

Designing a Cross-organizational Case Management System using Dynamic Condition Response Graphs

Thomas Hildebrandt
IT University of Copenhagen
Rued Langgaardsvej 7
2300 Copenhagen, Denmark
hilde@itu.dk

Raghava Rao Mukkamala
IT University of Copenhagen
Rued Langgaardsvej 7
2300 Copenhagen, Denmark
rao@itu.dk

Tijs Slaats
IT University of Copenhagen and
Exformatics A/S
2100 Copenhagen, Denmark
tslaats@itu.dk

Abstract—We present a case study of the use of Dynamic Condition Response (DCR) Graphs, a recently introduced declarative business process model, in the design of a cross-organizational case management system being developed by Exformatics A/S, a Danish provider of knowledge and workflow management systems. We show how DCR Graphs allow to capture directly both the behavioral constraints identified during meetings with the customer and the operational execution as markings of the graph. In comparison, imperative models such as BPMN, Petri Net, UML Sequence or Activity diagrams are only good at describing the operational way to fulfill the constraints, leaving the constraints implicit. In particular, we point out that the BPMN ad-hoc sub process activity, intended to support more loosely structured goal driven ad-hoc processes, is inconsistently described in the final version of the BPMN 2.0 standard. The case study motivated an extension of the DCR Graphs model to nested graphs and the development of graphical design and simulation tools to increase the understanding of the models. The study also revealed a number of challenges for future research in techniques for model-driven design of cross-organizational process-aware information systems combining declarative and imperative models.

Index Terms—Case Study, Declarative Workflow, Model-driven Design

I. INTRODUCTION

The purpose of a Case Management System as used in for instance Human Resource (HR) departments, hospitals, financial, and governmental institutions, is to guide case workers to perform the right tasks and to record the history of the case.

Since the initial work on office automation and workflow systems [9], [10], [32] it has been advocated to base the implementation of such systems, subsequently referred to as process-aware information systems [8], on explicit process descriptions described in some high-level process notation such as Petri Net or UML activity diagrams. The key motivations for using explicit process models are to allow the system to be more easily adapted to different work processes and to make the rules governing the system more visible to the users.

Authors listed alphabetically. This research is supported by the Danish Research Agency through a Knowledge Voucher granted to Exformatics (grant #10-087067, www.exformatics.com), the Trustworthy Pervasive Healthcare Services project (grant #2106-07-0019, www.trustcare.eu) and the Computer Supported Mobile Adaptive Business Processes project (grant #274-06-0415, www.cosmobiz.dk).

The rise of web service standards such as SOAP, WSDL and WS-BPEL has given new momentum to process-aware information systems. SOAP and WSDL standardize how to access external IT systems as web services in a service oriented architecture and WS-BPEL provides a standard high-level programming language for combining individual service calls into process flows, also referred to as a process orchestration. Following WS-BPEL, the BPEL4People [1] and WS-HumanTask [20] specifications were the first attempt to standardize the inclusion of human tasks into BPEL to encompass workflows. Moreover, W3C started in 2004 developing the Web Services Choreography Description Language (WS-CDL) [29] which can be used to provide a global view of the intended interactions between different actors of a system, similar to the view of interactions provided by UML sequence diagrams. Within the last 5 years focus has moved from WS-BPEL and BPEL4People to the development of Business Process Model and Notation (BPMN) [19] which standardizes the graphical notation used for business processes, encompassing both human and automated tasks, and including both notations for orchestrations and choreographies.

However, as pointed out in e.g. [10], [27], the imperative process notations with explicit control and message flows underlying all of the above models describe the operationalization of business process goals and constraints, and not the goals and constraints themselves. Consequently, the notations are best suited for well-defined, rigid and repeatable workflows following a predefined sequence of service invocations and human tasks and one need to use ad hoc annotations to record the constraints and goals of the process. Moreover, it has proven to be non-trivial to support changes of the processes on-the-fly [26]. This does not match well the typical more ad-hoc nature of case work where it is often needed to redo and skip tasks and possibly adapt the set of tasks and their mutual constraints dynamically [25].

An alternative approach studied by several research groups is the use of declarative process models [3], [6], [22], [23], [27], [28], which describes the temporal constraints on process flows, not how to fulfill them. As part of the PhD project of the second author within the Trustworthy Pervasive Healthcare Services (TrustCare) research project [11] we have developed a declarative process model called Dynamic Condition Re-

sponse Graphs (DCR Graphs) [12]–[14], [17]. The model is both a generalization of the Process Matrix model [16], [18] developed by Resultmaker, a danish provider of workflow and case-management systems and the classical event structure model for concurrency [30], [31]. DCR Graphs relate to DECLARE [28], which is a graphical notation that allows any temporal constraint pattern expressible as Linear-time Temporal Logic (LTL) formulas. However, instead of allowing the generality of expressing any constraint expressible in LTL, DCR Graphs only has a fixed handful of constraints which can be understood without reference to LTL. Still it maintains the full expressive power of LTL (described in a follow up paper). Moreover, it is possible to give an operational semantics expressed directly as transitions between a novel type of markings of the tasks. In this way DCR Graphs combine the declarative view (constraints between tasks) with the imperative view (markings of tasks) allowing to trace the constraints even at run-time.

In the present paper we first briefly review the definition of DCR Graphs in Sec. II and then in Sec. III describe a case study of applying DCR graphs in the design phase of the development of a cross-organizational case management system. In Sec. IV we briefly describe the current status of our development of tools for supporting design, simulation and verification of DCR Graphs. Finally, in Sec. VI we outline challenges identified in the case study and the proposal for the continued development of the DCR Graphs model, technologies and tools to make them applicable to component and model based design of distributed process-aware information systems.

II. DYNAMIC CONDITION RESPONSE GRAPHS

A Dynamic Condition Response Graph as introduced in [13] and extended in [14] consists of a set of events, a *marking* defining the execution state, and five binary relations between the events defining the conditions for the execution of events, the required responses and a novel notion of dynamic inclusion and exclusion of events. Hereto comes a set of actions, a labeling function assigning an action to each event, a set of roles, a set of principals and a relation assigning roles to actions and principals.

Notation: For a set A we write $\mathcal{P}(A)$ for the power set of A . For a binary relation $\rightarrow \subseteq A \times A$ and a subset $\xi \subseteq A$ of A we write $\rightarrow \xi$ and $\xi \rightarrow$ for the set $\{a \in A \mid (\exists a' \in \xi \mid a \rightarrow a')\}$ and the set $\{a \in A \mid (\exists a' \in \xi \mid a' \rightarrow a)\}$ respectively.

Formally we define a DCR Graph as follows.

Definition 1: A Dynamic Condition Response Graph is a tuple $(E, M, \rightarrow \bullet, \bullet \rightarrow, \rightarrow \diamond, \pm, \text{Act}, l, R, P, \text{as})$, where

- (i) E is the set of *events*
- (ii) $M \in \mathcal{M}(G)$ is the *marking*, and $\mathcal{M}(G) =_{def} \mathcal{P}(E) \times \mathcal{P}(E) \times \mathcal{P}(E)$
- (iii) $\rightarrow \bullet \subseteq E \times E$ is the *condition* relation
- (iv) $\bullet \rightarrow \subseteq E \times E$ is the *response* relation
- (v) $\rightarrow \diamond \subseteq E \times E$ is the *milestone* relation

- (vi) $\pm : E \times E \rightarrow \{+, \%\}$ is a partial function defining the *dynamic inclusion and exclusion* relations by $e \rightarrow + e'$ if $\pm(e, e') = +$ and $e \rightarrow \% e'$ if $\pm(e, e') = \%$
- (vii) Act is the set of *actions*
- (viii) $l : E \rightarrow \text{Act}$ is a *labeling* function mapping events to actions.
- (ix) R is a set of *roles*,
- (x) P is a set of *principals* (e.g. actors, persons, processors, services) and
- (xi) $\text{as} \subseteq (P \cup \text{Act}) \times R$ is the *role assignment* relation to principals and actions.

An event labelled with an action, e.g. **Create Case**, thus represents an execution of a (human or automated) task/activity/action **Create Case** in the workflow process. There may be several events with the same label, but in all our examples a label is assigned to a unique event, and thus we simply assume the set of events to be identical to the set of actions.

By default an event may be executed at any time and any number of times. However, the marking (defining the run-time state of the graph) and the five relations defined in (iii)-(vi) constrain the execution. The marking $M = (Ex, Re, In) \in \mathcal{M}(G)$ consists of three sets of events, capturing respectively which events have *previously been executed* (Ex), which events are *pending responses required to be executed* (Re), and finally which events are currently *included* (In). Only events $e \in In$, i.e. that are currently included, can be executed, and only if all currently included condition events e' , as specified by the condition relation $e' \rightarrow \bullet e$, have been executed and no currently included events e' which are milestones for e , as specified by the milestone relation $e' \rightarrow \diamond e$, are pending responses.

When an event e is executed, it is added to the set of executed events (Ex) of the marking and all response events e' , as specified by the response relation $e \bullet \rightarrow e'$, are added to the set of pending responses Re . Moreover, the set of included events is updated by adding (removing) all events e' included (excluded) by e as specified by the inclusion (exclusion) relation $e \rightarrow + e'$ ($e \rightarrow \% e'$).

The execution semantics of DCR Graphs is defined [12], [14] as a labelled transition system between markings as follows.

Definition 2: For a DCR Graph $G = (E, M, \rightarrow \bullet, \bullet \rightarrow, \rightarrow \diamond, \pm, l, \text{Act}, R, P, \text{as})$, we define the corresponding labelled transition systems $T(G)$ to be the tuple $(\mathcal{M}(G), M, \rightarrow \subseteq \mathcal{M}(G) \times \mathcal{L}(G) \times \mathcal{M}(G))$ where $\mathcal{L}(G) =_{def} E \times (P \times \text{Act} \times R)$ is the set of transition labels, $M = (Ex, Re, In) \in \mathcal{M}(G)$ is the initial marking, $\rightarrow \subseteq \mathcal{M}(G) \times \mathcal{L}(G) \times \mathcal{M}(G)$ is the transition relation given by $M' \xrightarrow{(e, (p, a, r))} M''$

where

- (i) $M' = (Ex', Re', In')$ is the marking before transition
- (ii) $M'' = (Ex' \cup \{e\}, Re'', In'')$ is the marking after transition
- (iii) $e \in In'$, $l(e) = a$, p as r , and a as r ,
- (iv) $(\rightarrow \bullet e \cap In') \subseteq Ex'$,
- (v) $(\rightarrow \diamond e \cap In') \cap Re' = \emptyset$,

- (vi) $In'' = (In' \cup e \rightarrow) \setminus e \rightarrow \%$,
- (vii) $Re'' = (Re' \setminus \{e\}) \cup e \bullet \rightarrow$,

We define a run $(e_0, (p_0, a_0, r_0)), (e_1, (p_1, a_1, r_1)), \dots$ of the transition system to be a sequence of labels of a sequence of transitions $M_i \xrightarrow{(e_i, (p_i, a_i, r_i))} M_{i+1}$ starting from the initial marking. We define a run to be accepting if $\forall i \geq 0, e \in Re_i, \exists j \geq i. (e = e_j \vee e \notin In_j)$. In words, a run is accepting if no response event is pending forever, i.e. it must either happen at some later state or become excluded.

Condition (iii) in the above definition expresses that, only events e that are currently included, can be executed, and to give the label (p, a, r) the label of the event must be a , p must be assigned to the role r , which must be assigned to a . Condition (iv) requires that all condition events to e which are currently included should have been executed previously. Condition (v) states that the currently included events which are milestones to event e must not be in the set of pending responses (Re'). Condition (vi) and (vii) are the updates to the sets of included events and pending responses respectively. Note that an event e' can not be both included and excluded by the same event e , but an event may trigger itself as a response.

In this paper we only consider finite runs. In this case, the acceptance condition degenerates to requiring that no pending response is included at the end of the run. This corresponds to defining all states where $Re \cap In = \emptyset$ to be accepting states and define the accepting runs to be those ending in an accepting state. Infinite runs are also of interest especially in the context of reactive systems and the LTL logic. The execution semantics and acceptance condition for infinite runs are captured by mapping to a Büchi-automaton with τ -event as formalized in [12], [17].

During the case study (Sec. III), we realized the need to extend our model with nested sub-graphs to allow for modeling of hierarchical sub structures. To address this need, so-called Nested DCR Graphs were introduced in [14]. It can be defined as an incremental extension to DCR Graph given in Def. 1 above as follows.

Definition 3: A Nested dynamic condition response graph is a tuple $(E, \triangleright, M, \rightarrow \bullet, \bullet \rightarrow, \rightarrow \diamond, \pm, Act, l, R, P, as)$, where $\triangleright : E \rightarrow E$ is a partial function mapping an event to its super-event (if defined) and $(E, M, \rightarrow \bullet, \bullet \rightarrow, \rightarrow \diamond, \pm, Act, l, R, P, as)$ is a DCR Graph, subject to the condition that the marking $M = (Ex, Re, In) \subseteq \text{atoms}(E) \times \text{atoms}(E) \times \text{atoms}(E)$ where $\text{atoms}(E) = \{e \mid \forall e' \in E. \triangleright(e') \neq e\}$ is the set of *atomic* events.

A nested DCR Graph can be mapped to a flat DCR Graph by extending all relations to the sub events and by preserving only the atomic events. This flattening of a nested DCR Graph into a DCR Graph is defined formally in [14]. In particular, the semantics of a Nested DCR Graph is given as the labelled transition semantics for its corresponding flattened DCR Graph.

III. CASE STUDY: A CROSS-ORGANIZATIONAL CASE MANAGEMENT SYSTEM

In this section we demonstrate how we have applied DCR Graphs in practice within a project that our industrial partner

Exformatics carried out for one of their customers. In the process, we acted as consultants, applying DCR Graphs in meetings with Exformatics and the customer to capture the requirements in a declarative way, accompanying the usual UML sequence diagrams and prototype mock-ups. Sequence diagrams typically only describe *examples* of runs, and even if they are extended with loops and conditional flows they do not capture the constraints explicitly.

The customer of the system is *Landsorganisationen i Danmark* (LO), which is the overarching organization for most of the trade unions in Denmark. Their counterpart is *Dansk Arbejdsgiverforening* (DA), which is an overarching organization for most of the Danish employers organizations.

At the top level, the workflow to be supported is that a case worker at the trade union must be able to create a case, e.g. triggered by a complaint by a member of the trade union against her employer. This must be followed up by a meeting arranged by LO and subsequently held between case workers at the trade union, LO and DA. After being created, the case can at any time be managed, e.g. adding or retrieving documents, by case workers at any of the organizations.

Fig. 1 shows the graphical representation of a simple DCR Graph capturing these top level requirements of our case study.




Figure 1. Top level requirements as a DCR Graph

Four top-level events were identified, shown as boxes in the graph labelled **Create case**, **Manage case**, **Arrange meeting** and **Hold meeting**. For the top-level events we identified the following requirements:

- 1) A case is created by a union case worker, and only once.
- 2) The case can be managed at the union, LO and DA after it has been created.
- 3) After a case is created, LO can and must arrange a meeting between the union case worker, the LO case worker and the DA case worker.
- 4) After a meeting is arranged it must be held (organized by LO).

The requirements translate to the following DCR Graph role assignments (shown as "ears" on the event boxes) and relations shown as different types of arrows between the events in Fig. 1:

- 1) Create case has assigned role U and excludes itself.
- 2) Create case is a condition for Manage case, which has assigned role U, LO and DA.
- 3) Create case has Arrange meeting as response, which has assigned role LO.
- 4) Arrange meeting has Create case as a condition and Hold meeting as response, which has assigned role LO.

For example, the U on Create case indicates that only a case worker at the trade union (U) can create a case, and the U, LO, DA on Manage case indicate that both the trade union, LO and DA can manage the case.

The arrow Create case \rightarrow Manage case denotes that Manage case has Create case as a (pre) condition. This simply means that Create case must have happened before Manage case can happen. Dually, Arrange meeting has Hold meeting as response, denoted by the arrow Arrange meeting $\bullet \rightarrow$ Hold meeting. This means that Hold meeting must eventually happen after Arrange meeting happens. Finally, the arrow Create case $\rightarrow \%$ Create case denotes that the event Create case excludes itself.

In the subsequent meetings, we came to the following additional requirements:

- 1) a) To create a case, the case worker should enter meta-data on the case, inform about when he/she is available for participating in a meeting and then submit the case.
- b) When a case is submitted it may get a local id at the union, but it should also subsequently be assigned a case id in LO.
- c) When a case is submitted, LO should eventually propose dates.
- 2) a) Only after LO has assigned its case id it is possible to manage the case and for LO to propose dates.
- b) Manage case consists of three possible activities (in any order): editing case meta data, upload documents and download documents. All activities can be performed by LO and DA. Upload and download documents can also be performed by the Union.
- 3) a) The meeting should be arranged in agreement between LO and DA: LO should always propose dates first - and then DA should accept, but can also propose new dates. If DA proposes new dates LO should accept, but can also again propose new dates. This could in principle go on forever.
- b) The union can always update information about when they are available and edit the metadata of the case.
- 4) a) No meeting can be held while LO and DA are negotiating on a meeting date. Once a date has been agreed upon a meeting should eventually be held.

These requirements led to the extension of the model allowing nested events as recalled in the previous section and given in full detail in [14].

The requirements could then be described by first adding the following additional events to the graph: A new super event Edit (E) which has the sub events: Metadata (E-M) and Dates available (E-D) and is itself a sub event to Create case (CC).

The Create case (CC) event has two sub events: Submit (SC) and Assign case Id (ACI). The Manage case (MC) event has two sub events: Edit metadata (EM) and Document (D), which in turn has two sub events: Upload (D-U) and Download (D-D). The Arrange meeting (AM) event has four sub events: Propose dates-LO (PLO), Propose dates-DA (PDA), Accept LO (ALO) and Accept DA (ADA). The Hold meeting (HM) event remains an atomic top-level event.

Subsequently, the relations was adapted to the following (Nested) DCR Graph relations, as shown in Fig 2:




Figure 2. Case Handling Process

- 1) Edit is a condition to Submit and is assigned role U.
- 2) Within the Create case superevent:
 - a) Submit is a condition to Assign case Id and also requires it as a response.
 - b) Assign case Id is a condition for Manage case (and therefore also all it's sub events).
 - c) Assign case Id is now the condition for Propose dates-LO and Submit requires it as a response.
- 3) Within the Arrange meeting superevent:
 - a) Arrange meeting still has Hold meeting as response, but is now also required as a milestone for Hold meeting
 - b) Propose dates-LO is a condition for Propose dates-DA
 - c) Propose dates-LO includes Accept DA and requires it as a response
 - d) Propose dates-DA includes Accept LO and requires it as a response
 - e) Accept LO excludes itself and Accept DA
 - f) Accept DA excludes itself and Accept LO
- 4) Within the Manage case superevent:

- a) Edit metadata has roles LO and DA assigned to it.
- b) Upload and Download have been grouped under a superevent Document with roles U, LO and DA assigned to it.
- c) Upload is a condition for Download.

The case handling process described above and shown in figure 2 can be represented formally as follows.

$G = (E, \triangleright, M, \rightarrow \bullet, \bullet \rightarrow, \rightarrow \diamond, \pm, \text{Act}, I, R, P, \text{as})$, where
 $\text{Act} = \text{atoms}(E) = \{E\text{-M}, E\text{-D}, \text{SC}, \text{ACI}, \text{EM}, \text{D-U}, \text{D-D}, \text{PLO}, \text{PDA}, \text{ALO}, \text{ADA}, \text{HM}\}$
 $E = \{\text{CC}, \text{AM}, \text{MC}, \text{E}, \text{D}\} \cup \text{atoms}(E)$
 $\triangleright = \{(E\text{-M}, E), (E\text{-D}, E), (E, \text{CC}), (\text{SC}, \text{CC}), (\text{ACI}, \text{CC}), (\text{PLO}, \text{AM}), (\text{PDA}, \text{AM}), (\text{ALO}, \text{AM}), (\text{ADA}, \text{AM}), (\text{ALO}, \text{D-D}), (\text{D-U}, \text{D}), (\text{D}, \text{MC}), (\text{EM}, \text{MC})\}$
 $M = (\emptyset, \emptyset, \text{atoms}(E) \setminus \{\text{ALO}, \text{ADA}\})$
 $\rightarrow \bullet = \{(E, \text{SC}), (\text{SC}, \text{ACI}), (\text{ACI}, \text{MC}), (\text{ACI}, \text{PLO}), (\text{D-U}, \text{D-D}), (\text{PLO}, \text{PDA})\}$
 $\bullet \rightarrow = \{(\text{SC}, \text{ACI}), (\text{SC}, \text{PLO}), (\text{PLO}, \text{ADA}), (\text{PDA}, \text{ALO}), (\text{AM}, \text{HM})\}$
 $\rightarrow \diamond = \{(\text{AM}, \text{HM})\}$
 $\rightarrow + = \{(\text{PLO}, \text{ADA}), (\text{PDA}, \text{ALO})\}$
 $\rightarrow \% = \{(\text{SC}, \text{SC}), (\text{ALO}, \text{ALO}), (\text{ALO}, \text{ADA}), (\text{ADA}, \text{ADA}), (\text{ADA}, \text{ALO})\}$
 $I = \{e \in \text{atoms}(E) \mid (e, e)\}$
 $R = \{U, LO, DA\}$ and $P = \{U, LO, DA\}$
 $\text{as} = \{(\text{SC}, U), (E, U), (D, U), (\text{ACI}, LO), (\text{EM}, LO), (D, LO), (\text{PLO}, LO), (\text{ALO}, LO), (\text{HM}, LO), (\text{EM}, DA), (D, DA), (\text{PDA}, DA), (\text{ADA}, DA), (U, U), (LO, LO), (DA, DA)\}$

During the case study it became clear that it would be useful to have design tools allowing to quickly create and simulate models. In the following section we describe the tools developed so far. In Sec. VI we describe the plans for future development of tools along with the challenges for extending the theory identified in the case study.

IV. PROTOTYPE TOOLS

To support designing with DCR Graphs, making the model available to a wider audience and allow interested parties to experiment with the notation, we are developing prototype implementations of various tools for DCR Graphs. Development up to this point includes:

- 1) A process repository; a service which can be used to store and retrieve DCR processes and process instances.
- 2) An execution host; a service which can be used to execute DCR process instances.
- 3) A windows-based graphical editor; which can be used to model DCR Graphs and run simple simulations on them.
- 4) A windows-based desktop client for executing process instances.
- 5) A platform independent web client; which can also be used to execute process instances. In the future we aim to support the creation of processes through this webinterface as well. (Fig. 3)

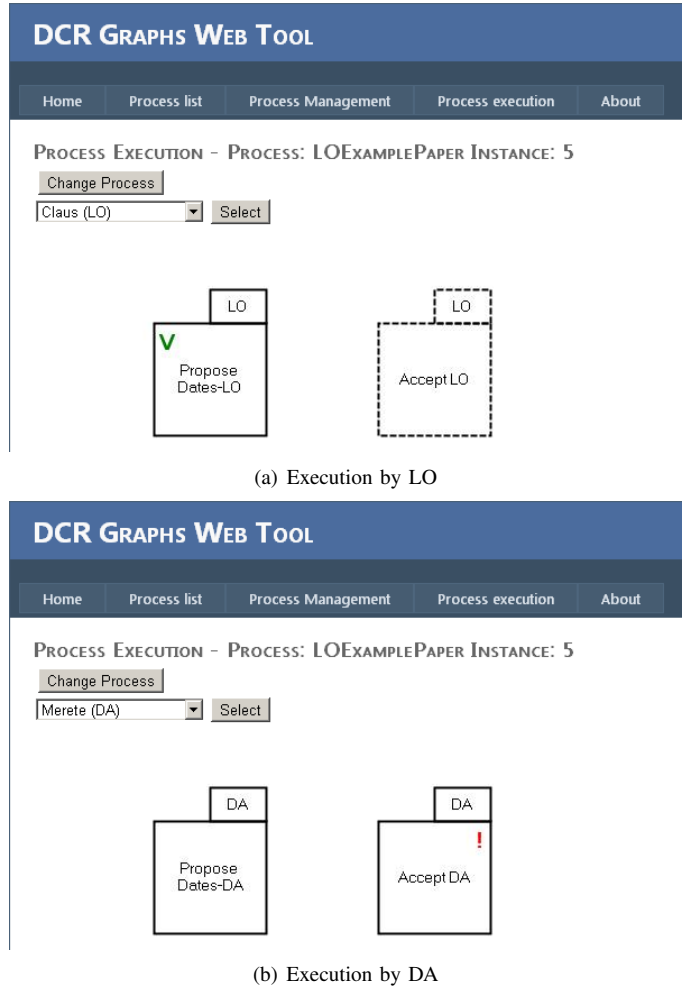


Figure 3. Execution in the Web Tool

- 6) A model checking and runtime verification tool; which interfaces to SPIN [24] and ZING [21] model checkers for model checking.
- 7) A runtime-monitor that can subscribe to the execution host and verify that the execution of processes adheres to given properties.

Fig. 4 shows how these tools interact: Usually, a process modeller will first create a process in the graphical editor (Fig. 5), which will be stored in the process repository. The process modeller can use the verification tool to check if his process adheres to the properties that he desires. Both safety and liveness properties on models can be verified with the help of SPIN [24] model checker, where as only safety properties on DCR Graphs can be verified using ZING [21] model checker, as ZING does not support liveness properties. A user can login to the web or desktop-client and select the process for execution. The client will request that the process repository start a new instance and the repository will provide the client with the description of the process and runtime information on the process instance. Execution requests are made to the execution server, which handles these requests atomically,

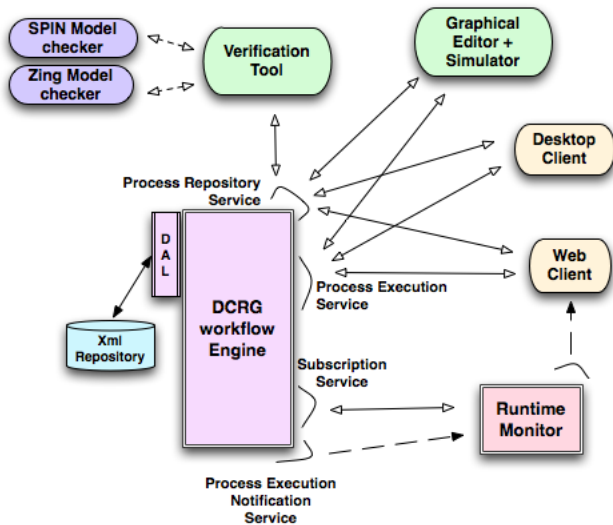


Figure 4. Prototype Architecture

making updates to the instance stored on the repository. If a request is invalid, the execution server will notify the user and leave the process instance in its original state. The runtime monitor can subscribe to the execution server and will get notified of every execution request. It will then check if the execution of the process follows the properties described for it.

Listing 1 shows a brief overview of the XML format of DCR Graphs that is being used in all prototype tools. A single XML format is used to contain information about both the specification and the runtime of a DCR Graph. The resources section of the specification contains information about roles, principals, events and actions, whereas the access controls section contains the mapping of principals and actions to roles. The last part of the specification contains the binary relations between the events. Note that the XML format supports nesting of events and the binary relations in between them and that flattening of nested events and their relations will be done at the beginning of executing a DCR Graph.

The second part of the XML format for a DCR Graph holds the runtime information, which primarily contains the execution trace and information about the current state. The execution trace records the actual sequence of events executed and the current state holds the information about the current marking which contains sets of included, executed and pending response events. In addition to the marking, the current state also holds additional information such as index of state copy, state accepted to support the acceptance condition for infinite computations that were characterized by mapping to Büchi-automata in [12], [17].

Listing 1. Overview of DCR Graph Xml

```
<?xml version="1.0" encoding="utf-8"?>
<dcr:process xmlns:dcr="http://itu.dk/trustcare/dcr/2011/">
  <dcr:specification processId="" modelName="">
```

```
<dcr:resources>
  <dcr:roles>....</dcr:roles>
  <dcr:principals>....</dcr:principals>
  <dcr:events>....</dcr:events>
  <dcr:actions>....</dcr:actions>
</dcr:resources>

<dcr:accessControls>
  <dcr:rolePrincipalAssignments>....</
  dcr:rolePrincipalAssignments>
  <dcr:actionRoleAssignments>....</
  dcr:actionRoleAssignments>
</dcr:accessControls>

<dcr:constraintSets>
  <dcr:constraintSet type="condition">...</
  dcr:constraintSet>
  <dcr:constraintSet type="response">...</
  dcr:constraintSet> ....
</dcr:constraintSets>
</dcr:specification>

<dcr:runtime processInstanceId="">
  <dcr:executionTrace> </dcr:executionTrace>
  <dcr:currentState stateId="">

    <dcr:eventsIncluded>....</dcr:eventsIncluded>
    <dcr:eventsExecuted>....</dcr:eventsExecuted>
    <dcr:eventsPendingResponses>....</
    dcr:eventsPendingResponses>
    <dcr:stateAccepting> </dcr:stateAccepting>
    <dcr:stateIndex> </dcr:stateIndex>
    <dcr:eventsEnabled>....</dcr:eventsEnabled>
    </dcr:currentState>
  </dcr:runtime>
</dcr:process>
```

The specification section of the XML document for the Case Handling Process shown in the figure 2 is given in listing 2.

Listing 2. DCRG specification in Xml

```
<dcr:specification>
  <dcr:resources>
    <dcr:roles>
      <dcr:role>U</dcr:role>
      <dcr:role>LO</dcr:role>
      <dcr:role>DA</dcr:role>
    </dcr:roles>
    <dcr:principals>
      <dcr:principal>u</dcr:principal>
      <dcr:principal>lo</dcr:principal>
      <dcr:principal>da</dcr:principal>
    </dcr:principals>
    <dcr:events>
      <dcr:event eventId="0" name="Create case" actionId="
      Create case">
      <dcr:event eventId="1" name="Submit" actionId="
      Submit" />
      <dcr:event eventId="2" name="Assign case Id"
      actionId="Assign case Id" />
      <dcr:event eventId="3" name="Edit" actionId="Edit">
      <dcr:event eventId="4" name="Metadata"
      actionId="Metadata" />
      <dcr:event eventId="5" name="Dates available"
      actionId="Dates available" />
      </dcr:event>
    </dcr:events>
  </dcr:specification>
```

```

<dcr:event eventId="9" name="Upload"
  actionId="Upload" />
<dcr:event eventId="10" name="Download"
  actionId="Download" />
</dcr:event>
</dcr:event>
<dcr:event eventId="11" name="Arrange Meeting" actionId
="Submit">
  <dcr:event eventId="12" name="Propose dates-LO"
    actionId="Propose dates-LO" />
  <dcr:event eventId="13" name="Accept LO" actionId="
    Accept LO" />
  <dcr:event eventId="14" name="Accept DA" actionId="
    Accept DA" />
  <dcr:event eventId="15" name="Propose dates-DA"
    actionId="Propose dates-DA" />
</dcr:event>
<dcr:event eventId="16" name="Hold meeting" actionId="
  Hold meeting" />
</dcr:events>
<dcr:actions>
  <dcr:action actionId="Create case" />
  <dcr:action actionId="Submit" />
  <dcr:action actionId="Edit" />
  <dcr:action actionId="Metadata" />
  <dcr:action actionId="Dates available" />....
</dcr:actions>
</dcr:resources>
<dcr:accessControls>
  <dcr:rolePrincipalAssignments>
    <dcr:rolePrincipalAssignment role-name="U">
      <principal>u</principal>
    </dcr:rolePrincipalAssignment>
    <dcr:rolePrincipalAssignment role-name="LO">
      <principal>lo</principal>
    </dcr:rolePrincipalAssignment>
  </dcr:rolePrincipalAssignments>
  <dcr:actionRoleAssignments>
    <dcr:actionRoleAssignment actionId="Submit">
      <dcr:role>U</dcr:role>
    </dcr:actionRoleAssignment>
    <dcr:actionRoleAssignment actionId="Document">
      <dcr:role>U</dcr:role>
      <dcr:role>LO</dcr:role>
      <dcr:role>DA</dcr:role>
    </dcr:actionRoleAssignment> ....
  </dcr:actionRoleAssignments>
</dcr:accessControls>
<dcr:constraintSets>
  <dcr:constraintSet type="condition">
    <dcr:constraint source="1" target="2" />
    <dcr:constraint source="3" target="1" />....
  </dcr:constraintSet>
  <dcr:constraintSet type="response">
    <dcr:constraint source="1" target="2" />....
  </dcr:constraintSet> ....
</dcr:constraintSets>
</dcr:specification>

```

All the prototype tools support the basic DCR Graph notation containing condition, response, include and exclude relations. We are currently working on extending the prototype to support milestone relations and nested events. In Fig. 6, 7, 8, we have illustrated how the execution state of the case-handling process may be visualized in the simulator in the future.

The graph in the figure. 6 shows the state after a run where the union started by creating a case: they edited meta-data, indicated the dates they were available and submitted. When LO received the case they assigned their own case ID to it. Some time later LO proposed possible dates for a meeting to DA. DA did not agree with these dates and responded by proposing some of their own. In the graph both Accept LO and Accept DA are included and have a pending response because both LO and DA have proposed dates. Because of




Figure 5. The Graphical Editor




Figure 6. Case Handling Process Runtime

these pending responses Hold meeting is disabled. Because no files have been uploaded to the document yet, Download is also disabled. The listing 3 shows the runtime information for the case handling process from the figure 6.

```

Listing 3. DCRG Runtime in Xml
<dcr:runtime processInstanceId="">
  <dcr:executionTrace>4,5,4,1,2,12,15</
  dcr:executionTrace>
  <dcr:currentState stateId="S6">

```

```

<dcr:eventsIncluded>2,4,5,7,9,10,12,13,14,15,16</
dcr:eventsIncluded>
<dcr:eventsExecuted>1,2,4,5,12,15</
dcr:eventsExecuted>
<dcr:eventsPendingResponses>13,14,16</
dcr:eventsPendingResponses>
<dcr:stateAccepting>0</dcr:stateAccepting>
<dcr:stateIndex>0</dcr:stateIndex>
<dcr:eventsEnabled>1,2,4,5,7,8,12,13,14,15</
dcr:eventsEnabled>
</dcr:currentState>
</dcr:runtime>

```

The graph in the figure. 7 shows the runtime state after the union has uploaded an agenda for the meetings. Note that, since the union has uploaded a file to the case, Download is now enabled. But at the same time, Accept LO and Accept DA still remain the same as the previous graph, as the proposed dates have not been accepted yet by either LO or DA.




Figure 7. Case Handling Process Runtime After Upload Document

Figure 8 shows the graph representing the state after LO has accepted one of the dates proposed by DA. Note that both Accept LO and Accept DA are excluded due to the mutual exclude relation between them. Even though there is a pending response on Accept DA, it is not considered relevant as it is excluded and Hold meeting has become pending because of the response relation. Continuing by executing Hold meeting as LO will cause the graph to reach an accepting state, as there will be no included pending responses.

V. COMPARISON TO OTHER APPROACHES

As already mentioned in the introduction, our approach is closely related to the work on DECLARE [27], [28]. In particular the condition and response relations are also considered in [27], [28], and we have used the same graphical notation as loc. cit. The crucial difference is that we focus on a few core constraints allowing to describe the state and operational semantics of processes as a labelled transitions between simple

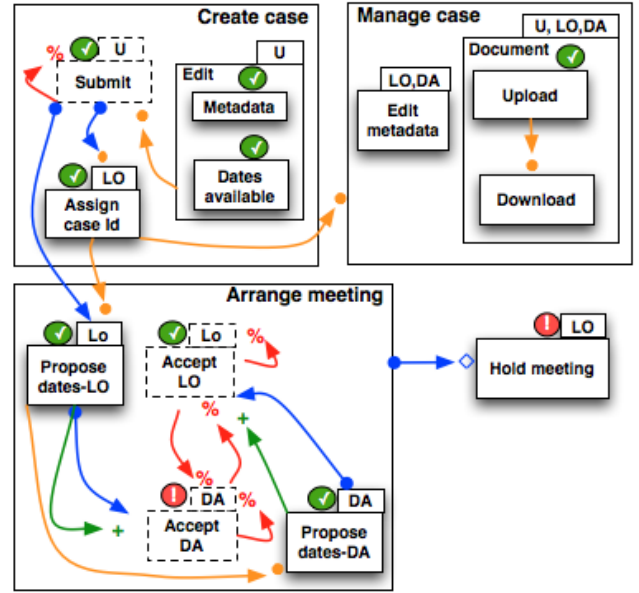


Figure 8. Case Handling Process Runtime After Accept Dates

markings consisting of three sets of respectively executed, included and required events. As also pointed out in [27], [28], the generality of LTL offers much flexibility with respect to specifying execution constraints but makes it more complex to execute processes given in DECLARE and to describe and understand their run-time state. It typically requires a translation of the constraints to LTL and subsequent using the standard mapping of an LTL formula to a Büchi-automaton. In particular, there is no obvious way to trace the graphical constraints in DECLARE to the states of the Büchi-automaton. Moreover, we show in a follow up paper that every set of traces expressible in LTL (and thus DECLARE) can also be expressed using DCR Graphs.

We have shown in [18] it is possible, but much more complex to represent in LTL the interplay between dynamic inclusion/exclusion and the other relations. Neither this novel notion of dynamic inclusion/exclusion relations nor nesting are considered in [27], [28].

The DCR Graphs model also relates to the independent work on the *Guard-Stage-Milestone* model [15] by Hull et al presented as an invited talk at the WS-FM 2010 workshop and part of the work on artifact-centric business processes [2], [5], [7].

Finally, BPMN 2.0 includes the *ad-hoc sub-process activity* which allows one to group a set of activities that can be carried out in an ad-hoc way. Fig. 9 below shows how one may attempt to describe the top level requirements described in Fig. 1 as a BPMN 2.0 ad-hoc sub-activity. According to the informal description of BPMN 2.0 ad-hoc sub-process activity in the current BPMN 2.0 specification ([19]) Create case is a *condition* for Manage case since the latter cannot start without the *data object case* as input, which is produced

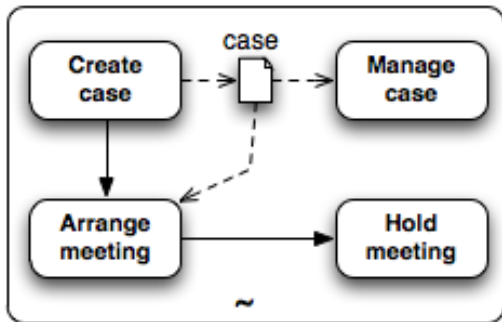


Figure 9. BPMN 2.0 ad-hoc sub-process activity

by the former. Moreover, quoting from the specification, the sequence flow between **Create case** and **Arrange meeting** (and similarly between **Arrange meeting** and **Hold meeting**): “creates a dependency where the performance of the first Task *MUST* be followed by a performance of the second Task. This does not mean that the second Task is to be performed immediately, but there *MUST* be a performance of the second task after the performance of the first Task.”. This seems exactly to correspond to the response relation in DCR Graphs. However, when reading the semantics section of the specification ([19], Sec.13.2.5, 445-446) it appears that the sequence flow introduces just a standard precondition. Thus, the specification is not consistent in the description of sequence flows within ad-hoc sub-activities. Also, it is not clear how to specify roles on actions (swim lanes seem not to be allowed within ad-hoc sub-activities) nor how to specify that an activity within an ad-hoc sub activity only can be executed once. In particular, **Create case** can be executed any number of times in the above process.

VI. CONCLUSIONS AND FUTURE WORK

Our case study showed that the DCR Graphs model is well suited to give a global description of the temporal constraints between the individual tasks which is helpful in capturing the requirements of the overall system.

However, there are still many points for future developments.

First of all there is the need to extend the expressiveness of DCR Graphs. In the ongoing PhD project of the second author we intend to extend the DCR Graphs model to be able to express relevant features such as multi-instance sub-graphs (allowing the dynamic creation of sub-graphs representing dynamic sub process instantiation), time, exceptions and data. Along with this we intend to continue developing the technology for model checking and run time verification and apply it within case studies.

Second, our industrial partner Resultmaker who already use a declarative process model based on the primitives in the DCR Graphs model expects to investigate the use of the formalization to support *safe* dynamic changes to the process constraints at run time.

Thirdly, the DCR Graphs model presently describe a global view of the process. Through our discussions with Exformatics during the case study we identified the wish to be able to automatically synthesize distributed views of the process. In particular, they wanted to be able to derive descriptions of communication protocols and message exchange between the individual local components in a distributed implementation of the system.

Derivations of descriptions of communication protocols between local components from a global model is been researched for the imperative choreography language WS-CDL in the work on structured communication-centred programming for web services by Carbone, Honda and Yoshida [4]. Put briefly, the work formalizes the core of WS-CDL as the global process calculus and define a formal theory of end point projections projecting the global process calculus to abstract descriptions of the behavior of each of the local “end-points” given as pi-calculus processes typed with session types.

We are currently working on the challenge of synthesizing a distributed view of a DCR Graph as a set of interacting DCR Graphs, thus providing a declarative notion of end-point projections. As a challenge for future work we propose to provide a formal map between DCR Graphs and imperative choreographies formalized in the global process calculus [4].

REFERENCES

- [1] Active Endpoints, Adobe Systems, BEA Systems, IBM, Oracle, SAP. Ws-bpel extension for people (bpel4people) version 1.0, 2007. http://www.adobe.us/content/dam/Adobe/en/devnet/livecycle/pdfs/bpel4people_spec.pdf.
- [2] Kamal Bhattacharya, Cagdas Gerede, Richard Hull, Rong Liu, and Jianwen Su. Towards formal analysis of artifact-centric business process models. In *In preparation*, pages 288–304, 2007.
- [3] Christoph Bussler and Stefan Jablonski. Implementing agent coordination for workflow management systems using active database systems. In *Research Issues in Data Engineering, 1994. Active Database Systems. Proceedings Fourth International Workshop on*, pages 53–59, Feb 1994.
- [4] Marco Carbone, Kohei Honda, and Nobuko Yoshida. Structured Communication-Centred Programming for Web Services. In *16th European Symposium on Programming (ESOP’07)*, LNCS, pages 2–17. Springer, 2007.
- [5] David Cohn and Richard Hull. Business artifacts: A data-centric approach to modeling business operations and processes. *IEEE Data Eng. Bull.*, 32(3):3–9, 2009.
- [6] Hasam Davulcu, Michael Kifer, C. R. Ramakrishnan, and I.V. Ramakrishnan. Logic based modeling and analysis of workflows. In *Proceedings of ACM SIGACT-SIGMOD-SIGART*, pages 1–3. ACM Press, 1998.
- [7] Alin Deutsch, Richard Hull, Fabio Patrizi, and Victor Vianu. Automatic verification of data-centric business processes. In *Proceedings of the 12th International Conference on Database Theory, ICDT ’09*, pages 252–267, New York, NY, USA, 2009. ACM.
- [8] Marlon Dumas, Wil M. van der Aalst, and Arthur H. ter Hofstede. *Process Aware Information Systems: Bridging People and Software Through Process Technology*. Wiley-Interscience, 2005.
- [9] Clarence A. Ellis and Gary J. Nutt. Office information systems and computer science. *ACM Comput. Surv.*, 12:27–60, March 1980.
- [10] Clarence A. Ellis and Gary J. Nutt. Workflow: The Process Spectrum. In Amit Sheth, editor, *Proceedings of the NSF Workshop on Workflow and Process Automation in Information Systems*, pages 140–145, May 1996.
- [11] Thomas Hildebrandt. Trustworthy pervasive healthcare processes (TrustCare) research project. Webpage, 2008. <http://www.trustcare.dk/>.
- [12] Thomas Hildebrandt and Raghava Rao Mukkamala. Declarative event-based workflow as distributed dynamic condition response graphs. In *Post-proceedings of PLACES 2010*, 2010.

- [13] Thomas Hildebrandt and Raghava Rao Mukkamala. Distributed dynamic condition response structures. In *Pre-proceedings of International Workshop on Programming Language Approaches to Concurrency and Communication-cEntric Software (PLACES 10)*, March 2010.
- [14] Thomas Hildebrandt, Raghava Rao Mukkamala, and Tijs Slaats. Nested dynamic condition response graphs. In *Fundamentals of Software Engineering Conference 2011 (to appear)*, April 2011.
- [15] Richard Hull. Formal study of business entities with lifecycles: Use cases, abstract models, and results. In *Proceedings of 7th International Workshop on Web Services and Formal Methods*, volume 6551 of *Lecture Notes in Computer Science*, 2010.
- [16] Karen Marie Lyng, Thomas Hildebrandt, and Raghava Rao Mukkamala. From paper based clinical practice guidelines to declarative workflow management. In *Proceedings ProHealth 08 workshop*, 2008.
- [17] Raghava Rao Mukkamala and Thomas Hildebrandt. From dynamic condition response structures to büchi automata. In *Proceedings of 4th IEEE International Symposium on Theoretical Aspects of Software Engineering (TASE 2010)*, August 2010.
- [18] Raghava Rao Mukkamala, Thomas Hildebrandt, and Janus Boris Tøth. The resultmaker online consultant: From declarative workflow management in practice to LTL. In *Proceeding of DDBP*, 2008.
- [19] Object Management Group BPMN Technical Committee. Business Process Model and Notation, version 2.0, 2010. <http://www.omg.org/spec/BPMN/2.0/>.
- [20] Organization for the Advancement of Structured Information Standards (OASIS). Web services human task (ws-humantask) specification, version 1.1, 2009. <http://docs.oasis-open.org/bpel4people/ws-humantask-1.1-spec-cd-06.pdf>.
- [21] Microsoft Research. Zing model checker. Webpage, 2010. <http://research.microsoft.com/en-us/projects/zing/>.
- [22] Pinar Senkul, Michael Kifer, and Ismail H. Toroslu. A logical framework for scheduling workflows under resource allocation constraints. In *In VLDB*, pages 694–705, 2002.
- [23] Munindar P. Singh, Greg Meredith, Christine Tomlinson, and Paul C. Attie. An event algebra for specifying and scheduling workflows. In *Proceedings of DASFAA*, pages 53–60. World Scientific Press, 1995.
- [24] Spin. On-the-fly, ltl model checking with spin. Webpage, 2008. <http://spinroot.com/spin/whatispin.html>.
- [25] Keith D. Swenson. *Mastering the Unpredictable: How Adaptive Case Management Will Revolutionize the Way That Knowledge Workers Get Things Done*. Meghan-Kiffer Press, 2010.
- [26] Wil M. P. van der Aalst and S Jablonski. Dealing with workflow change: Identification of issues and solutions. *International Journal of Computer Systems, Science, and Engineering*, 15(5):267–276, 2000.
- [27] Wil M. P. van der Aalst, Maja Pesic, and Helen Schonenberg. Declarative workflows: Balancing between flexibility and support. *Computer Science - R&D*, 23(2):99–113, 2009.
- [28] Wil M.P van der Aalst and Maja Pesic. A declarative approach for flexible business processes management. In *Proceedings DPM 2006*, LNCS. Springer Verlag, 2006.
- [29] W3C. Web services choreography description language, version 1.0, 2005. <http://www.w3.org/TR/ws-cdl-10/>.
- [30] Glynn Winskel. Event structures. In Wilfried Brauer, Wolfgang Reisig, and Grzegorz Rozenberg, editors, *Advances in Petri Nets*, volume 255 of *Lecture Notes in Computer Science*, pages 325–392. Springer, 1986.
- [31] Glynn Winskel and Mogens Nielsen. Models for concurrency. pages 1–148, 1995.
- [32] M. D. Zisman. *Representation, Specification and Automation of Office Procedures*. Philadelphia, Pa.: University of Pennsylvania, Wharton School, Department of Decision Sciences, Ph.D. Thesis, Sep, 1977.