

RESEDA: Declaring Live Event-driven Computations as REactive SEmi-structured DATA

João Costa Seco NOVA LINCS FCT, NOVA University of Lisbon Lisbon, Portugal joao.seco@fct.unl.pt	Søren Debois Dept. of Computer Science IT University of Copenhagen Copenhagen, Denmark debois@itu.dk	Thomas T. Hildebrandt Dept. of Computer Science University of Copenhagen Copenhagen, Denmark hilde@di.ku.dk	Tijs Slaats Dept. of Computer Science University of Copenhagen Copenhagen, Denmark slaats@di.ku.dk
---	--	---	--

Abstract—Enterprise computing applications generally consists of several inter-related business processes linked together via shared data objects and events. We address the open challenge of providing formal modelling and implementation techniques for such enterprise computing applications, introducing the declarative, data-centric and event-driven process language RESEDA for REactive SEmi-structured DATA. The language is inspired by the computational model of spreadsheets and recent advances in declarative business process modelling notations. The key idea is to associate either input events or reactive computation events to the individual elements of semi-structured data and declare reactive behaviour as explicit reaction rules and constraints between these events. Moreover, RESEDA comes with a formal operational semantics given as rewrite rules supporting both formal analysis and persistent execution of the application as sequences of rewrites of the data. The data, along with the set of constraints, thereby at the same time constitutes the specification of the data, its behaviour and the run-time execution component. This key contribution of the paper is to introduce the RESEDA language, its formal execution semantics and give a sufficient condition for liveness of programs. We also establish Turing-equivalence of the language independently of the choice of underlying data expressions and exemplify the use of RESEDA by a running example of an online store. A prototype implementation of RESEDA and the examples of the paper are available on-line at <http://dcr.tools/reseda>.

Index Terms—declarative, event-driven computations, semi-structured data, liveness, enterprise applications, business processes

I. INTRODUCTION

Realistic enterprise computing applications contain several inter-related business processes defining the behaviour of the application and which are often linked together via shared data, such as customers, orders and shipments in a shop. The data and the activities of the processes are two sides of the same coin: Data classes have corresponding activities in the processes for the creation and change of the data objects. Moreover, cardinality relations between data elements in the data model are mirrored by behavioural orderings of activities in the processes. For instance, the data model may specify that a shipment is related to one or more orders belonging to the

same customer. This is reflected by a behavioural constraint, that the activity for creating a shipment can only be executed after the activities of creating a customer and the first order of the customer.

The management of customers, orders and shipments of orders are likely to be independently defined and developed business processes, which are not necessarily hierarchically embedded in each other. However, as pointed out in [19], [21], standard business process notations such as BPMN [13] are not ideal for modelling nor implementing enterprise business applications with dependencies and shared data between activities belonging to multiple instances of such independent processes. The same holds for standard case management notations such as CMMN [14].

We propose in the present paper a novel, declarative data-centric process language RESEDA based on REactive SEmi-structured DATA as a unified specification of data, behaviour and run-time execution state. RESEDA is inspired by and generalises both the declarative Dynamic Condition Response (DCR) graph process notation [4], [5], [10] and the widely used reactive model of data and behaviour used in spreadsheets [17].

In a spreadsheet, data is organised in a matrix. The fields of the matrix are either empty and expect a data element to be *input* by the user or contain an expression, possibly referring to data of other fields, and expected to be *computed* by the spreadsheet processor. The computation of an expression happens *initially*, when the expression is defined and *reactively*, whenever a data element referenced by the expression changes its value.

In RESEDA we categorise data elements as either input or computation data elements similarly to spreadsheets, but instead of the matrix structure used in spreadsheets, RESEDA employs a semi-structured model for the data elements. Semi-structured data is widely used in practice, notably via XML and JSON, and allows more easily than matrices the modelling of inter-related, unbounded and dynamically growing and shrinking collections of data elements using the containment relation, attributes with unique ids (UIDs) and locational referencing of data elements via so-called path expressions.

Input data elements in RESEDA are thus data elements, for which the assignment of a value is triggered by an input event

from the outside context of the process. In addition to an empty field of a spreadsheet, other analogies for such input data elements can be found in input fields of web forms, input parameters of web-services, or even IoT sensors e.g. measuring the location of a business object using a geo-location capability. Similarly to input fields in a web form being either required or optional, RESEDA allows for data elements to be declared as either required or optional. A required input data element must eventually be provided a value for the program to reach an accepting state. Similarly, a required computation data element must eventually be evaluated by the run-time system, for the program to reach an accepting state. Also, we allow for the declaration of additional rules, e.g. for dynamically creating new data and determining additional ordering constraints, e.g. that a new input or re-computation of a value for a data element is required as a response to another input or computation. A computation of a process is then only considered complete if every required input or computation of a data value is eventually executed.

The execution semantics of RESEDA makes the language directly applicable for execution as a sequence of rewrites of the data. The semantic foundation for describing both the reactive computation of expressions and the creation of new data obtained from the declarative DCR* process notation introduced in [4], [5] and generalising DCR graphs [10] to allow dynamic creation of sub processes. The creation of new data (with its own behavioural constraints and reaction rules) is represented directly by the DCR* rule for dynamic spawning of new sub processes. RESEDA can thereby also be seen as an extension of DCR* to allow for associating events with semi-structured data and expressions, and using path expressions to indirectly refer to other events by their location instead of only their identity. Our approach in this way harmoniously integrates the dynamic construction of semi-structured data and concurrent declarative control flow.

The concurrent event-driven reactive computation raises issues of possible deadlocks and livelocks, as an event may in response lead to an infinite sequence of re-computations and creations of new data elements. As a main technical result, we provide sufficient conditions for deadlock and livelock freedom. In addition, we prove that Reseda generalises the DCR graphs process formalism, [4], [5] and derive Turing completeness of the language as a corollary.

For want of space, proofs and trivial or tedious definitions have been relegated to the extended on-line version of this paper, available at [16].

II. RESEDA

We illustrate and motivate the constructs of our language, RESEDA, by gradually designing the shopping process introduced in [22].

The data model shown in Figure 1 describes (using UML notation) classes of customers, orders, orderlines, products and deliveries. The behavioural model shown in Figure 2 describes behavioural constraints between key activities of the business process using a variant of DECLARE [15] notation used in

the so-called Object-Centric Behavioural Constraint (OCBC) notation of [21], [22]. Briefly, the process describes a shop in which registered customers can create orders (by the activity `create order` (1), to which they must add one or more order lines (by the activity `pick item` (2), each referring to a unique product. Each product mentioned in a order line must be wrapped (by the activity `wrap item` (3)) and eventually delivered to the customer (by the activity `deliver items` (4)). A delivery may contain items from different orders but restricted to only one customer.

The relations (5) and (6) from `create order` to `pick item` in Figure 2 define behavioural constraints denoting respectively that any `create order` event is followed by one or more `pick item` events, and that every `pick item` event is preceded by a unique `create order` event. The relation (7) between `pick item` and `wrap item` denote a one-to-one correspondence between events related to the two activities. Finally, the relations (8) and (9) define respectively that every `wrap item` must be followed by exactly one `deliver items` event, and each `deliver items` event must be preceded by one or more `wrap item` events.

The OCBC notation of [21], [22] makes it possible to combine the data model of Figure 1 and the declarative behavioural model of Figure 2. But OCBC does not provide notation for specifying the run-time data objects and process instances. In the following we describe a RESEDA process doing exactly that. We represent data items as named leaf elements of the tree-structured semi-structured data, and model the structure of the data model either by regular nesting or by linking using unique identifiers like in document bases databases.

In this section we use the concrete RESEDA syntax, used in our implementation. However, in the formal developments ahead we define a core calculus, using a slightly different, more abstract notation, for the sake of readability of the results. We invite the reader to experiment with the syntax and semantics of RESEDA in the on-line tool accompanying this paper, available at <http://dcr.itu.dk/reseda/>. For convenience, all major examples of the present section are pre-loaded in that tool for experimentation.

A. Static data

We begin by showing how to model static data with no behaviour in Example 1. Concretely, we model two data objects of class Customer of Figure 1.

Example 1 (Modelling structured data using constant values).

```

1 customer[] {
2   customer_id(0) [],
3   customer_name('John') []
4 },
5 customer[] {
6   customer_id(1) [],
7   customer_name('Mary') []
8 }
```

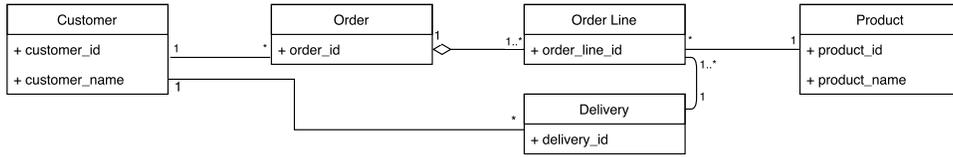


Fig. 1. Data model of the shopping example from [22]

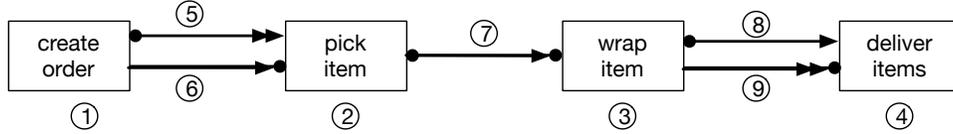


Fig. 2. Behavioural constraints of the shopping example from [22]

In Reseda, the basic building block are instances of data elements, with the general form, $M \text{ name}(V) [E]P$, where name defines a label for the data element, V defines the current value for the element, E is an expression that defines the value for the element if computed, M defines the behavioural state of the data element (to be explained later), and P defines inner (nested) data elements and control elements (behavioural constraints and reaction rules, described ahead). The element label is the only required part of a data element specification.

In this example, we model two data elements with label `customer`, which contain data elements modelling the fields `customer_id` and `customer_name`. The structure of data items and associated inner data items and values can be interpreted as an object structure found in document- or XML-based databases. Notice that the data element `customer` has structuring purposes only — intuitively it does not have a value nor expression for computing one, formally we model it holding the value unit (1) and omitted from the specification.

B. Dynamically created data

The key mechanism for obtaining dynamic behaviour in RESEDA is that whenever the value of a data element is updated, that update may have side effects triggered by reaction rules inherited from the DCR* process language [5]. A rule can e.g. specify the dynamic creation of new data elements or requiring other data elements to be updated in response.

We next illustrate how to declare a rule for dynamically creating instances of the class `Customer` of Figure 1.

Example 2 (Dynamically created customers collection).

```

1  customers[]{
2    create_customer[?]
3    ;
4    create_customer -->> {
5      customer[]{
6        !customer_id[freshid()],
7        !customer_name[?:string]
8      }
9    }
10 }
```

The dynamically created collection of customers is enclosed by a `customers[]` data element, that nests a list of data elements, like the ones defined in Example 1. Initially, the

element contain no such customer data elements. Instead, we have introduced an input data element `create_customer`. When the corresponding input event happens, an additional `customer` data element is created as a side effect. This side-effect is specified via a rule using the “spawn relation” (\rightarrow) introduced in [5]. Syntactically, we distinguish between the part of a process that specifies data elements, and the part that specifies rules by a `;` symbol.

The reaction rule declares that, whenever the left operand — a data item named `create_customer` — is provided a value, a new data item named `customer` is created in the local context, i.e. inside the collection `customers`. When no data type is provided, as for the `create_customer` data item, we assume it is of unit type, which in the analogy of the web forms would just be a button being clicked. The value of the computation data element `customer_id` inside each new `customer` data element is given by the predefined function `fresh_id` similar to the `generate_id` function in XPath [7]. The input data item `customer_name` gets its value of string type from the outside context. The result of executing a `create customer` event will thus be the RESEDA program shown in Example 3

Example 3 (Freshly created customer instance).

```

1  customers[]{
2    create_customer[?],
3    customer[]{
4      !customer_id[freshid()],
5      !customer_name[?:string]
6    }
7  ;
8  create_customer -->> {
9    customer[]{
10     !customer_id[freshid()],
11     !customer_name[?:string]
12   }
13 }
14 }
```

Notice that for any newly created `customer` record, the computation data element `customer_id` needs to be updated and the input data element `customer_name` element is missing a data value. We indicate that the events providing values to the data elements inside a new `customer` data element are initially *pending* as in DCR Graphs, by the exclamation mark

! used as prefix. A pending event is an event that is required to happen eventually. This means, for computation data elements like `customer_id`, that they are staged to be automatically computed. The corresponding computation event will happen unless it depends on some other data elements for which values have not yet been computed or provided, or other behavioural constraints have been defined as exemplified in later examples. Since there are no dependencies for the computation event `customer_id`, the event will be triggered by the RESEDA processor, thereby computing a fresh id and assigning it as value to the data element, resulting in the RESEDA program shown in Example 4.

Example 4 (Customer id computed).

```

1  customers[] {
2    create_customer[?],
3    customer[] {
4      customer_id(0)[freshid()],
5      !customer_name[?:string]
6    }
7  ;
8  create_customer -->> {
9    customer[] {
10     !customer_id[freshid()],
11     !customer_name[?:string]
12   }
13 }
14 }
```

For input data elements, the pending state means that the input should be done *eventually* for the process to progress. Using the analogy of web forms, this would correspond to an input field being marked as mandatory.

The result of executing a `customer_name` event providing the string `John` as value will be the RESEDA program shown in Example 5

Example 5 (Customer name provided as input event).

```

1  customers[] {
2    create_customer[?],
3    customer[] {
4      customer_id(0)[freshid()],
5      customer_name('John')[?:string]
6    }
7  ;
8  create_customer -->> {
9    customer[] {
10     !customer_id[freshid()],
11     !customer_name[?:string]
12   }
13 }
14 }
```

Note that a new `customer_name` input event can happen again, analogous to a web form field being editable, but not mandatory.

The same structure can be followed to create the collection corresponding to class `Product` depicted in Figure 1, by the code of Example 6.

Example 6 (Dynamically created products and trigger path expression).

```

1  products[] {
2    create_product[?:string]
3  ;
4  create_product -->> {
5    product[] {
6      !product_id[freshid()],
```

```

7      !product_name[@trigger:value]
8    }
9  }
10 }
```

Here we also illustrate how to access the data element whose update caused the dynamic creation of data using the reserved word `trigger` in a *path expression* (identified by the prefix `@`) for a computation data element `product_name`. We get the value(s) of the element(s) identified by the path expression by adding the suffix `:value`.

The result of executing a `create_product` event providing the string `iPhone X` as value will be the RESEDA program shown in Example 7. Note that the two computation data elements are pending and will thus the corresponding computation events subsequently be computed and assign values to the data elements.

Example 7 (Freshly created iPhone X product).

```

1  products[] {
2    create_product[?:string],
3    product[] {
4      !product_id[freshid()],
5      !product_name['iPhone X']
6    }
7  ;
8  create_product -->> {
9    product[] {
10     !product_id[freshid()],
11     !product_name[@trigger:value]
12   }
13 }
14 }
```

We can observe by now that RESEDA expressions represent both the state and code of the system, which relates to live programming and execution environments such as `smalltalk` virtual machines.

C. Regulating behaviour: Condition and exclusion rules

Next, to define the collection of orders we must establish relations between classes `Order`, `Customer` and `Product` (via class `OrderLine`) as depicted in Figure 1. Consider Example 8 below introducing data element `orders` with an input data element `create_order`, used to introduce new elements in the collection.

Example 8 (Orders collection definition).

```

1  orders[] {
2    create_order[
3      ?:@/customers/customer/customer_id:value
4    ]
5  ;
6  create_order -->> {
7    order[] {
8      !order_id[freshid()],
9      !customer_id[@trigger:value],
10     !pick_item[
11       ?:@/products/product/product_id:value
12     ]
13   }
14 ;
15 pick_item -->> {
16   order_line[] {
17     !order_line_id[freshid()],
18     !item[@trigger:value],
19     !wrap_item[
20       ?:@/deliveries/
         delivery[
```

```

21         not (deliver_items:executed)
22         ]/delivery_id:value
23     ]
24     ;
25     wrap_item -->* order_line_id,
26     item -->* order_line_id,
27     wrap_item -->% wrap_item
28 }
29 }
30 }
31 }
32 }

```

Observe the type annotation for the input of data element `create_order` in line 2–4, given by the path expression `@/customers/customer/customer_id:value`. The expression is evaluated with relation to its definition context to a set of data items and dereferenced by the attribute label value to obtain all values carried by the selected data items. The initial `/` of the path expression means that the expression is evaluated at the root of the process.

This kind of annotations in input fields thereby acts as a *dependent type* [12] for the input, with the meaning, that only values in the computed set are accepted as valid input value for the data element `create_order`. Further, the input event updating the data element is only possible if there is at least one valid data value. The type in this case thereby ensures that a valid customer id is provided as input for the create order data element. In the analogy of the web form, this would correspond to a data-dependent input validation of the input field or the values in a dropdown UI widget.

As in the previous examples, the unique identifier of an order is automatically computed using the `freshid()` expression in line 8, and the customer explicitly identified by the input value for the `create_order` event is copied to a locally available computation data element (`customer_id`) using the `@trigger:value` path expression in line 9. We thereby establish the relation between a customer and an order.

The pending input data element `pick_item` in line 10–12 allows items to be added to a created order. The fact that it is initially pending means that at least one item must eventually be added to the order. Combined with the fact that the input data element is created with the new order we fulfil the constraints (5) and (6) in Figure 2. Again, every time a value is provided as input for the data element `pick_item`, by the user using a web form or an external service, the dependent type in line 11 guarantees that the value is a valid product identifier value.

Now, the rule for dynamic creation of data (line 14–28) triggered by the input event for the `pick_item` data element specifies that an `order_line` data element is created inside the structure of the current order. Using the `@trigger` path expression in line 17, the product identifier provided as value to the `pick_item` is copied to the `item` data element of the order line and we thereby establish the relation between an order line and a product. The `wrap_item` input data element in line 18 takes an input value which is again constrained by a path expression. In this case, the path expression constrains the values to be delivery identifiers for deliveries, for which the

event corresponding to the `deliver_items` data element has not yet been executed. This implements that the wrapped item is associated with a delivery which has not yet been delivered. We will look into the details of the delivery data class and behaviour in the next example.

Notice the nested structure created here to represent the containment relation of objects from classes `Order` and `OrderLine` from Figure 1 and the use of identifiers to represent relations from class `Order` to class `Customer` and from class `OrderLine` to the classes `Product` and `Delivery`.

The example introduces two new kinds of behavioural rules in line 25–27 in the definition of `order_line`, inherited from DCR Graphs. The first two behavioural rules (line 25–26) are standard condition constraint rules, as also found in DECLARE, specifying that the `order_line_id` computation data element can only be updated after `wrap_item` and `item` have been updated. As we will see in the next example, the reason for this constraint is that the update of the `order_line_id` has the side effect in the `delivery` class that the item is added to the delivery identified by the `wrap_item` data element. For this reason we want it only to happen once all data elements has been computed. Finally, the rule `wrap_item -->% wrap_item` (line 27), causes the data element `wrap_item` to be dynamically excluded from future computation once a value has been provided to it. In the analogy of the web form, the button will no longer be visible, but its data can still be accessed by the code. This ensures that the `wrap_item` input event can only happen once. Combined with the fact that the data element is initially pending implements constraint (7) in Figure 2.

D. Initially excluded data elements aka hidden fields

We now explain how to specify the collection of `delivery` data elements corresponding to the class `Delivery` from Figure 1. Refer to Example 9 below.

Example 9 (Deliveries collection definition).

```

1 deliveries[] {
2   create_delivery[?:
3     @/customers/customer/customer_id:value
4   ]
5   ;
6   create_delivery -->> {
7     delivery[] {
8       !delivery_id[freshid()],
9       !customer_id[@trigger:value],
10      !%deliver_items[?]
11     ;
12     /orders/order[
13       @customer_id:value==@trigger:value
14     ]/order_line[
15       @wrap_item:value==@rule/delivery_id:value
16     ]/order_line_id
17       -->> { !item_id[@trigger:value] },
18     item_id -->+ deliver_items,
19     deliver_items -->% deliver_items
20   }
21 }
22 }
23

```

As for the previous examples, a `delivery` data element is created by assigning a value, in this case a valid customer

id, to an input data element `create_delivery` (line 2–4). The delivery has data elements, just as the other classes, updated with a fresh id (line 8) and a copy of the customer id (line 9) provided when it was created. Each `delivery` data element moreover contains a `deliver_items` input data element (line 10), which is initially both pending and excluded, the exclusion being declared by the prefix `%`. Initially excluded data elements are another feature inherited from DCR Graphs. Excluded data elements may be included using an *include* reaction rule. In the analogy with web forms, excluded data elements correspond to hidden fields, which may be made dynamically visible.

The Path expression in line 12–16 introduces the use of filters constraining which data elements will trigger the rule when they are updated by a computation event. The first filter, `@customer_id:value==@trigger:value` ensures that only orders of the customer with id matching the id of the delivery are considered. The second filter, `@wrap_item:value==@rule/delivery_id:value` ensures that only items in the order line wrapped for that delivery are considered. (Recall that a delivery id was added when an item was wrapped). The path expression thereby make sure that only order lines with the designated delivery id and for the customer assigned to the delivery triggers the creation of an `item_id` data element. Once the computation event for the created pending `item_id` computation data element is executed, the rule `item_id ->+ deliver_items` in line 20 causes the `deliver_items` data element to be dynamically included. Finally, once the computation event for the `deliver_items` element happens, the `deliver_items` data element is again excluded by the `deliver_items ->% deliver_items` rule. These rules together with the constraint of the `wrap_item` input events that items can only wrapped for deliveries that have not been delivered implement the constraints (8) and (9) of Figure 2, i.e. that a delivery can only happen after an item has been added, and that an item can only be delivered once.

This example demonstrates that behavioural constraints of business processes as well as their underlying data model can be declared harmoniously specified together in RESEDA. In the next sections we provide the formal syntax and semantics.

III. SYNTAX

In this section we introduce the formal syntax of the RESEDA calculus, leaving out types to simplify the exposition. RESEDA *processes* are the syntactic entities generated by the syntax in Figure 3. (We will relate the formal syntax to the concrete ASCII-syntax used in the previous section shortly.)

A RESEDA process P consists of a vector \vec{D} of (structured) *data terms*, followed by a vector \vec{R} of *relation terms*. Both data and relation terms are built over an *expression language*. As already described above, in the ASCII syntax, we separate data terms and rules by a symbol `;`, and nested subprocesses are enclosed in braces.

$P ::= \vec{D}; \vec{R}$	processes
$D ::= n[E]:M\{P\} \mid n[?]:M\{P\}$	structured data
$M ::= (h, i, r, v)$	markings
$R ::= \phi \rightarrow \bullet \phi \mid \phi \rightarrow \diamond \phi \mid \phi \bullet \rightarrow \phi$ $\mid \phi \rightarrow + \phi \mid \phi \rightarrow \% \phi \mid \phi \rightarrow \rightarrow P$	relations
$\phi ::= / \psi \mid \psi$	path expressions
$\psi ::= \alpha \mid \psi / \alpha \mid \psi[E]$	
$\alpha ::= . \mid .. \mid n \mid *$	
$E ::= c \mid f(E_1, \dots, E_n) \mid \phi \mid \text{trigger}$ $\mid \text{this} \mid \text{rule} \mid E:\text{attr} \mid \perp \mid \mathbf{1}$	data expressions
$\text{attr} ::= \text{value} \mid \text{executed}$ $\mid \text{included} \mid \text{pending}$	

Fig. 3. Syntax of RESEDA.

A. Expression language

Expressions in RESEDA are built from the mutually recursive classes of *data expressions* E and *path expressions* ϕ .

Data expressions E are built over an unspecified set of constants c , assumed to include the positive integers; function symbols f ; the undefined value \perp and the unit value $\mathbf{1}$, and three constructs with particular meaning in RESEDA. Data expressions are meant to be evaluated always in the context of a given subterm of a process p . Given such a context, the remaining constructs are:

- 1) ϕ , a path expression, identifying a set of subterms of P ;
- 2) `this`, the subterm identified by the current context; and
- 3) `rule`, if used in a rule, the subterm identified by the context in which the rule is defined; and
- 4) `trigger`, the subterm identifying the data element that caused a spawning effect; and
- 5) $E:\text{attr}$, the (set of) attribute values at the current context.

Path expressions ϕ navigate between subprocesses in a manner that should be superficially familiar from, e.g., the UNIX filesystem or XPath expressions [7]: A path expression is a name (n), an instruction to move up (`.`), a wildcard (`*`) matching any child process, or the concatenation of such steps. Moreover, like in XPath, a path can be filtered by an expression along the way ($\psi[E]$).

In the concrete syntax, we distinguish path expressions from data expressions by prefixing the former with an `@`, and we allow abbreviations of the attribute selectors `value`, `executed`, etc. into simply `v`, `e`, etc.

Example 10. In Example 9, we see path expressions on lines 3, 9, 12-16 and 18. In line 3, for instance, the path expression `@/customers/customer/customer_id:value` starts with a `/` meaning that it selects data elements in the upper context (see Example 8). In line 12-16 we have an example of a filtered path expression, that only matches `order_line` data elements inside orders with a particular value for the `customer_id` and `wrap_item`. The path expression `customer_id:value` used

in the first filter (line 13) matches data elements in deeper context levels, in this case the nested elements of the order element, while the path expression `rule/delivery_id:value` used in the second filter (line 15) matches in the context where the rule is defined, i.e. inside the `delivery` data element.

B. Structured data

A data term d consists of either a *computation data element* $n[E]:M\{P\}$ or an *input data element* $n[?]:M\{P\}$. In either case, the data element is associated with a *marking* M and a subprocess P . Computation data elements further have an associated expression E , indicating what computation they carry out when executed (updated). Input data elements are syntactically distinguished from computation data elements by having a question mark ‘?’ instead of such an expression.

A marking M is a 4-tuple comprising:

- 1) a boolean h , indicating whether the data element has been executed at least once before,
- 2) a boolean i indicating whether the data element is currently included,
- 3) a boolean r indicating whether the data element is currently pending (i.e. required to be updated again if included) and finally
- 4) a value v indicating the current value of the data element.

A special undefined value is used for a data element that has yet to compute/receive a value for the first time and has no initial value.

In the ASCII variant of this syntax used in previous sections and in the online tool, we admit a number of shorthands:

- 1) Markings are by default (false, true, true, \perp), that is, unless otherwise specified, data elements that have not been executed, are included, not pending, and have the “undefined” value.
- 2) Data elements are marked “executed”, “excluded” respectively “pending” in the ASCII syntax by prefixing them with `:`, `%`, respectively `!`.
- 3) Data elements with no square brackets are considered computation data elements computing the unit value, `1`.

Example 11. In Example 9, we see on lines 8–10 pending data elements, and in line 10 also an excluded data element.

C. Relations

A relation definition R consists of one of six relations:

- 1) the condition $\rightarrow\bullet$ (ASCII `->*`), indicating that data elements on the right cannot execute unless each data on the left is either marked not included or executed;
- 2) the milestone $\rightarrow\circ$ (ASCII `-<>`), indicating that elements on the right cannot execute unless each element on the left is marked not included or not pending;
- 3) the response $\bullet\rightarrow$ (ASCII `*->`), indicating that whenever some data element on the left executes, all data elements on the right become marked pending;
- 4) the dynamic inclusion $\rightarrow+$ (ASCII `->+`), indicating that whenever some element on the left executes, all elements on the right become marked included;

- 5) the dynamic exclusion $\rightarrow\%$ (ASCII `->%`), indicating that whenever some element on the right executes, all elements on the right become excluded; and finally,
- 6) the spawn $\rightarrow\rightarrow$ (ASCII `->>`), indicating that whenever some data element on the left execute, new sub processes (data elements and rules) as indicated on the left are merged into the current process.

For the first five relations, both the domain and range of the relation is a path expression indicating which data elements (left-hand side) are being related to which other data elements (right-hand side). For the latter relation, the right-hand side indicates the process to be spawned.

Example 12. We see an example of the spawn relation in Example 9 starting on lines 6 and 12 respectively, where a process is spawned for each execution of respectively the `create_delivery` data element inside the `delivery` data element and the `order_line_id` data element.

IV. SEMANTICS

The semantics of RESEDA is given as a transition system, where states are the data elements combined with information about which elements are pending for update and currently included. The transitions are events corresponding to executions (i.e. value updates) of either input data elements or computation data elements, specified as a path to the corresponding data element, and, if the event is an input event, a value to input. To define this transition system, we must define two functions:

- 1) One that determines, for given a RESEDA term, which data elements are currently executable in that term; and
- 2) One that determines, for given a RESEDA term and a data element of that term, what is the next state (term) after executing this data element.

Because data elements in a process are not necessarily identified uniquely by name, we identify a data element by a path from the root to it. Such a path can be uniquely represented as a sequence of integers, each integer representing the left-to-right index of the subterm to descend into. In the sequel, we let ρ range over such sequences of numbers, calling such a ρ a *location*. A path expression evaluates to the set of locations it identifies; for a path expression ϕ evaluated from the location ρ in a process P , we write this set $\llbracket\phi\rrbracket_{\rho}^P$. Similarly, the value of a data expression E we write $\langle E \rangle_{\rho}^P$. It is straightforward to give a mutually recursive definition of these two evaluation functions; refer to [16].

A. Enabledness

The notion of *enabled* data element ρ is defined in Figure 6. For a data element ρ to be enabled, its parents must also be, viz. the recursive definition of the enabled function. Moreover, the data element ρ that we are considering for enabledness must satisfy that (enabled’):

- 1) every data element ρ' that is a condition for ρ must be either excluded or previously executed,

- 2) every data element ρ' that is a milestone for ρ must be either excluded or not pending, and finally,
- 3) the data element ρ must itself be included.

The $\text{enabled}'$ function makes this check. Note how in Figure 6 it is always invoked by enabled at the root P : Both for the data element ρ we are checking for enabledness and *all of its parents*, we must traverse the entire process, looking for relations that might prevent it from being enabled.

B. Transitions

Finally, we define the transition system of a process P . Recall that we intend states to be processes P and transitions to be either locations ρ (identifying data elements), or pairs of locations ρ and (input) values v .

We begin by defining the set of *effects* executing a data element ρ on a process P' in the context of a global process P ; we write this set $\text{effects}_P(P', \rho)$, defined by cases in Figure 4. The first case described defines that the effects are computed from the relations present in the term and inductively computed in the structure of the data element. Next, we treat each base case separately, by registering which data element may cause what particular effect (pend, incl, excl), and to what other data elements it is applied to.

Finally, we define the labelled transition system by the transition rules in Figure 5, making a distinction between computation events (rule COMPUTE) and input events (rule UPDATE). The former transitions are labelled by a data element location ρ , whereas the latter is labelled $\rho(v)$ by both a data element location ρ and an input value v . In both cases the selected data element ρ must be enabled, and the effects of the data element's execution δ are applied to a process P' where the target data element is already (re)computed/updated.

Computing a transition involves in either case three steps:

- 1) In either case, the data element in question must be *enabled* as defined in Figure 6 for the transition to apply.
- 2) Then, for an input data element ρ , the value of ρ is updated with the value v input ($\text{compute}_\rho^P(P, \rho)$). For a computation data element, the computation is executed ($\text{update}_\rho^P(P, \rho, v)$).
- 3) Finally, the effects of executing ρ are computed and applied ($\text{effects}_{P'}(P', \varepsilon)(\rho)$).

See [16] for full definitions of compute, update, and effects.

This concludes our introduction of RESEDA semantics. We proceed in the following Sections to study static properties of the language: its expressiveness and its relation to the notion of liveness.

V. EXPRESSIVE POWER

We report in this section that Reseda processes embed DCR* processes [5]¹.

Theorem 1. *Let G be a DCR* process. Then there exists a Reseda process $\llbracket G \rrbracket$ s.t. the languages of G and $\llbracket G \rrbracket$ are identical. Moreover, every data expression in $\llbracket G \rrbracket$ is trivial.*

¹Restating the syntax and semantics of DCR* is out of scope for the present paper; we refer instead the reader to [5].

We can use this encoding to study the expressive power of RESEDA. In [5] it is shown how to encode counter machines in DCR* and from that proving that the problem determining whether a given sequence of events is accepted by a DCR* process is undecidable.

Corollary 1. *The problem of determining whether the language of a RESEDA process P includes a string s of events is undecidable. This problem is undecidable already in the fragment of RESEDA where every data expression is trivial.*

VI. LIVENESS

We aim for the central property that a pending data-element can always be either executed or excluded in the future. When this is the case we say that the program is *live*.

Definition 1. Let P be a process. We say that P is *live* iff whenever Q is reachable from P and an event ρ is included and pending in Q , then there exists an R reachable from Q s.t. either ρ was executed on the way from Q to R , or ρ is not included in R .

We will not be able to achieve this property for arbitrary RESEDA processes. In particular, we rule out RESEDA processes with subprocesses whose bodies spawn copies of themselves. We formalise processes not including this behaviour as *bounded* processes.

Definition 2 (Runs). Let P be a process. The set $R(P)$ of *runs* of P is the set

$$R(P) = \{\rho_1, \dots, \rho_n \mid P = P_1 \xrightarrow{\rho_1} P_2 \xrightarrow{\rho_2} \dots \xrightarrow{\rho_n} P_n\}.$$

The *non-repeating runs* $R^!(P)$ of P is the subset of $R(P)$ defined by $\rho_1, \dots, \rho_n \in R^!(P)$ iff the ρ_i are pairwise distinct, i.e., for $1 \leq i < j \leq n$ we have $\rho_i \neq \rho_j$.

Definition 3 (Bounded process). Let P be a process. We say that P is *bounded* iff for every Q reachable from P , $R^!(Q)$ is finite.

Definition 4. Let P be a process, and suppose P contains as subterm a relation $\phi_b \rightarrow_{x \in E} Q$, and in turn that Q contains as subterm an event a . Let b be any name mentioned in ϕ . In this case we say there is a *potential spawn dependency* (PSD) from b to a in P .

The following Lemma is then immediate.

Lemma 1. *Let P be a process and suppose the graph of PSDs in P is acyclic. Then P is bounded.*

Now that we have identified the class of bounded processes, we sketch how to ensure liveness; for full proofs and definitions, refer to full version of this paper [16]. The basic idea follows that of [2] in constructing a static approximation of the dependencies between events that may be incurred by conditions and milestones, referred to as the *static inhibitor graph*, and then require the static inhibitor graph to be acyclic. To execute a given event, we may simply execute “from the bottom” of that acyclic graph, eventually reaching the desired event at the top.

Definition 5 (Static inhibitors). Let P be a process, and let ρ be an event of P . The *static inhibitors of ρ in P* , written $SI_P(\rho)$ is the set of events ρ' of P s.t. there is in P a condition or milestone $\rho' \rightarrow \bullet \rho$ or $\rho' \rightarrow \diamond \rho$. The *static inhibitor graph of ρ in P* , written $SI_P^*(\rho)$ is the graph arising as the closure under $SI_P(-)$ of ρ , further closed under the rule that if $\rho' \bullet \rightarrow \rho''$ or $\rho' \rightarrow + \rho''$ in P with $\rho', \rho'' \in SI_P^*(\rho)$ then there is an edge ρ', ρ'' in $SI_P^*(\rho)$.

While the core of the idea of [2] remains in the present development, we also need to account for events that are generated dynamically as a result of spawn relations. We do so by requiring that the static inhibitor graph can at any time be embedded in a fixed, “global” graph, referred to as the *dependency control*.

Definition 6 (Dependency control). Let P be a process, and let ρ be an event of P . A *dependency control* (G_ρ, ι_ρ^-) for ρ is a finite non-empty acyclic graph G_ρ and a family of graph homomorphisms $\iota_\rho^Q : SI_Q^*(\rho) \rightarrow G$, for Q reachable from P , s.t. for any Q reachable from P and Q' with $Q \xrightarrow{\rho''} Q'$:

- 1) if $\iota_\rho^Q(x)$ defined, then $\iota_\rho^Q(x) = \iota_\rho^{Q'}(x)$, and
- 2) for $\rho' \in SI_{Q'}^*(\rho) \setminus SI_Q^*(\rho)$ we have $\iota_\rho^{Q'}(\rho') > \iota_\rho^Q(\rho'')$.

Given a dependency control, we can find appropriate first actions to execute. Notice that, exactly because events can be dynamically generated, the planned actions to execute *must* be derived from the global dependency control, and not from the particular static inhibitors graph: It may be that a dependency has yet to materialise; in that case, it will not be present in the static inhibitors graph, but it *will* be represented in the dependency control.

Theorem 2. *Let P be a bounded process, let Q be a process reachable from P , and let ρ be an event of Q . If for all $\rho' \in \text{events}(P)$ the dependency graphs $G_{\rho'}$ of P is acyclic, then so is the dependency graph for ρ in Q .*

Corollary 2. *Let P be a bounded process, and let P' be the unfolding of P . If for every node ρ of P' the dependency graph G_ρ of P' is acyclic, then P is live.*

We conclude by noting that the prerequisites of Corollary 2 are efficiently computable; the following proposition forms the basis of the implementation at <http://dcr.itu.dk/reseda>.

Proposition 1. *Let P be a process. We can compute in polynomial time (a) whether P is bounded, (b) the unfolding P' of P , (c) the dependency graphs for the events of P and (d) whether these dependency graphs are acyclic.*

VII. DISCUSSION & RELATED WORK

RESEDA relates to so-called data-centric approaches, including *case handling* [14], [23], *business artefacts* [3], [11] and other artefact- and object-centric notations [9], [21], [22]. All these approaches give more equal emphasis on the modelling of data and behaviour than process-centric approaches such as BPMN [13]. The case handling paradigm is similar to RESEDA by being data-driven, based on forms, and distinguishing

mandatory and optional data fields. However, as pointed out in [21], [22], the artefact-centric process notations tend not to provide a good overview of the overall data and behavioural model, since the behavioural modelling is distributed on artefacts. The authors of [21] propose instead the Object-Centric Behavioural Constraints (OCBC) approach to provide combined presentations of the data-model and behavioural constraints. The OCBC notation also include relations between the two models (not shown in Figure 1 and Figure 2). In this example, the cross model relations indicate a one-to-one relationship between activities *create order* and *pick item* (marked by (1) and (2) respectively) and the elements of the classes Order and Order Line, indicating that these activities represent the activities for creating elements of these classes. This is reflected in the RESEDA process (cf. Example 8) by the spawn rules triggered by the `create_order` and `pick_item` input data elements.

None of the data-centric approaches nor the OCBC approach provide a formalised representation of data and behaviour unifying design and run-time as it is the case for RESEDA. The business artefacts approach models life cycles of data objects as finite state machines and not initially supported by an operational semantics for execution [9], [20]. This was provided with the GSM approach. The operational semantics for GSM is however considerably more complex than the one defined for RESEDA in the present paper and it is non-trivial to relate GSM to the declarative constraints present in DECLARE and DCR on which RESEDA and OCBC is based, e.g., see [8] for a comparison between GSM and DCR.

XML-based approaches following [1] build interesting dynamic and nested databases by extending a local schema with external sources. Although structurally similar, our approach provides a fully expressive language to build and change the structure of nested data objects. Finally, RESEDA relates to the work on *reactors* [18], which combines relational databases with behaviour specified declaratively as actors, but has a focus on predictable performance and relational databases.

VIII. CONCLUSIONS

We introduced a novel executable, declarative data-centric process language RESEDA unifying data and behavioural modelling as reactive semi-structured data, inspired in equal parts by declarative business process modelling notations and reactive spreadsheet programming. We demonstrated the use of RESEDA on a process given in [21], [22]. We provided a formal semantics for RESEDA defining an execution semantics as rewrite rules of processes and outlined the proof of a sufficient, polynomially decidable criteria for liveness of bounded processes. Finally we noted that Turing-completeness of RESEDA can be derived from the fact that RESEDA embeds the DCR* language. We leave for future work to investigate generalising the results of safe refinement in [6] to RESEDA. The RESEDA language is supported by a prototype implementation available both in source form and on-line at <http://dcr.itu.dk/reseda>. All major examples of the present paper are included in the tool.

$$\begin{aligned}
\text{effects}_P(\vec{D}; \vec{R}, \rho) &\triangleq \bigcup_{R \in \vec{R}} \text{effects}'_P(R, \rho) \cup \bigcup_{1 \leq i < n} \text{effects}_P(D_i, \rho_i) \\
&\quad \text{where } P_i = n_i[.]:M_i\{P_i\} \\
\text{effects}'_P(\phi \bullet \rightarrow \phi', \rho) &\triangleq \{(\rho', (\text{pend}, \rho'')) \mid (\rho', \rho'') \in \llbracket \phi \rrbracket_\rho^P \times \llbracket \phi' \rrbracket_{\rho'}^P\} \\
\text{effects}'_P(\phi \rightarrow + \phi', \rho) &\triangleq \{(\rho', (\text{incl}, \rho'')) \mid (\rho', \rho'') \in \llbracket \phi \rrbracket_\rho^P \times \llbracket \phi' \rrbracket_{\rho'}^P\} \\
\text{effects}'_P(\phi \rightarrow \% \phi', \rho) &\triangleq \{(\rho', (\text{excl}, \rho'')) \mid (\rho', \rho'') \in \llbracket \phi \rrbracket_\rho^P \times \llbracket \phi' \rrbracket_{\rho'}^P\} \\
\text{effects}'_P(\phi \rightarrow \rightarrow P', \rho) &\triangleq \{(\rho', (\text{spawn}(P' \{ \rho':v \} / \text{trigger}:v), \rho')) \mid \rho' \in \llbracket \phi \rrbracket_\rho^P\}
\end{aligned}$$

Fig. 4. Effects of Terms

$$\begin{aligned}
&\frac{\text{enabled}_P(\rho) \quad P' = \text{compute}_\rho^P(P, \rho) \quad \delta = \text{effects}_{P'}(P', \varepsilon)(\rho)}{P \xrightarrow{\rho} P' \triangleleft \delta} \quad (\text{COMPUTE}) \\
&\frac{\text{enabled}_P(\rho) \quad P' = \text{update}_\rho^P(P, \rho, v) \quad \delta = \text{effects}_{P'}(P', \varepsilon)(\rho)}{P \xrightarrow{\rho(v)} P' \triangleleft \delta} \quad (\text{UPDATE})
\end{aligned}$$

Fig. 5. Labelled Transition System

$$\begin{aligned}
\text{enabled}'(\rho, P, \vec{D}; \vec{R}, \rho') &= \forall \phi \bullet \rightarrow \phi' \in R. \rho \in \llbracket \phi \rrbracket_{\rho'}^P \implies \llbracket \phi[\text{included}] \rrbracket_{\rho'}^P = \llbracket \phi[\text{included} \wedge \text{executed}] \rrbracket_{\rho'}^P \\
&\quad \wedge \forall \phi \rightarrow \phi' \in R. \rho \in \llbracket \phi \rrbracket_{\rho'}^P \implies \llbracket \phi[\text{included}] \rrbracket_{\rho'}^P = \llbracket \phi[\text{included} \wedge \neg \text{pending}] \rrbracket_{\rho'}^P \\
&\quad \wedge \forall i \in \text{size}(\vec{D}). \text{enabled}'(\rho, P, P_i, \rho'_i) \quad (\text{where } \vec{D} = P_1, \dots, P_n \text{ and } P_i = n_i[.]:M_i\{P'_i\}) \\
\text{enabled}_P(\rho_i) &= \text{enabled}_P(\rho) \wedge \text{enabled}'(\rho_i, P, P, \varepsilon) \quad \text{and} \quad \text{enabled}_P(\varepsilon) = \text{true}
\end{aligned}$$

Fig. 6. Enabledness.

REFERENCES

- [1] Serge Abiteboul, Omar Benjelloun, and Tova Milo. The active xml project: an overview. *The VLDB Journal*, 17(5):1019–1040, Aug 2008.
- [2] David A. Basin, Søren Debois, and Thomas T. Hildebrandt. In the nick of time: Proactive prevention of obligation violations. In *IEEE 29th Computer Security Foundations Symposium, CSF 2016, Lisbon, Portugal, June 27 - July 1, 2016*, pages 120–134, 2016.
- [3] K. Bhattacharya, R. Hull, and J. Su. *A Data-Centric Design Methodology for Business Processes*, pages 503–531. IGI Global, 2009.
- [4] Søren Debois, Thomas Hildebrandt, and Tijs Slaats. Hierarchical declarative modelling with refinement and sub-processes. In *International Conference on Business Process Management*, pages 18–33. Springer, Cham, 2014.
- [5] Søren Debois, Thomas T. Hildebrandt, and Tijs Slaats. Safety, liveness and run-time refinement for modular process-aware information systems with dynamic sub processes. In *FM 2015*, pages 143–160, 2015.
- [6] Søren Debois, Thomas Hildebrandt, and Tijs Slaats. Safety, liveness and run-time refinement for modular process-aware information systems with dynamic sub processes. In Nikolaj Bjørner and Frank de Boer, editors, *FM 2015: Formal Methods*, pages 143–160, Germany, 2015. Springer.
- [7] Denise Draper, Peter Fankhauser, Mary Fernandez, Ashok Malhotra, Kristoffer Rose, Michael Rys, Jarome Simon, and Philip Wadler. XQuery 1.0 and XPath 2.0 formal semantics. *W3C recommendation*, 23, 2007.
- [8] R. Eshuis, S. Debois, T. Slaats, and T. Hildebrandt. Deriving consistent gsm schemas from dcr graphs. In *Service-Oriented Computing. ICSOC 2016.*, volume 9936 of *Lecture Notes in Computer Science*. Springer, Cham, 2016.
- [9] Hull R. et al. Introducing the guard-stage-milestone approach for specifying business entity lifecycles. In *Web Services and Formal Methods. WS-FM 2010*, volume 6551 of *Lecture Notes in Computer Science*, Berlin, Heidelberg, 2011. Springer.
- [10] Thomas T Hildebrandt and Raghava Rao Mukkamala. Declarative event-based workflow as distributed dynamic condition response graphs. *arXiv:1110.4161*, 2011.
- [11] Richard Hull. Artifact-centric business process models: Brief survey of research results and challenges. In *OTM '08. Part II on On the Move to Meaningful Internet Systems, OTM '08*, pages 1152–1163. Springer, 2008.
- [12] James McKinna. Why dependent types matter. In *Conference Record of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '06, pages 1–1, New York, NY, USA, 2006. ACM.
- [13] Object Management Group BPMN Technical Committee. *Business Process Model and Notation, Version 2.0*, 2013.
- [14] Object Management Group CMMN Technical Committee. *Case Management Model and Notation, Version 1.1*, 2016.
- [15] M. Pesic, H. Schonenberg, and W.M.P. van der Aalst. DECLARE: Full Support for Loosely-Structured Processes. In *EDOC '07*, pages 287–. IEEE, 2007.
- [16] João Costa Seco, Søren Debois, Thomas Hildebrandt, and Tijs Slaats. RESEDA: Declaring live concurrent event-driven computations as RE-active SEmi-structured DATA. Technical Report TR-2018-202, IT University of Copenhagen, 2018. Available on-line at <http://www.itu.dk/~debois/reseda-tr.pdf>.
- [17] Peter Sestoft. *Spreadsheet implementation technology. Basics and extensions*. MIT Press, United States, 2014.
- [18] Vivek Shah and Marcos Antonio Vaz Salles. Reactors: A case for predictable, virtualized actor database systems. In *Proceedings of the 2018 International Conference on Management of Data, SIGMOD '18*, pages 259–274, New York, NY, USA, 2018. ACM.
- [19] Jianwen Su, Lijie Wen, and Jian Yang. From data-centric business processes to enterprise process frameworks. In *21st IEEE International Enterprise Distributed Object Computing Conference, EDOC 2017, Quebec City, QC, Canada, October 10-13, 2017*, pages 1–9, 2017.
- [20] Yutian Sun, Wei Xu, and Jianwen Su. Declarative choreographies for artifacts. In *Service-Oriented Computing - 10th International Conference, ICSOC 2012, Shanghai, China, November 12-15, 2012. Proceedings*, pages 420–434, 2012.
- [21] W. van der Aalst, A. Artale, Montali, M., and S. Tritini. Object-centric behavioral constraints: Integrating data and declarative process modelling. In *DL 2017, International Workshop on Description Logics*, volume 1879 of *CEUR Workshop Proceedings*, 2017.
- [22] Wil M. P. van der Aalst, Guangming Li, and Marco Montali. Object-centric behavioral constraints. *CoRR*, abs/1703.05740, 2017.
- [23] Wil M. P. van der Aalst and Mathias Weske. Case handling: A new paradigm for business process support. *Data Knowl. Eng.*, 53(2):129–162, May 2005.