

From Dynamic Condition Response Structures to Büchi Automata

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

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

Abstract—Recently we have presented *dynamic condition response structures* (DCR) as a declarative process model conservatively generalizing labelled event structures to allow for finite specifications of repeated, possibly infinite behavior and acceptance conditions for infinite runs. The key ideas are to split the causality relation of event structures in two dual relations: the condition relation and the response relation, to split the conflict relation in two relations: the dynamic exclusion and dynamic inclusion, and finally to allow configurations to be multi sets of events. In the present abstract we recall the model and show how to formalize the acceptance condition for infinite runs by giving a map to Büchi-automata as a first step towards automatic verification of processes specified as DCR structures.

I. INTRODUCTION

A key difference between declarative and imperative process languages is that the control flow for the first kind is defined *implicitly* as a set of constraints or rules, and for the latter is defined *explicitly*, e.g. as a flow diagram or a sequence of state changing commands. There is a long tradition for using declarative logic based languages to schedule transactions in the database community. Several authors have noted that it could be an advantage to also use a declarative approach to specify workflow and business processes [3], [7], [8]. An important motivation for considering a declarative approach is to achieve more flexible process descriptions [8]. The increased flexibility is obtained in two ways: Firstly, imperative descriptions tend to over-constrain the control flow, since one does not think of all possible ways of fulfilling the intended constraints, where as the declarative modes specify the process description to the minimum extent, leaving more flexibility to the users. Secondly, adding a new constraint to an imperative process description may require that the process code is completely rewritten, while the declarative approach just requires the extra constraint to be added.

A drawback of the declarative approach however, is that the implicit definition of the control flow makes the flow less easy to perceive for the user or compute by the execution engine. At each state, one has to solve the set of constraints to figure out what are the next possible events. It becomes even worse if you are not only interested in knowing the immediate next event, but also want to get an overview of the complete run of the process.

This motivates researching the problem of finding an expressive declarative process model language that can be easily visualized by the end user, allows an effective run-time scheduling and can be mapped easily to a state based model if an overview of the flow graph is needed. In [4] we have proposed a new such declarative process model language called *dynamic condition response structures* (DCR). The model is inspired by the declarative *process matrix* model language [5], [6] used by our industrial partner and (labelled) prime event structures [9]. Indeed, it is formally a conservative generalization and strict extension of both event structures and the core primitives of the process matrix model language. In [4] we gave a labelled transition semantics for DCR structures and formalized acceptance for finite runs. In the present abstract we first recall the model and then show how to formalize the acceptance condition for infinite runs by giving a map to Büchi-automata. This is a first step towards automatic verification of processes specified as DCR structures.

II. DYNAMIC CONDITION RESPONSE STRUCTURES

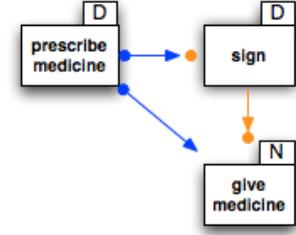
An (labelled, prime) event structure [9] can be regarded as a minimal, declarative model for concurrent processes. It consists of a set of events, a causality (partial order) relation between events stating which events are caused by the previous events (or dually, which events are preconditions for the execution of an event), a conflict relation stating which events can not happen in the same execution and finally a labeling function describing the observable action name of each event. To use it as an execution language for concurrent reactive systems, several aspects are missing. For example, in event structures, each event can only be executed once and they don't have a notion of acceptance condition. In the paper [4], we have consider three of these aspects: Firstly, we allow each event to happen many times and replace the symmetric conflict relation by an asymmetric relation which *dynamically* determines which events are included in or excluded from the structure. Secondly, the causality relation is split in two relations (not necessarily partial orders): A *condition* relation stating which events must have happened before an event and a *response* relation stating which events must happen after (as a response to) an event. We can then define runs to be acceptable if no response event from some point in

the execution is executable continuously without ever being executed. This relates to the elegant definition of fair runs in true concurrency models investigated in [2]. Finally, since we wanted to apply the model to workflow processes we defined *distributed* DCR structures by adding a set of roles assigned to persons/processors and actions.

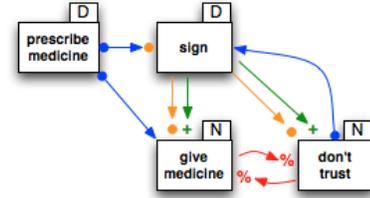
Being based on essentially only four relations between events, our model can be simply visualized as a directed graph with events (labelled by activities and roles) as nodes and four different kinds of arrows. We found that our condition and response relations were two of the core LTL templates used in [8] and thus decided to base our graphical notation on the one suggested in [8].

Before giving the formal definitions let us consider a small part of hospital workflow extracted from a study of paper-based oncology workflow at danish hospitals [5], with a rule stating that the doctor must sign after adding medical prescription to the patient record. A naive imperative process description may instruct the doctor first to prescribe medicine and then sign it. In this way the possibility of adding several prescriptions before or after signing is lost, even if it is perfectly legal according to the declaratively given rule. With respect to the second type of flexibility, consider adding the rule that a nurse should give the prescribed medicine to the patient, but it is not allowed before the patient record has been signed. A simple imperative solution may be to just add a command in the end of the program instructing the nurse to give the medicine. Perhaps we remember to insert a loop to allow that the nurse give the medicine repeatedly. But the nurse should be allowed to give medicine as soon as the first signature is put and the doctor should also be allowed to add new prescriptions after or even at the same time as the nurse gives the medicine. So, the most flexible imperative description should in fact spawn a new thread for the nurse after the first signature has been given.

The example described above is modeled using DCR graphical notation in figure 1(a). It contains three events uniquely labelled (and thus identified) by the actions: `prescribe medicine`, `sign` and `give medicine`. The events are also labelled by the assigned roles (D for Doctor and N for Nurse). The arrow $\bullet \rightarrow \bullet$ between `prescribe medicine` and `sign` indicates that the two events are related by both the condition relation and the response relation. The condition relation ($\rightarrow \bullet$) means that the `prescribe medicine` event must happen at least once before the `sign` event. The response relation ($\bullet \rightarrow$) enforces that, if the `prescribe medicine` event happens, subsequently at some point of time, the `sign` event must happen for the flow to be accepted. Similarly, the response relation between `prescribe medicine` and `give medicine` enforces that, if the `prescribe medicine` event happens, subsequently later the `give medicine` event must happen for the flow to be accepted. Finally, the condition relation between `sign` and `give medicine` enforces that the `sign` event must have happened before the medicine can be given. Note the nurse can give medicine many times, and that the doctor can at any point chose to prescribe new medicine



(a) Prescribe Medicine Example



(b) Prescribe Medicine Example With Check

Figure 1. DCRS example in graphical notation

and sign again. (This will not block the nurse from continue to give medicine. The interpretation is that the nurse may have to keep giving medicine according to the previous prescription).

The dynamic inclusion and exclusion of events is illustrated by an extension to the scenario (also taken from the real case study): If the nurse distrusts the prescription by the doctor, it should be possible to indicate it, and this action should force either a new prescription followed by a new signature or just a new signature. As long the new signature has not been added, medicine must not be given to the patient. This scenario is modeled in Figure 1(b), where one more action `don't trust` is added. Now, the nurse have a choice to indicate distrust of prescription and thereby avoid `give medicine` until the doctor re-execute `sign` action. Executing the `don't trust` action will exclude `give medicine` and makes the `sign` as pending response. So the only way to execute `give medicine` action is to re-execute `sign` action which will then include `give medicine`. Here the doctor may choose to re-do `prescribe medicine` followed by `sign` actions (new prescription) or simply re-do `sign`.

We are now ready for the formal definitions.

Definition 1: A *dynamic condition response structure* (DCR) is a tuple $D = (E, \text{Act}, \rightarrow \bullet, \bullet \rightarrow, \pm, l)$ where

- (i) E is the set of events
- (ii) Act is the set of actions
- (iii) $\rightarrow \bullet \subseteq E \times E$ is the *condition* relation
- (iv) $\bullet \rightarrow \subseteq E \times E$ is the *response* relation
- (v) $\pm : E \times E \rightarrow \{+, \%, *\}$ is the *dynamic inclusion/exclusion* relation.
- (vi) $l : E \rightarrow \text{Act}$ is a labelling function mapping events to actions.

Finally, we define *distributed* dynamic condition response structures by adding roles and principals.

Definition 2: A *distributed dynamic condition response structure*(DDCR) is a tuple

$$(E, \text{Act}, \rightarrow, \bullet, \bullet \rightarrow, \pm, l, R, P, \text{as})$$

where $(E, \text{Act}, \rightarrow, \bullet, \bullet \rightarrow, \pm, l)$ is a dynamic condition response structure, R is a set of *roles*, P is a set of *principals* (e.g. persons/processors/agents) and $\text{as} \subseteq (P \cup \text{Act}) \times R$ is the role assignment relation to executors and actions.

For a *distributed DCR*, the role assignment relation indicates the roles of principals and which roles gives permission to executed which actions. As an example, if *Peter as Doctor* and *Sign as Doctor* (for $Peter \in P$ and $Doctor \in R$, then *Peter* can do the *Sign* action having the role as *Doctor*.

We now define the acceptance criteria for infinite runs by giving a mapping to a Büchi-automaton as follows. The definition is dependent on a fixed order of the finite set of events E of the DCR structure. For an event $e \in E$ we write $\text{rank}(e)$ for its rank in that order and for a subset of events $E' \subseteq E$ we write $\text{min}(E')$ for the event in E' with the minimal rank.

Definition 3: For a finite distributed DCR $D = (E, \text{Act}, \rightarrow, \bullet, \bullet \rightarrow, \pm, l, R, P, \text{as})$ where $E = \{e_1, \dots, e_n\}$ we define the corresponding Büchi-automaton $\text{Aut}(D)$ to be the tuple $(S, s, \rightarrow \subseteq S \times \text{Act} \times S, F)$ where $S = \mathcal{P}(E) \times \mathcal{P}(E) \times \mathcal{P}(E) \times \{1, \dots, n\} \times \{0, 1\}$ is the set of states and $s = (\emptyset, E, \emptyset, 1, 1) \in S$ is the initial state and $F = \mathcal{P}(E) \times \mathcal{P}(E) \times \mathcal{P}(E) \times \{1, \dots, n\} \times \{1\}$ is the set of accepting states. $\rightarrow \subseteq S \times (P \times \text{Act} \times R) \times S$ is the transition relation given by

- $$(E, I, R, i, j) \xrightarrow{(p, a, r)} (E \cup \{e\}, I', R', i', j')$$
- (i) $e \in I, l(e) = a, p \text{ as } r$, and $a \text{ as } r$
 - (ii) $\{e' \in I \mid e' \rightarrow \bullet e\} \subseteq E$
 - (iii) $I' = (I \cup \{e' \mid \pm(e, e') = +\}) \setminus \{e' \mid \pm(e, e') = \% \}$
 - (iv) $R' = (R \setminus \{e\}) \cup \{e' \mid e \bullet \rightarrow e'\}$
 - (v) $j' = 1$ if
 - a) $I' \cap R' = \emptyset$ or
 - b) $\text{min}(M) \in (I \cap R \setminus (I' \cap R')) \cup \{e\}$ or
 - c) $M = \emptyset$ and $\text{min}(I \cap R) \in (I \cap R \setminus (I' \cap R')) \cup \{e\}$
 otherwise $j = 0$.
 - (vi) $i' = \text{rank}(\text{min}(M))$ if $\text{min}(M) \in (I \cap R \setminus (I' \cap R')) \cup \{e\}$ or else
 - (vii) $i' = \text{rank}(\text{min}(I \cap R))$ if $M = \emptyset$ and $\text{min}(I \cap R) \in (I \cap R \setminus (I' \cap R')) \cup \{e\}$ or else
 - (viii) $i' = i$ otherwise.
- for $M = \{e \in I \cap R \mid \text{rank}(e) \} i\}$.

The set E in each state of the automaton records the events that have been executed. The set I records the events that are currently included. The set R records the pending responses. The index i is used to make sure that no event stays forever included and in the response set without being executed. Finally, the flag j indicates if the state is accepting or not.

Condition (i) captures that the executed event must be currently included (i.e. in the set I), and record in the label the principal, action and role.

Condition (ii) captures that all the currently included conditions for the executed event must have been executed.

Condition (iii) captures the dynamic inclusion and exclusion of events.

Condition (iv) removes the currently executed event from the pending response set R and adds new pending responses (which may include the currently executed event as we will see below).

Condition (v) defines when a state is accepting. Either (va) there are no pending responses in the resulting state which are also included. Or (vb), the included pending response with the minimal rank above the index i was excluded or executed. Or (vc), there is no included pending response with rank above the index i and the included pending response with the minimal rank was excluded or executed.

Condition (vi) records the new rank if the resulting state is accepting according to condition (vb). Condition (vii) records the new rank if the resulting state is accepting according to condition (vc).

To give a simple example of the mapping consider the DCR in Fig. 2(a) and the corresponding generalized Büchi-automaton in Fig. 2(b). The structure consists of two events, a and b , having themselves as responses. The accepting runs are thus all infinite runs containing both an infinite number of a events and an infinite number of b events.

The key thing to note is that the automaton enters an accepting state if a pending response is executed, but *only if it is the minimal ranked event according to the index i* . So, if a is executed in state $S4$ we do not enter an accepting state, even if a is a pending response, because event b is also a pending response and it is the one to be executed according to the rank i . Dually, in state $S7$ only an a event will take us to an accepting state, even though b is also a pending response.

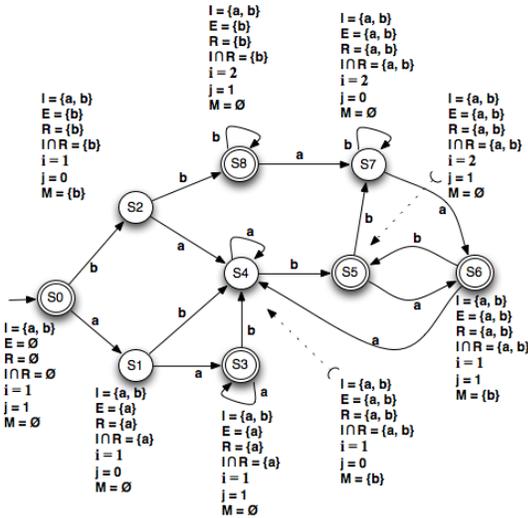
Fig. 2(c) shows an example accepting run, dividing the state sets according to the rank i in order to emphasize the role of the rank in guaranteeing that both a and b events get executed infinitely often.

III. CONCLUSION AND FUTURE WORK

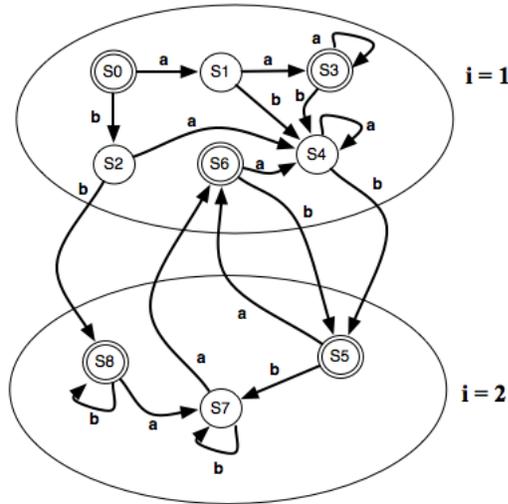
We have shown how to formalize the acceptance condition for infinite runs of DCR structures introduced in [4], a declarative process model derived as a simple generalization of labelled event structures that adds the possibility to specify repeated, possibly infinite behavior and acceptance conditions for infinite runs by allowing events to be executed several times by default, generalizing the conflict relation to a dynamic inclusion/exclusion relation, and employing a response relation dual to that of precedence. We believe that the DCR model and its graphical presentation is much simpler to understand than Büchi-automata or full Linear-time-temporal logic. For future work we will consider a more detailed comparison between DCR and existing models for concurrency, including the relation to the work in [2] and the work on time Petri Nets in [1]. We also plan to study distributed scheduling and extensions to the model, notably with time, nested sub structures, soft constraints, and compensation/exceptions.



(a) Example for infinite runs



(b) State diagram



(c) example accepting run

Figure 2. DCRS example in graphical notation

REFERENCES

- [1] B. Bérard, F Cassez, S. Haddad, D. Lime, and O.H. Roux. Comparison of the expressiveness of timed automata and time petri nets. In *Proceedings of 3rd Int. Conf. on Formal Modeling and Analysis of Timed Systems (FORMATS'05)*, pages 211–225, 2005.
- [2] Allan Cheng. Petri nets, traces, and local model checking. In *Proceedings of AMAST*, pages 322–337, 1995.
- [3] 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.
- [4] 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)*, Paphos, Cyprus, March 2010.
- [5] Karen Marie Lyng, Thomas Hildebrandt, and Raghava Rao Mukkamala. From paper based clinical practice guidelines to declarative workflow management. In *Proceedings ProHealth 08 workshop*, Milan, Italy, 2008.
- [6] 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.
- [7] 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.
- [8] 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.
- [9] 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.