $\mathbf{P}[-] : \text{Predicates} \rightarrow (2^\Sigma)^\Sigma$, where Predicates is the set of all parse trees with root P . For a scope $\sigma \in \Sigma$, we define the semantics $\mathbf{P}[-](\sigma)$ of a predicate P as the set of all extensions of σ that satisfy P . We define $\mathbf{P}[-](\sigma)$ inductively as follows:

1. $\mathbf{P}[V \in T](\sigma) = \begin{cases} \{\sigma\{x \mapsto t_i\} \mid 1 \leq i \leq n\} & \text{if } \mathbf{V}[V](\sigma) = x \notin \text{dom}(\sigma), \mathbf{T}[T](\sigma) = [t_1, \dots, t_n] \\ \{\sigma\} & \text{if } \mathbf{T}[V](\sigma) \in \mathbf{T}[T](\sigma) \\ \emptyset & \text{otherwise} \end{cases}$,
2. If P is $R(L)$, we define $\mathbf{P}[R(L)](\sigma) = \{\sigma' \in \Sigma \mid \sigma'\sigma = \sigma', \mathbf{L}[L](\sigma') \in \mathbf{R}[R]\}$;
3. If P is $\neg P$, we define $\mathbf{P}[\neg P](\sigma) = \{\sigma' \in \Sigma \mid \sigma'\sigma = \sigma', \neg \mathbf{P}[P](\sigma')\}$;
4. If P is $P_0 \wedge P_1$, we define $\mathbf{P}[P_0 \wedge P_1](\sigma) = \mathbf{P}[P_0](\sigma) \cap \mathbf{P}[P_1](\sigma)$;
5. If P is $P_0 \vee P_1$, we define $\mathbf{P}[P_0 \vee P_1](\sigma) = \mathbf{P}[P_0](\sigma) \cup \mathbf{P}[P_1](\sigma)$;
6. If P is (P) , we define $\mathbf{P}[(P)](\sigma) = \mathbf{P}[P](\sigma)$.

For set and list comprehensions, we can iterate over only a finite subset of scopes $\mathbf{P}[P](\sigma)$ of P . We ensure this by restricting the set of scopes to those solutions that are minimal with respect to inclusion of domains. Formally, we write $\text{min}\mathbf{P}[P](\sigma)$ for the set of all scopes that are minimal with respect to \leq defined as $\sigma_1 \leq \sigma_2$ iff $\text{dom}(\sigma_1) \subseteq \text{dom}(\sigma_2)$, for all $\sigma_1, \sigma_2 \in \mathbf{P}[P](\sigma)$.

The semantics of component instances is a map $\mathbf{C}[-] : \text{Components} \rightarrow (\mathcal{C}_T \cup \{\downarrow\})^\Sigma$, where Components is the set of parse trees with root C . Recall that Treo views components as values ($\mathcal{C}_T \subseteq \mathcal{V}$). Given a scope $\sigma \in \Sigma$, we define $\mathbf{C}[-](\sigma)$ inductively as follows:

1. $\mathbf{C}[V](\sigma) = \begin{cases} \sigma(x) & \text{if } \mathbf{V}[V](\sigma) = x \in \text{dom}(\sigma) \text{ and } \sigma(x) \in \mathcal{C}_T \\ \downarrow & \text{otherwise} \end{cases}$;
2. $\mathbf{C}[A](\sigma) = \mathbf{A}[A]$, where $\mathbf{A}[-] : \text{Atoms} \rightarrow \mathcal{C}_T$ is the semantics of primitive components;
3. $\mathbf{C}[C_0 C_1](\sigma) = \begin{cases} \mathbf{C}[C_0](\sigma) \wedge_T \mathbf{C}[C_1](\sigma) & \text{if } \mathbf{C}[C_i](\sigma) \in \mathcal{C}_T, \text{ for } i \in \{0, 1\} \\ \downarrow & \text{otherwise} \end{cases}$;
4. $\mathbf{C}[\{C : P\}](\sigma) = \begin{cases} \top_T & \text{if } \text{min}\mathbf{P}[P](\sigma) \text{ is empty or infinite} \\ C_1 \wedge_T \dots \wedge_T C_k & \text{if } \text{min}\mathbf{P}[P](\sigma) = \{\sigma_1, \dots, \sigma_k\} \neq \emptyset, \mathbf{C}[C](\sigma_i) = C_i \in \mathcal{C}_T \\ \downarrow & \text{otherwise} \end{cases}$

$$5. \mathbf{C}[\langle D \rangle(U)](\sigma) = \begin{cases} \mathbf{D}[\langle D \rangle(\sigma)(\mathbf{L}[L](\sigma), \mathbf{U}[U](\sigma))] & \text{if defined} \\ \dagger & \text{otherwise} \end{cases}$$

The semantics of component definitions is a map $\mathbf{D}[-] : \text{Definitions} \rightarrow (\mathcal{D} \cup \{\dagger\})^\Sigma$, where Definitions is the set of all parse trees with root D . For a scope $\sigma \in \Sigma$, we define $\mathbf{D}[-](\sigma)$ as follows:

1. $\mathbf{D}[\langle V \rangle](\sigma) = \begin{cases} \sigma(\mathbf{V}[V](\sigma)) & \text{if } \mathbf{V}[V](\sigma) = x \in \text{dom}(\sigma) \text{ and } \sigma(x) \in \mathcal{D}; \\ \dagger & \text{otherwise} \end{cases}$;
2. If D is a component $\langle U_0 \rangle(U_1)\{C\}$, then for an array of parameter values $t \in \mathcal{V}^\square$ and an array of nodes $q \in \mathcal{N}^\square$, we define $\mathbf{D}[\langle U_0 \rangle(U_1)\{C\}](\sigma)(t, q)$ as follows: Recall from Section 3 that the number of parameters and nodes can implicitly define variables. Suppose that there exists a unique ‘index-defining’ scope $\sigma' \in \Sigma$ such that for $m = \text{len}(t)$ and $n = \text{len}(q)$. Then we have
 - (a) $\mathbf{U}[U_0](\sigma') = [s_1, \dots, s_m] \neq \dagger$ satisfies $s_i \simeq t(i)$, for all $1 \leq i \leq m$;
 - (b) $\mathbf{U}[U_1](\sigma') = [p_1, \dots, p_n] \neq \dagger$ satisfies $p_i \simeq q(i)$, for all $1 \leq i \leq n$;
 - (c) $\text{flatten}([s_1, \dots, s_m, p_1, \dots, p_n]) \in \mathcal{N}^\square$ has no duplicates;
 - (d) $\text{dom}(\sigma') \subseteq \mathcal{N}$ is minimal such that properties (a)-(c) are satisfied.

We evaluate the body C of the component definition to the component $\mathbf{C}[C](\sigma\sigma')$, where $\sigma\sigma'$ is the composition of σ and σ' . Define the map $r : \text{supp}(\mathbf{C}[C](\sigma\sigma')) \rightarrow \mathcal{N}$ as

$$r(x) = \begin{cases} t_i(k_1) \cdots (k_l) & \text{if } x = s_i(k_1) \cdots (k_l) \\ q_i(k_1) \cdots (k_l) & \text{if } x = p_i(k_1) \cdots (k_l) \\ v \text{ fresh} & \text{otherwise} \end{cases}$$

Map r is well-defined, because $\text{flatten}([s_1, \dots, s_m, p_1, \dots, p_n]) \in \mathcal{N}^\square$ has no duplicates. Note that r is finite, since we assume that $\text{supp}(\mathbf{C}[C](\sigma\sigma'))$ is finite. We define $\mathbf{D}[\langle U_0 \rangle(U_1)\{C\}](\sigma)(t, q)$ as the simultaneous substitutions $\mathbf{C}[C](\sigma\sigma')[r(x)/x : x \in \text{dom}(f)]$. If such ‘index-defining’ scope σ' does not exist or is not unique, then we simply define $\mathbf{D}[\langle U_0 \rangle(U_1)\{C\}](\sigma)(t, q) = \dagger$.

We define the semantics of files as a map $\mathbf{K}[-] : \text{Files} \rightarrow \Sigma \cup \{\dagger\}$, where Files is the set of parse trees with root K . Let $\mathbf{I}[-] : \text{Imports} \rightarrow \Sigma$ be the semantics of imports. For a scope $\sigma \in \Sigma$, we define $\mathbf{K}[-](\sigma)$ inductively as follows:

1. $\mathbf{K}[\langle I \rangle](\sigma) = \mathbf{I}[\langle I \rangle]$;
2. $\mathbf{K}[\langle KND \rangle](\sigma) = \begin{cases} \sigma_0\{x \mapsto c\} & \text{if } \sigma_0 = \mathbf{K}[\langle K \rangle](\sigma) \neq \dagger, x = \mathbf{N}[\langle N \rangle], \text{ and } c = \mathbf{D}[\langle D \rangle](\sigma_0) \neq \dagger \\ \dagger & \text{otherwise} \end{cases}$.

6 Input/output nodes

As mentioned in Section 3, nodes of primitive component definitions require input/output annotations. Treo regards such port type annotations as attributes of the primitive component. For a semantic sort T , we model the input nodes and output nodes of its instances via two maps $I, O : \mathcal{C}_T \rightarrow 2^{\mathcal{N}}$ satisfying $\text{supp}(C) = I(C) \cup O(C)$, for all $C \in \mathcal{C}_T$. If $x \in I(C) \cap O(C)$, then we call x a *mixed* node.

Example 6.1 (Mixed nodes). Recall the I/O component sort from Example 4.3. Let $P_1 = (C_1, \{x\}, \{y\})$, $P_2 = (C_2, \{y\}, \emptyset)$, and $P_3 = (C_3, \{z\}, \{y\})$ be three primitive I/O components. Figure 2(a) shows a graphical representation of composition of P_1 , P_2 , and P_3 . In this figure, an arrow from a node a to a component




Figure 2: Surgery on an I/O-component to remove mixed nodes.

P indicates that a is an input node of P . An arrow from a component P to a node a indicates that a is an output node of P . Node y is an output node of P_1 and P_3 , and it is an input node of P_2 . Thus, y is a mixed node in the composition $P_1 \cdot P_2 \cdot P_3$, where \cdot is sequential composition of I/O components. \triangle

Most semantic sorts that distinguish input and output nodes assume *well-formed* compositions: each shared node in a composition is an output of one component and an input of the other.

Definition 6.1 (Well-formedness). A composition $C_1 \wedge_T \dots \wedge_T C_n$, with $n \geq 0$, is well-formed if and only if $|\{i \in \{1, \dots, n\} \mid x \in I(C_i)\}| \leq 1$ and $|\{i \in \{1, \dots, n\} \mid x \in O(C_i)\}| \leq 1$, for all $x \in \mathcal{N}$.

For well-formed compositions, the behavior of the composition naturally corresponds to the composition of Reo connectors. However, specification of complex components as well-formed compositions is quite cumbersome, because it requires explicit verbose expression of the ‘merge-replicate’ behavior of every Reo node in terms of a suitable number of binary mergers and replicators. Reo nodes abstract from such detail and yield more concise specifications. Like Reo, Treo does not impose any restriction on the nodes of constituent components in a composition. Indeed, the denotational semantics of components $\mathbf{C}[-]$ in Section 5 unconditionally computes the composition. To define the semantics of $\mathbf{C}[-]$ for a semantic sort T where \wedge_T requires well-formedness, parsing a (non-well-formed) Treo composition needs the degree (i.e., the number of coincident input and output channel ends) of each node to correctly express the ‘merge-replicate’ semantics of that node. The degree of every node used in a definition can be known only at the end of that definition. The Treo compiler could discover the degree of every node via two-pass parsing.

Alternatively, Treo can delay applying composition \wedge_T in T until parsing completes, Treo accomplishes this by interpreting a Treo program over the I/O-component sort IO_T , as defined in Example 4.3, wherein compositions consist of lists of primitive components. First, Treo wraps each primitive component $C \in \mathcal{C}_T$ within a primitive I/O-component $(C, I(C), O(C)) \in \mathcal{P}_T$. Using Section 5, Treo parses the Treo program over the semantic sort IO_T as usual, and obtains a single I/O-component $P_1 \dots P_n \in \text{IO}_T$.

However, the resulting composition $P_1 \dots P_n$ may not be well-formed. Therefore, the Treo compiler applies some surgery on $P_1 \dots P_n$ to ensure a well-formed composition. This surgery consists of splitting all shared nodes in X , and reconnecting them by inserting a *node component*. We model these node components (over semantic sort T) as a map node: $(2^{\mathcal{N}})^2 \times \mathcal{N} \rightarrow \mathcal{C}_T$. For sets of names $I, O \subseteq \mathcal{N}$ and a default name $x \in \mathcal{N}$, the component $\text{node}(I, O, x) \in \mathcal{C}_T$ has input nodes I (or $\{x\}$, if I is empty) and output nodes O (or $\{x\}$, if O is empty).

Definition 6.2 (Surgery). The surgery map $\text{surg} : \text{IO}_T \rightarrow \text{IO}_T$ is defined as $\text{surg}(P_1 \dots P_n) = P'_1 \dots P'_n \cdot \prod_{x \in \text{supp}(P_1 \dots P_n)} N_x$, where $P'_i = P_i[x_i/x : x \in \text{supp}(P_i)]$, for all $1 \leq i \leq n$, and $N_x = (\text{node}(I_x, O_x, x), I_x, O_x)$, with $I_x = \{x_i \mid x \in O(P_i)\}$ and $O_x = \{x_i \mid x \in I(P_i)\}$. The composition \prod is ordered arbitrarily.

Intuitively, the surgery map takes a possibly non-well-formed composition and produces a well-formed composition by inserting node components. Although initially, multiple components may pro-

duce output at the same node. After applying the surgery map, these components offer data for the same node component via different ‘ports’.

Example 6.2 (Surgery). Figure 2(b) shows the result of applying the surgery map to the I/O-component $P_1 \cdot P_2 \cdot P_3$ from Example 6.1. The surgery map consists of two parts. First, the surgery map splits every node $a \in \{x, y, z\}$ by renaming a to a_i in P_i , for every $1 \leq i \leq n$. Second, the surgery map inserts at every node $a \in \{x, y, z\}$ a node component N_a . Clearly, $\text{surg}(P_1 \cdot P_2 \cdot P_3)$ is a well-formed composition. \triangle

7 Related work

The Treo syntax offers a textual representation for the graphical Reo language [1, 2]. We propose Treo as a syntax for Reo that (1) provides support for parameterization, recursion, iteration, and conditional construction; (2) implements basic design principles of Reo more closely than existing languages; and (3) reflects its declarative nature. The graphical Reo editor implemented as an Eclipse plugin [8] does not support parameterization, recursion, iteration, or conditional construction. RSL (with CARML for primitives) [3, 19] is imperative, while Reo is declarative. FOCAML [12], supports only constraint automata [4], while Treo allows arbitrary user-defined semantic sorts for expressing the behavior of Reo primitives.

Since Treo leaves the syntax for primitive subsystems (i.e., semantic sorts) as user-defined, Treo is a “meta-language” that specifies compositional construction of complex structures (using the common core language defined in this paper) out of primitives defined in its arbitrary, user-defined sub-languages. As such, Treo is not directly comparable to any existing language. We can, however, compare the component-based system composition of Treo with the system composition of an existing language.

Treo components are similar to proctype declarations in Promela, the input language for the SPIN model checker developed by Holzmann [11]. However, the focus of Promela is on imperative definitions of processes, while Treo is designed for declarative composition of processes.

SysML is a graphical language for specification of systems [9]. SysML offers 9 types of diagrams, including activity diagrams and block diagrams. Each diagram provides a different view on the same system [20]. Diagram types in SysML are comparable to semantic sorts in Treo. The main difference between the two, however, is that Treo requires a well-defined composition operator, using which it allows construction of more complex components, while diagram composition is much less prominent in SysML.

A component model is a programming paradigm based on components and their composition. Our Treo language can be viewed as one such component model with a concrete syntax. Over the past decades, many different component models have been proposed. For example, CORBA [21] is a component model that is flat in the sense that every CORBA component is viewed as a black box, i.e., it does not support composite components. Fractal [5] is an example of a component model that is hierarchical, which means a component can be a composition of subcomponents. Concrete instances of Fractal consist of libraries (API’s) for a variety of programming languages, such as Java, C, and OMG IDL [5]. Treo components and Fractal component differ with respect to interaction: Treo components interact via shared names, while Fractal component interact via explicit bindings.

8 Conclusion

We propose Treo as a textual syntax for Reo connectors that allows user-defined semantic sorts, and incorporates Reo’s predefined node behavior. These features are not present in any of the existing alter-

native languages for Reo. We provided an abstract syntax for Treo and its denotational semantics based on this abstract syntax. We identify three possible directions for future work.

First, since our semantics disallows recursion, a component in Treo is currently restricted to consist of a composition of finitely many subsystems. Consequently, we cannot, for instance, express the construction of a primitive with an unbounded buffer, B_ω , from a set of primitives with buffer capacity of one, B_1 . It seems, however, possible to use simulation and recursion to define B_ω in terms of B_1 : B_ω is the smallest (with respect to simulation) component that simulates B_1 and is stable under sequential composition with B_1 . These assumptions readily imply that B_ω simulates a primitive with buffer of arbitrary large capacity. Semantically, the unbounded buffer would then be defined as a least fixed point of a certain operator on components. An extension of Treo semantics that allows such fixed point definitions would provide a powerful tool to define complex ‘dynamic’ components.

Second, the current semantics in Section 5 does not support components with an identity. If we instantiate a component definition twice with the same parameters, we obtain two instances of the same component. Ideally, component instantiation should return a component instance with a fresh identity. Allowing components with identities in Treo enables programmers to design systems more realistically.

Finally, a semantic sort T from Definition 4.1 consists of a single composition operator \wedge_T . Generally, a semantic sort consists of multiple composition operators (each with its own arity). For example, we may need both sequential composition as well as parallel composition. Extending Treo with (a variable number of) composition operators would enable users to model virtually all semantic sorts.

References

- [1] Farhad Arbab (2004): *Reo: a channel-based coordination model for component composition*. *Mathematical Structures in Computer Science* 14(3), pp. 329–366, doi:10.1017/S0960129504004153.
- [2] Farhad Arbab (2011): *Puff, The Magic Protocol*. In Gul Agha, Olivier Danvy & José Meseguer, editors: *Formal Modeling: Actors, Open Systems, Biological Systems - Essays Dedicated to Carolyn Talcott on the Occasion of Her 70th Birthday*, *Lecture Notes in Computer Science* 7000, Springer, pp. 169–206, doi:10.1007/978-3-642-24933-4_9.
- [3] Christel Baier, Tobias Blechmann, Joachim Klein & Sascha Klüppelholz (2009): *A Uniform Framework for Modeling and Verifying Components and Connectors*. In John Field & Vasco Thudichum Vasconcelos, editors: *Proceedings of COORDINATION 2009*, *Lecture Notes in Computer Science* 5521, Springer, pp. 247–267, doi:10.1007/978-3-642-02053-7_13.
- [4] Christel Baier, Marjan Sirjani, Farhad Arbab & Jan J. M. M. Rutten (2006): *Modeling component connectors in Reo by constraint automata*. *Sci. Comput. Program.* 61(2), pp. 75–113, doi:10.1016/j.scico.2005.10.008.
- [5] Eric Bruneton, Thierry Coupaye, Matthieu Leclercq, Vivien Quéma & Jean-Bernard Stefani (2006): *The FRACTAL component model and its support in Java*. *Softw., Pract. Exper.* 36(11-12), pp. 1257–1284, doi:10.1002/spe.767.
- [6] Philipp Chrszon, Clemens Dubslaff, Christel Baier, Joachim Klein & Sascha Klüppelholz (2016): *Modeling Role-Based Systems with Exogenous Coordination*. In Erika Ábrahám, Marcello M. Bonsangue & Einar Broch Johnsen, editors: *Theory and Practice of Formal Methods - Essays Dedicated to Frank de Boer on the Occasion of His 60th Birthday*, *Lecture Notes in Computer Science* 9660, Springer, pp. 122–139, doi:10.1007/978-3-319-30734-3_10.
- [7] Kasper Dokter & Farhad Arbab (2018): *Rule-Based Form for Stream Constraints*. In: *Proceedings of COORDINATION 2018*. To appear.
- [8] *Extensible Coordination Tools ECT*. <http://reo.project.cwi.nl>. Accessed: 2018-03-23.

- [9] Sanford Friedenthal, Alan Moore & Rick Steiner (2014): *A practical guide to SysML: the systems modeling language*. Morgan Kaufmann.
- [10] C. A. R. Hoare (1978): *Communicating Sequential Processes*. *Commun. ACM* 21(8), pp. 666–677, doi:10.1145/359576.359585.
- [11] Gerard J. Holzmann (2004): *The SPIN Model Checker - primer and reference manual*. Addison-Wesley.
- [12] S.-S. T. Q. Jongmans (2016): *Automata-theoretic protocol programming*. Ph.D. thesis, Leiden University.
- [13] Sung-Shik T. Q. Jongmans & Farhad Arbab (2012): *Overview of Thirty Semantic Formalisms for Reo*. *Sci. Ann. Comp. Sci.* 22(1), pp. 201–251, doi:10.7561/SACS.2012.1.201.
- [14] Sung-Shik T. Q. Jongmans & Farhad Arbab (2015): *Take Command of Your Constraints!* In Tom Holvoet & Mirko Viroli, editors: *Proceedings of COORDINATION 2015, Lecture Notes in Computer Science* 9037, Springer, pp. 117–132, doi:10.1007/978-3-319-19282-6_8.
- [15] Sung-Shik T. Q. Jongmans & Farhad Arbab (2016): *Data optimizations for constraint automata*. *Logical Methods in Computer Science* 12(3), doi:10.2168/LMCS-12(3:11)2016.
- [16] Sung-Shik T. Q. Jongmans & Farhad Arbab (2016): *PrDK: Protocol Programming with Automata*. In Marsha Chechik & Jean-François Raskin, editors: *Proceedings of TACAS 2016, Lecture Notes in Computer Science* 9636, Springer, pp. 547–552, doi:10.1007/978-3-662-49674-9_33.
- [17] Sung-Shik T. Q. Jongmans & Farhad Arbab (2018): *Centralized coordination vs. partially-distributed coordination with Reo and constraint automata*. *Sci. Comput. Program.* 160, pp. 48–77, doi:10.1016/j.scico.2017.06.004.
- [18] Sung-Shik T. Q. Jongmans, Sean Halle & Farhad Arbab (2014): *Automata-Based Optimization of Interaction Protocols for Scalable Multicore Platforms*. In Eva Kühn & Rosario Pugliese, editors: *Proceedings of COORDINATION 2014, Lecture Notes in Computer Science* 8459, Springer, pp. 65–82, doi:10.1007/978-3-662-43376-8_5.
- [19] Sascha Klüppelholz (2012): *Verification of Branching-Time and Alternating-Time Properties for Exogenous Coordination Models*. Ph.D. thesis, Dresden University of Technology. Available at [http://www.qucosa.de/recherche/frontdoor/?tx_slubopus4frontend\[id\]=8621](http://www.qucosa.de/recherche/frontdoor/?tx_slubopus4frontend[id]=8621).
- [20] Philippe Kruchten (1995): *The 4+1 View Model of Architecture*. *IEEE Software* 12(6), pp. 42–50, doi:10.1109/52.469759.
- [21] OMG (2006): *CORBA Component Model, v4.0, OMG document formal/06-04-01*. Available at <http://www.omg.org/spec/CCM>. Visited 11-13-2017.
- [22] Terence Parr (2013): *The definitive ANTLR 4 reference*. Pragmatic Bookshelf.
- [23] *ReoLanguage GitHub repository*. <https://github.com/ReoLanguage/Reo>. Accessed: 2018-03-23.
- [24] David A. Schmidt (1986): *Denotational Semantics: A Methodology for Language Development*. William C. Brown Publishers, Dubuque, IA, USA.