

W.M.P. van der Aalst · M. Pesic · H. Schonenberg

Declarative Workflows

Balancing Between Flexibility and Support

Eingegangen: date / Angenommen: date

Zusammenfassung Insert your German abstract here. Include up to five keywords.

Schlüsselwörter First keyword · Second keyword · More

Abstract Today's process-aware information systems tend to either support business processes *or* provide flexibility. Classical workflow management systems offer good process support as long as the processes are structured and do not require much flexibility. Information systems that allow for flexibility have a tendency to lack process-related support. If systems offer guidance, then they are typically also inclined to "enforce guidelines" and are perceived as inflexible. Moreover, implementing flexible systems is far from trivial. This paper will show that using a more declarative approach can assist in a better balance between flexibility and support. This is demonstrated by presenting the DECLARE framework that aims to take care of the full spectrum of flexibility while at the same time support the user using recommendations and other process-mining-based diagnostics.

Keywords Workflow management · Business Process Management · Flexibility · Process mining

CR Subject Classification H.4.1 · D.2.2 · D.2.11

1 Introduction

Process-aware information systems (PAISs) support operational business processes by combining advances in information technology with recent insights from management science [21]. Workflow Management Systems (WFMSs)

W.M.P. van der Aalst · M. Pesic · H. Schonenberg
Eindhoven University of Technology, P.O. Box 513, NL-5600 MB, Eindhoven, The Netherlands.
Tel.: +31-40-2472733
Fax: +123-40-2463992
E-Mail: {w.m.p.v.d.aalst,m.pesic,h.schonenberg}@tue.nl

are typical examples of such systems. However, many other types of information systems are also "process aware" even if their processes are hard-coded or only used implicitly (e.g., ERP systems). The shift from data orientation to process orientation has increased the importance of PAISs. Moreover, advanced analysis techniques ranging from simulation and verification to process mining and activity monitoring allow for systems that support process improvement in various ways.

When backing business processes with IT there is a difficult trade-off to be made. On the one hand, there is a desire to control processes and to avoid incorrect or undesirable executions of these processes. On the other hand, users want flexible processes that do not constrain them in their actions. This apparent paradox has limited the application of WFMSs thus far, since, as indicated by many authors, WFMSs are too restrictive and have problems when dealing with change [5].

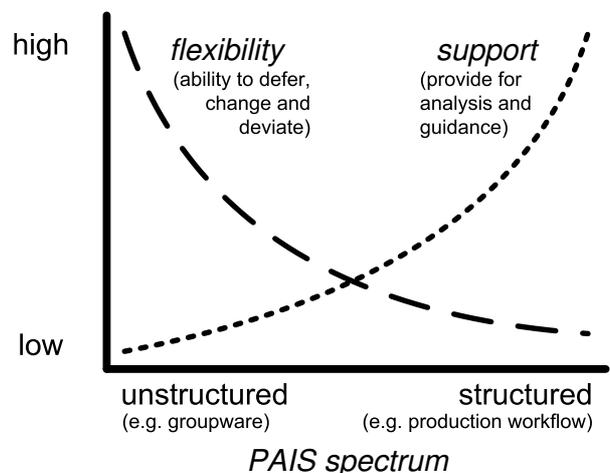


Abb. 1 Today's PAIS spectrum: Systems offer support or flexibility but not both.

Many approaches have been proposed to resolve the apparent paradox illustrated by Figure 1. Some of them try to avoid change, e.g. by generating implicit alternative paths [8; 12], or by deferring the selection of the desired behavior [11]. Others allow for changing the model for a single instance and/or changing a process model while migrating all instances [16; 22; 40; 48]. The migration of process instances from one model to another introduces many interesting problems [5; 16; 40; 48]. For example, the “dynamic change bug” originally described in [22] shows that it may be impossible to put the process instance into a suitable state of the new model without skipping or repeatedly executing activities.

The goal of the paper is to demonstrate that *by using a declarative approach it is possible to balance between support or flexibility*. Traditional approaches tend to use procedural process models to explicitly (i.e., step-by-step) specify the execution procedure. The declarative approach presented in this paper is based on *constraints*, i.e., anything is possible as long as it is not explicitly forbidden. Constraint-based models, therefore, implicitly specify the execution procedure by means of constraints: any execution that does not violate constraints is possible. We have implemented the DECLARE framework that provides for multiple constraint-based languages. By using such a declarative approach, different types of flexibility become possible as shown in the remainder

This paper is organized as follows. Section 2 presents a small “taxonomy of flexibility” using three basic types of flexibility. Moreover, the different types of flexibility are positioned in the overall life-cycle of a process. Using this as a basis, Section 3 discusses the support ultimately desired from a PAIS. Section 4 introduces the DECLARE framework and explains the concept of constraint-based workflow models. Section 5 relates DECLARE to the initial taxonomy presented in Section 2. Section 6 evaluates the support provided by DECLARE and Section 7 discusses the limitations of the current approach and implementation. Section 8 presents related work and Section 9 concludes the paper.

2 Flexibility

Flexibility has become one of the major research topics in the area of workflow management [5]. Today’s WFMSs and many other PAISs have problems providing for flexibility [33; 47]. As a result, these systems are not used to support dynamically changing business processes or the processes are supported in a rigid manner, i.e., changes are not allowed or handled outside of the system. These problems have been described and addressed extensively in literature [8; 10; 22; 35; 38; 40; 42; 48]. Nevertheless, many problems related to flexibility remain unsolved.

In this section, we provide an overview of the different types of flexibility using a taxonomy. This taxonomy is

based on [43]. See also [5; 9; 23; 26; 33; 39; 42; 47] for other classifications of flexibility.

To start, let us identify the different *phases* of a process (instance) in the context of a WFMS:

- *Design time*. At design time a generic process model is created. This model cannot be enacted because it is not connected to some organizational setting.
- *Configuration time*. At configuration time a generic model is made more specific and connected to some organizational context which allows it to be instantiated.
- *Instantiation time*. At instantiation time a process instance is created to handle a particular case (e.g., a customer order or travel request).
- *Run time*. At run time the process instance is executed according to the configured model. The different activities are being enacted as specified.
- *Auditing time*. At auditing time the process instance has completed, however, its audit trail is still available and can be inspected and analyzed.

Flexibility plays a role in most of these phases. At design time some modeling decisions can be postponed to run time. At run time one can decide to deviate from the model and at instantiation time one can change the process model used for the particular instance. When it comes to flexibility, we identify three *flexibility mechanisms*:

- *Defer, i.e., decide to decide later*. This flexibility mechanism deliberately leaves freedom to maneuver at a later phase. Examples are the use of a declarative process modeling language which allows for the “under specification” of processes and the use of late-binding, i.e., the process model has a “hole” that needs to be filled in a later phase.
- *Change, i.e., decide to change model*. Most researchers have addressed flexibility issues by allowing for change. Decisions made at an earlier phase may be revisited. For example, for premium customers the process may be adapted in an ad-hoc manner. The change may refer to the model for a single instance (*ad-hoc change*) or to all future cases (*evolutionary change*). In both cases, a change can create inconsistencies. For example, there may be process instances that cannot be migrated properly when conducting an evolutionary change [2; 16; 22; 38; 40; 48].
- *Deviate, i.e., decide to ignore model*. The third mechanism is to simply deviate from the model, e.g., activities are skipped even if the model does not allow for this to happen. In many environments it is desirable that people are in control, i.e., the system can only suggest activities but not force them to happen.

Figure 2 relates the two dimensions just mentioned. Based on the different phases and the three mechanisms, different types of flexibility are classified. Note that we did not mention any examples of flexibility at auditing

	<i>defer</i> (decide to decide later)	<i>change</i> (decide to change model)	<i>deviate</i> (decide to ignore model)
<i>design time</i>	e.g., defer to run-time by using late binding or declarative modeling	N/A	N/A
<i>configuration time</i>	e.g., defer configuration decisions	e.g., remodel parts of the process at configuration time	e.g., violate a configuration constraint
<i>instantiation time</i>	e.g., defer the selection of parameters or process fragments	e.g., modify model for a particular customer	N/A
<i>run time</i>	N/A	e.g., change model for running instance or migrate instance to new model	e.g., skip or redo a task while this is not specified
<i>auditing time</i>	N/A	N/A	N/A

Abb. 2 Classification of the different types of flexibility based on the phase and mechanism.

time. After the process instance completes, deferring, changing, or deviating is not possible anymore or would imply fraud. Also note that the actual mechanism in different cells can be the same, e.g., deferring a choice to run-time can be done at design time, configuration time and instantiation time.

Figure 2 can be used to characterize the support for flexibility of a concrete WFMS. Unfortunately, today's systems typically provide for only a few forms of flexibility. This is limiting the applicability of PAISs since within one organizations different forms of flexibility may be required.

In the remainder, we will show that it is relatively easy to care for the various types of flexibility by using a declarative approach. However, as indicated in the introduction there may be a trade-off between flexibility and support. Therefore, we first elaborate on the support desired.

3 Support

It is relatively easy to develop systems that that are extremely flexible while providing little process support. A nice example is an e-mail program like Outlook. Outlook is not forcing users to send e-mails at particular times or in a particular order. So it can be seen as a very flexible system as it allows users to execute activities (e.g., sending e-mails) in any way they want. However, Outlook is providing hardly any support and has no knowledge of the processes users are involved in. In this paper, we only consider PAISs where some process model is used during enactment. In fact, we focus on WFMSs. These systems use an explicit process design as a starting point. This design is used at run time to control and support users.

In this section we elaborate on support in the main two phases depicted in Figure 2: *design time* and *run time*.

3.1 Design-time support

At design time a model is constructed that will be used to enact the workflow at run time. Besides providing a good editor to design the workflow process, the WFMS should also provide good analysis tools. Here we elaborate on two types of analysis: *verification* and *performance analysis*.

Verification The goal is to design processes that are correct, e.g., there should not be potential deadlocks, etc. Using verification techniques it is possible to discover problems before the design is enacted. We distinguish between *syntactical* verification and *semantical* verification. Syntactical verification aims at the discovery of errors that do not require domain analysis. For example, a model that can deadlock is incorrect independent of the intent of the process. Similarly, it is not acceptable to have processes that may not terminate or where some activities can never be executed. The notion of *soundness* [1] is a typical example of a correctness property checked through syntactical verification. Semantical verification aims at the analysis of properties that can only be formulated based on domain knowledge. For example, when two drug types are incompatible, then it should not be allowed to make a model in such a way that an instance (i.e., a patient) can receive both drug types during execution. Semantical constraints over processes express domain knowledge. Such knowledge is vital when providing more flexibility, e.g., end users should not modify process models such that essential semantical constraints

get violated. See [31] for examples of semantical correctness notions relevant when adapting processes.

Performance analysis Process models should be correct, but also the performance of a process is relevant. The flow time may be too long or utilization levels and response times can be unacceptable. To predict such situations and to support process improvement, techniques for performance analysis can be used at design time. For example, simulation can be used to predict key performance indicators and to evaluate redesigns. In this paper, we will not elaborate on performance analysis. However, it is important to realize that such functionality may be important when justifying the use of a WFMS.

3.2 Run-time support

At run time, the model made at design time is enacted. Assuming that the model is correct and has a good performance, we distinguish various types of run-time support.

Enforcing correct execution The WFMS should use the model made at design time to enforce the correct execution, i.e., work-items should be offered based on the process model and the system should prevent the execution of activities that are not enabled according to the model.

Recommending effective execution The focus of classical WFMSs is on control, i.e., the system decides on the ordering of activities. When systems allow for more flexibility the decision freedom increases. For example the defer flexibility type (“decide to decide later”) may move choices from design time to run time. To support users when making decisions, the system may provide recommendations. These recommendations may depend on explicit domain knowledge. However, it is also possible to learn good strategies through process mining and then use these for guidance [44]. For example, it is possible to analyze choices made in the past with respect to flow time and then suggest for new cases the “fastest path” through the process. The user can still decide to ignore the recommendation, but at least some decision support is given.

Monitoring process instances The WFMSs should also support the monitoring of process instances. It should be possible to follow specific cases and to look at the workflow at a more aggregate level. Users should be able to get insight into the status of a single case, a group of cases, and the whole workflow.

Learning from processes The more flexible the system is, the more variations are possible when conducting work. Therefore, it is interesting to use process mining [7] to discover the real processes taking place. Starting point for process mining are the event logs, i.e., data on the actual execution of activities. The WFMS should support the recording of these events in a systematic manner and the analysis tools of the WFMS should be able to extract knowledge from these logs. The more variability that is possible, the more valuable such analysis is.

Enforcing correct changes Verification is not only relevant at design time, but also at run time when changes are applied. When using the flexibility type change (“decide to change model”), errors can be introduced into the model. Hence, again syntactical and semantical verification need to be provided to ensure the correct operation of the system.

Note that the above summary of required design-time and run-time support is far from complete. We emphasized aspects related to flexibility, e.g., the monitoring of instances becomes more relevant when users are not forced to work in a particular way.

4 Declare

The goal of this paper is to show that a more declarative approach to workflow management makes it easier to balance between flexibility and support. We will show that it is easy to provide for a wide variety of flexibility types without losing key functionalities at design and/or run time.

To illustrate the advantages of declarative workflows we present a concrete framework: DECLARE. DECLARE is a constraint-based WFMS that provides for multiple declarative languages (DecSerFlow [6], ConDec [34], etc.). In fact, DECLARE is extendible and without any programming it can be configured to support additional constraint-based languages. Instead of explicitly defining the ordering of activities in models, DECLARE models rely on constraints to implicitly determine the possible ordering of activities (any order that does not violate constraints is allowed).

DECLARE can be downloaded from [19]. The architecture of the DECLARE system is shown in Figure 3. The core of the system consists of the following basic components: *Designer*, *Framework* and *Worklist*. The *Designer* component is used for creating the so-called constraint templates, to design concrete process models, and to verify these model. The *Framework* enacts instances of process models. Moreover, it also conducts ad-hoc changes of running instances. While the *Framework* centrally manages the execution of all instances, each user uses his/her *Worklist* component to access active instances. Also, a user can execute activities in active instances in his/her *Worklist*.

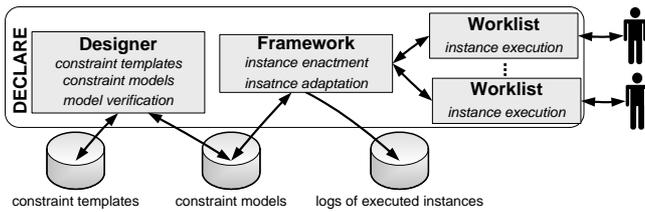


Abb. 3 The architecture of DECLARE.

Unlike most of the approaches, which offer a predefined set of constructs for defining dependencies between activities in process models (e.g., sequence, choice, parallelism, loops, etc.), DECLARE uses a customizable set of arbitrary constructs called *constraint templates*. A declarative language can be seen as a collection of constraint templates. This explains why the DECLARE system provides for multiple languages (e.g. DecSerFlow [6] and ConDec [34]). Templates are defined on the system level in the *Designer* component. Each template has (1) a unique name, (2) a graphical representation, and (3) a formal specification of its semantics in terms of Linear Temporal Logic (LTL) [17; 24]. LTL is a special type of logic which, in addition to classical logical operators, uses temporal operators such as: always (\square), eventually (\diamond), until (\sqcup), and next time (\circ) [17]. Figure 4 shows how the *response* template is defined. This template is graphically represented with a single line between two activities *A* and *B*, with a filled circle next to *A* and a filled arrow next to *B*. The semantics of the template is given as LTL formula $\square(A \Rightarrow \diamond(B))$, i.e., every time activity *A* is executed, *B* has to be executed afterwards at least once.

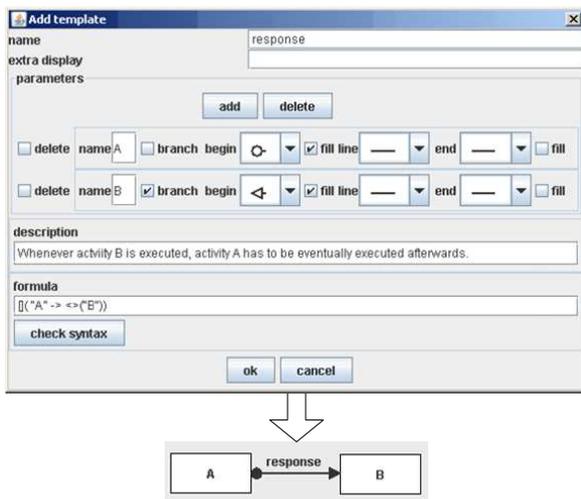


Abb. 4 Constraint template *response*.

In this paper, we use ConDec [34] as the modeling language. ConDec can be seen as a collection of constraint templates. Using these templates one can make concrete models. In this paper, we will refer to such a

ConDec model as “DECLARE model” or simply “model” because we do not elaborate on the different languages supported by DECLARE.

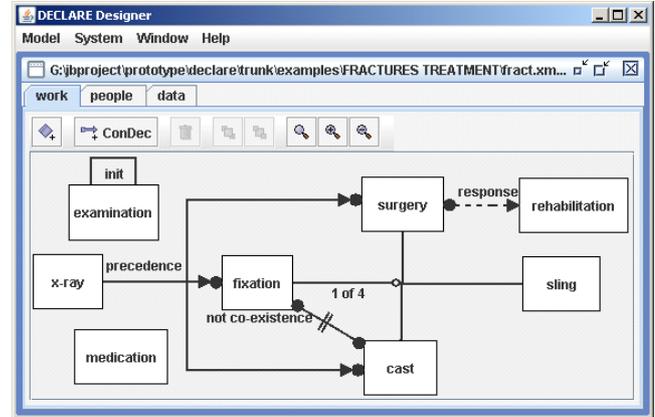


Abb. 5 Declarative model for handling patients using DECLARE.

In the remainder, we use the process for handling a patient at the first aid department in a hospital with a suspicion of a fracture as a running example. Figure 5 shows the model constructed using DECLARE. All activities are depicted as boxes, constraints are depicted by relations (lines) between activities. The model contains *mandatory constraints* (solid lines) and *optional constraints* (dashed lines). For example, the dashed arrow in Figure 5 shows an application of the *response* template defined in Figure 4. We use the following fonts to refer to activity (task) names and constraint template names respectively: *task name* and *constraint name*.

Initially, a specialist performs activity *examination* (constraint *init*), if necessary, additional diagnosis is done by *x-ray*. Depending on the absence, presence and type of fracture, there are several types of treatment available, such as *sling*, *fixation*, *surgery* and *cast*. Except for *cast* and *fixation*, which are mutually exclusive (constraint *not co-existence*), the treatments can be given in any combination and each patient receives at least one treatment (*1 of 4* constraint). Additional diagnosis (*x-ray*) is not necessary when the specialist diagnoses the absence of a fracture during *examination*. Without this additional diagnosis, the patient can only receive the *sling* treatment. All other treatments require *x-ray* to rule out the presence of a fracture, or to decide how to treat the fracture (constraint *precedence*). Simple fractures can be treated just by *cast*. For unstable fractures activity *fixation* may be preferred over activity *cast*. For patients who undergo *surgery* the specialist is advised to execute activity *rehabilitation* afterwards (optional constraint *response*). Moreover, the specialist can provide *medication*, e.g., pain killers or anticoagulants, at any stage of the treatment. Also additional examinations and x-rays can be done during the treatment.

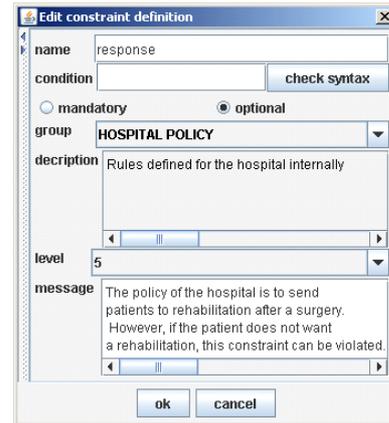
Tabelle 1 LTL expressions for constraints in Figure 5.

Template formula	Constraint	LTL expression
$init(A) = A$	<i>init</i>	<i>examination</i>
$precedence(A, B) = (!B)WA$	<i>precedence</i>	$(!(surgery \vee fixation \vee cast)Wx\text{-ray})$
$response(A, B) = \square(A \Rightarrow (\diamond B))$	<i>response</i>	$\square(surgery \Rightarrow (\diamond rehabilitation))$
$1of4(A, B, C, D) = \diamond(A \vee B \vee C \vee D)$	<i>1 of 4</i>	$\diamond(surgery \vee fixation \vee cast \vee sling)$
$not\ coexistence(A, B) = !((\diamond A) \wedge (\diamond B))$	<i>not-coexistence</i>	$!((\diamond fixation) \wedge (\diamond cast))$

Note that *init*, *precedence*, *1 of 4*, and *not co-existence* refer to constraint templates whose semantics are expressed in terms of LTL. Table 1 shows the relation between the constraints shown in Figure 5, the constraint templates, and LTL. The process should start with *examination*. This constraint is specified using the *init* template. Table 1 shows its definition: $init(A) = A$. Therefore, $init(examination) = examination$. Note that in LTL-terms this means that *examination* should be the current (i.e., first) action. The *precedence* constraint template is defined by the LTL formula $precedence(A, B) = (!B)WA$, i.e., *B* should *not* happen before *A* has happened. Note that *W* is a temporal operator similar to \sqcup (until). The “weak until” operator *W* in “ $(!B)WA$ ” says that *A* does not have to happen if *B* never happens. In Figure 5, the *precedence* constraint template is used with three *B*’s, i.e., $(!(surgery \vee fixation \vee cast)Wx\text{-ray})$ defines the semantics of this particular constraint. This means that the treatments *surgery*, *fixation*, and *cast* all require *x-ray* to rule out the presence of a fracture. However, *x-ray* is not needed if none of the treatment activities (*surgery*, *fixation*, and *cast*) occurs. Table 1 also defines the *1 of 4* and *not co-existence* constraints. $1of4(A, B, C, D) = \diamond(A \vee B \vee C \vee D)$ means that eventually (\diamond) at least one of the four activities should occur. $not\ coexistence(A, B) = !((\diamond A) \wedge (\diamond B))$ means that it cannot (!) be the case that eventually *A* occurs ($\diamond A$) and that eventually *B* occurs ($\diamond B$).

The process defined by Figure 5 allows for many execution paths. Unlike imperative languages, there is no need to include these execution paths explicitly. For example, the mutual exclusion constraint between *cast* and *fixation* is difficult to express in imperative languages, especially since the moment of choice between these two treatments is not fixed. In an imperative language one would need to decide on the moment of choice, specify the loop behavior, and determine the people making these choices. In DECLARE one can simply use the *not-coexistence* constraint with an intuitive graphical notation. In declarative languages only the rules that constrain the behavior need to be specified. Therefore, there is no need to enumerate the execution paths.

Constraint *response* between activities *surgery* and *rehabilitation* is optional as shown by the dashed arrow in Figure 5. Figure 6 shows the definition of the constraint that is using the *response* template. Note that for optional constraints a level and a warning message can be defined. In this particular case a warning of level “5”

**Abb. 6** Defining the optional response constraint.

is generated when the user is about to violate the constraint.

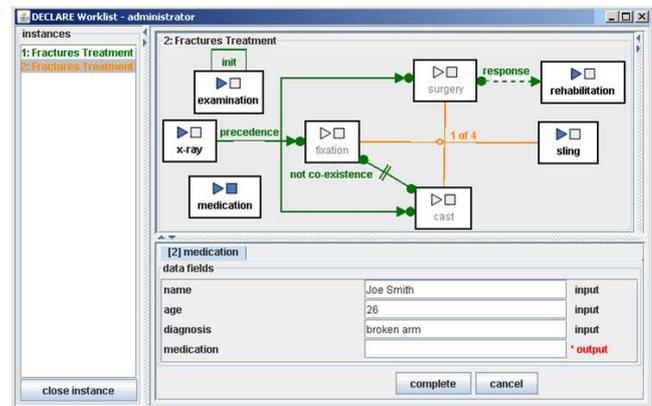
**Abb. 7** Executing an instance.

Figure 7 shows the *Worklist* component containing two active instances (active instances are presented in the list on the left-hand side of the screen). After executing activity *examination*, the user is currently executing activity *medication* for the second process instance. Activities *examination*, *x-ray*, and *medication* are enabled, i.e., can be executed. Activities *surgery*, *fixation*, and *cast* are disabled, i.e., cannot be executed (indicated by the gray color) due to the *precedence* constraint (activity *x-ray* is not executed yet in this instance).

The enabling and execution of activities is driven by constraints: Everything that does not violate the mandatory constraints is enabled for execution and all mandatory constraints must be *satisfied* at the end of the instance execution. For an active instance, each constraint is in one of three states: *satisfied*, *temporarily violated*, or *violated*. In the *violated* state there is no possible future that will satisfy the constraint. In the *satisfied* state there is no need to execute more activities to let the constraint evaluate to true. For example, the *init*, *precedence*, and *response* constraints are *satisfied* in the instance presented in Figure 7. In all other states, the constraint is not (yet) satisfied, but it is still possible to satisfy the constraint by executing the appropriate activities. These states are called *temporarily violated*. For example, constraint 1 of 4 is *temporarily violated* because none of the activities *surgery*, *fixation*, *cast* and *sling* are executed, but executing one of these four activities will bring this constraint into the state *satisfied*.

Optional constraints may become violated. However, this is not possible for mandatory constraints. In order to enforce this, DECLARE constructs an automaton for each constraint [6; 34]. To achieve this we use an approach similar to [17; 24]. Based on the automaton it is possible to see whether the execution of some activity will violate the constraint. Since there is an automaton per constraint, it is easy to see which constraints are violated by which actions. Moreover, DECLARE constructs an overall automaton that is based on the conjunction of all LTL-based constraints. This way it is possible to make sure that mandatory constraints are either in state *satisfied* or state *temporarily violated*. Optional constraints can also be in state *violated*. However, when an optional constraint is about to be violated the system issues a warning (cf. Figure 6).

So each constraint is translated into a finite state automaton that exactly represents all words (i.e., all executions) that satisfy the corresponding LTL formula. For each process instance, DECLARE creates (1) one automaton for the conjunction of LTL formulas of all constraints in the instance, i.e., the so-called instance automaton, and (2) one automaton for each constraint in the instance. As mentioned before, these automata are used for several purposes when executing the instance. First, the instance automaton is used for driving the execution: executing an activity in the instance triggers one or more transitions in the automaton and, thus, causes the state change of the automaton. Second, the instance automaton is used to determine which activities are enabled: only activities that can be triggered at the current state of the automaton are enabled. Third, the instance automaton is used to determine the state of the instance: if the current state of the automaton is accepting, then the instance is *satisfied*; if the current state of the automaton is non-accepting, then the instance is *temporarily violated*. Fourth, the automaton created for each constraint is used to determine the state of the constraint: if

the current state of the automaton is accepting, then the constraint is *satisfied*; if the current state of the automaton is non-accepting, then the constraint is *temporarily violated*. Details about the usage of automata for the execution of instances in DECLARE are out of scope of this paper. We refer the interested reader to [33].

5 Flexibility with Declare

In this section we show that DECLARE cares for all three flexibility types introduced in Section 2, i.e., (1) *defer*, (2) *change*, and (3) *deviate*. To show that this is indeed the case, we first summarize the main differences with imperative/procedural languages using Figure 8.

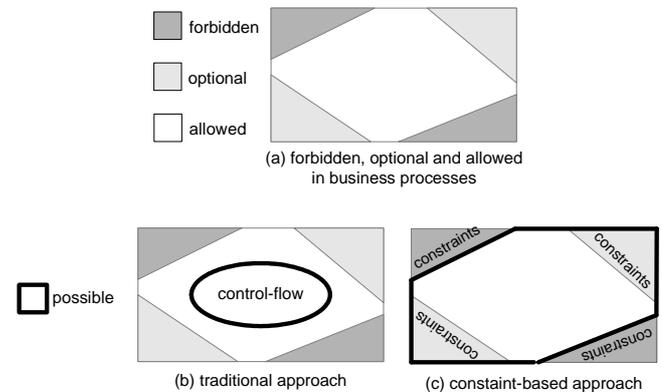


Abb. 8 Mandatory constraints restrict the set of possible behaviors while optional constraints further guide the user. Since traditional approaches explicitly specify the possible behaviors, declarative language tend to be more flexible [33].

The black oval in Figure 8(b) represents the behavioral boundary of a classical (i.e., imperative/procedural) process model that is defined using an “inside out” style of modeling where routing is modeled explicitly. The black thick boundary in Figure 8(c) represents the “outside in” style of modeling supported by DECLARE and the use of optional constraints. Comparison of the two approaches in Figure 8 suggests that a declarative approach indeed allows for more flexibility.

5.1 Defer (“decide to decide later”)

DECLARE facilitates flexibility by the possibility to easily defer choices to run time. Consider for example the *not-coexistence* constraint in Figure 5. This constraint specifies the mutual exclusion of the two activities *cast* and *fixation*. We only specify that if *cast* is done, then *fixation* cannot occur and vice versa. Hence, with little effort, many choices are deferred to run time. For example, *cast* and *fixation* may both occur multiple times or not at all. The only requirement is that for the same instance

they do no both happen. Note that there is no explicit choice between activities *cast* and *fixation*, i.e., there is no activity inserted to make decisions on the number of times each of these activities should be executed. Specifying such a constraint in an imperative language typically results in an “over specification” of the desired behavior, i.e., the user has to explicitly incorporate all execution paths and identify explicit decision points determined at design time.

Figure 5 shows several other constraints that would be difficult to capture in traditional languages without unnecessarily restricting the user at run time. For example, in Figure 5 a specialist can administer *medication* at any moment. The decision *how often* and *when* this can be done is deferred to run time and not included into the design. The only thing that is specified is that *medication* is an activity the specialist can execute in the fracture process and that there are no constraints on this activity.

These examples show that DECLARE definitely supports the flexibility type “decide to decide later” described in Section 2.

5.2 Change (“decide to change model”)

DECLARE allows for instance change at run time to support unforeseen situations or changed circumstances. The change may refer to a single instance (ad-hoc change) or all instances of a given process (evolutionary change). To illustrate the functionality provided by DECLARE, we consider Figure 5. Suppose that new guidelines would forbid specialists to prescribe the usage of *sling* without the execution of *x-ray*. This can be achieved by adding an additional branch to activity *sling* on the *precedence* constraint as shown in Figure 9. This branch can be added to the instance in an ad-hoc manner, i.e., during the execution of the instance. In addition, it is possible to request that the change should be applied to all instances (evolutionary change). This implies that all running instances of the model are migrated (if possible) from the model in Figure 5 to the model in Figure 9. Moreover, all new instances would start in the new model. All of this is possible in DECLARE illustrating that a declarative approach is able to address many of the problems described in literature [2; 22; 40].

DECLARE applies the change only to instances that do not become *violated* because of the migration to the new or adapted process. For example, assume that the ad-hoc change presented in Figure 9 is applied to an instance of the model from Figure 5 for which activities *examination*, *sling*, and *medication* have been executed. Because the added branch specifies that activity *sling* can be executed only after activity *x-ray* was executed, applying this change would lead to a violation because activity *sling* is already executed before activity *x-ray*. The instance automaton (cf. Section 4) is used to check whether an ad-hoc change is applicable for an instance.

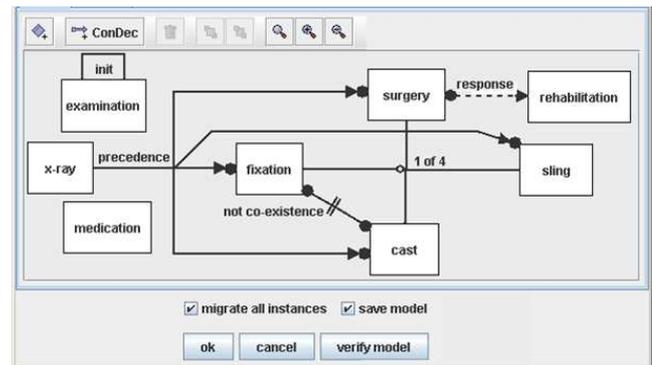


Abb. 9 The initial model is changed by extending the *precedence* constraint. This constraint now also includes activity *sling*.

The change is applicable if and only if the instance history can be “replayed” on the automaton generated for the new model. If the new model is not consistent with the instance history, then the automata generated for subsets of constraints in the new model prevents the ad-hoc change. Indeed, Figure 10 shows that the *precedence* constraint does not allow the change presented in Figure 9.

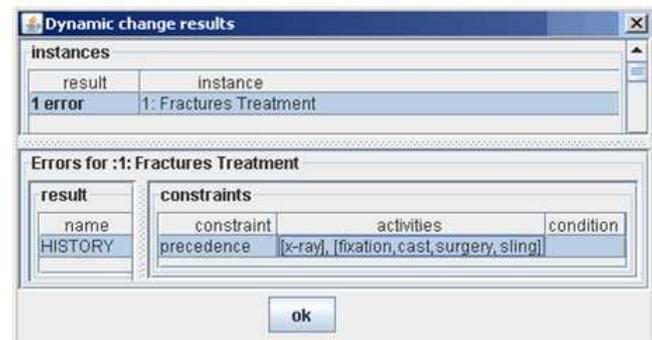


Abb. 10 Report for the ad-hoc change in Figure 9 showing that an instance cannot be migrated because of a violation of the new *precedence* constraint.

5.3 Deviate (“decide to ignore model”)

As mentioned earlier, DECLARE allows for optional constraints. In Figure 5 the *response* constraint is optional as indicated by the dashed line. Figure 6 showed the definition of this constraint indicating the warning that will be given if a users tries to violate this constraint. Figure 8(c) shows the basic idea behind such optional constraints, the user can deviate into the light-gray part of the figure and thus “decide to ignore model”. One can think of optional constraints as guidelines or “soft” constraints. DECLARE allows the designer to set a warning

level to indicate the severity of such a deviation. This way DECLARE also provides for the third type of flexibility identified in Section 2.

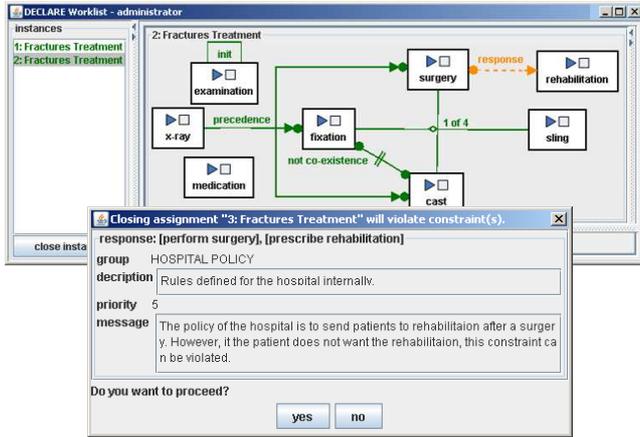


Abb. 11 A warning regarding the possible violation of the optional *response* constraint.

Figure 11 shows a warning shown to a user who is about to violate the optional *response* constraint. This constraint indicates that if *surgery* has been executed, then *rehabilitation* has to be executed afterwards. It is not enforced because when the patient is young and mobile and the severity of fracture is minor this is not strictly necessary. The instance presented in Figure 11 corresponds to a patient for which activity *surgery* is executed but activity *rehabilitation* is not executed yet. Therefore, the *response* constraint is *temporarily violated* (shown using the color orange). If the user would try to close this instance at this moment, this would permanently violate the *response* constraint. If this would be a mandatory constraint, closing the instance would not be possible. However, since the constraint is optional, the user gets an informative warning about the possible violation. Note that the warning shown in Figure 11 is indeed the warning defined in Figure 6.

6 Support by Declare

This section presents the support provided by DECLARE using the topics identified in Section 3. First, we discuss the design-time support offered by DECLARE. Then, we elaborate on DECLARE's run-time support.

6.1 Design-time support

6.1.1 Verification

Constraints in DECLARE models can interfere in subtle ways. This can cause *two types of errors*. First, it might

be that an activity in the model can never be executed. We refer to such an activity as a *dead activity*. Second, it might be that constraints are conflicting, i.e., it is not possible to execute a model in such a way that all mandatory constraints are *satisfied*. We refer to this kind of error as a *conflict*. Each dead activity and conflict is caused by a certain combination of mandatory constraints in the model (i.e., the so-called *cause of error*). Through verification, DECLARE can automatically detect the existence of dead tasks and conflicts while explicitly indicating the combination of mandatory constraints that cause each error. Consider, for example, the model shown in Figure 12. This model contains three errors. First, *fixation* is dead because of the *non co-existence* constraint and the $1..^*$ constraint on activity *cast*, i.e., *cast* should be performed at least once and therefore *fixation* can never occur. Second, *cast* is dead because of the *non co-existence* constraint and the $1..^*$ constraint on activity *fixation*. Third, there is a conflict because the combination of the *non co-existence* constraint, the $1..^*$ constraint on activity *fixation*, and the $1..^*$ constraint on activity *cast* leads to a contradiction.

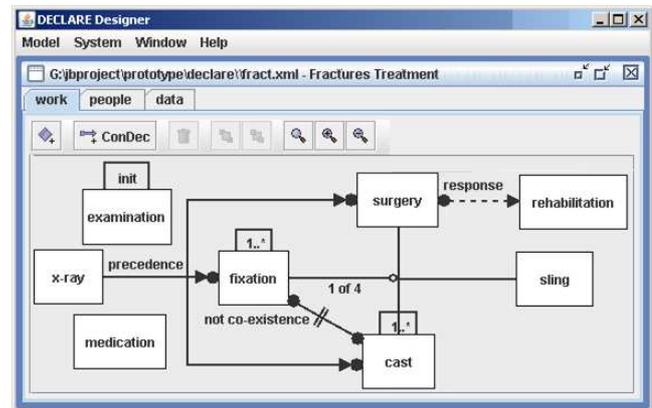


Abb. 12 A model with a conflict.

The three errors in the Figure 12 and the combination of constraints that cause them can be detected using the automata generated from LTL specifications of constraints. In order to find the cause of an error, an automaton is generated for the conjunction of all subsets of mandatory constraints. If an error is found in a minimal subset of constraints, then the error is reported and the root cause is given as a subset of constraints. Because transitions of an automaton enable the execution of activities, a dead activity is detected as an activity that cannot be triggered by any transition of the automaton. Because a model exhibiting a conflict can never be executed such that all related constraints are satisfied, this error is detected by checking whether the corresponding automaton is empty, i.e., if it has no states and no transitions, then the model has a conflict. DECLARE uses this automata-based method to detect errors and the cor-

responding root causes. Figure 13 shows the verification report of DECLARE for the model presented in Figure 12. Indeed, the three errors and combinations of constraints that cause them are detected.

result		constraints		
name	info	constraint	activities	condition
DEAD ACTIVITY	[cast]	not co-existence	[fixation], [cast]	
CONFLICT		1..*	[fixation]	
DEAD ACTIVITY	[fixation]			

(a) activity *cast* is dead

result		constraints		
name	info	constraint	activities	condition
DEAD ACTIVITY	[cast]	not co-existence	[fixation], [cast]	
CONFLICT		1..*	[fixation]	
DEAD ACTIVITY	[fixation]	1..*	[cast]	

(b) a conflict

result		constraints		
name	info	constraint	activities	condition
DEAD ACTIVITY	[cast]	not co-existence	[fixation], [cast]	
CONFLICT		1..*	[fixation]	
DEAD ACTIVITY	[fixation]			

(b) activity *fixation* is dead

Abb. 13 Verification result for the model in Figure 12.

Note that dead activities and conflicts correspond to syntactical errors (cf. Section 3.1). Using LTL it is also easy to check for semantical errors. As long as the (un)desired property can be expressed in terms of LTL, classical model checking techniques can be used to find the error. Note that this is not (yet) implemented in DECLARE. However, syntactical errors are verified both at design time and run time (see also Section 6.2).

6.1.2 Performance analysis

Declarative models are less suitable for performance analysis. The reason is that many execution paths are possible, i.e., there is a tendency to defer decisions to run time. Hence, arbitrary runs of DECLARE may not be representative for actual user behavior. In a procedural language, the users are driven by the process model and it is easy to define the behavior of users as their degree of freedom is limited. This implies that DECLARE could only support performance analysis if a lot of information is added about the expected behavior of users. Therefore, no support for performance analysis at design time is added.

6.2 Run-time support

6.2.1 Enforcing correct execution

As explained in Section 4, DECLARE constructs an instance automaton based on the conjunction of all mandatory constraints. This automaton forces the users to stay within the boundaries of the model. By adding lots of restrictive constraints, the behavior of the DECLARE engine becomes similar to that of a classical workflow engine. However, it is also possible to have just a few mandatory constraints and thus only force the users with respect to these constraints. The run-time view of DECLARE shows the activities that are enabled and indicates the status of each constraint (cf. Figure 7). An activity is not enabled if its execution would permanently violate a mandatory constraint.

6.2.2 Recommending effective execution

The link between the process mining tool PROM [4] and DECLARE enables the guidance of users through recommendations [44]. DECLARE records all events taking place, e.g., starting or completing an activity. These recorded events are shared with PROM. PROM is able to extract knowledge from DECLARE’s event logs using process mining techniques. For example, as shown in [44], PROM can find the paths that minimize costs, flow time, response time, resource utilization, etc. Based on this PROM provides a so-called *recommendation service*. For each running instance, DECLARE continuously asks the recommendation service for suggestions, i.e., all the enabled activities are ranked based on historic information. Note that users do not have to follow recommendations; they merely serve as an advice to support the user. Related to the *recommendation service* is the *prediction service* of PROM that predicts when a case is finished. This information may also be used when selecting an enabled activity.

6.2.3 Monitoring process instances

The monitoring of instances is important to manage the process and to make users aware of the context of their work. Many systems use work distribution mechanisms that lead to “context tunneling”, i.e., the user only sees one particular work-item in isolation without seeing the bigger picture. This is the reason that DECLARE shows the process model augmented with additional information. Using colors it is shown which constraints are satisfied and which activities are enabled (cf. Figure 7). As explained in Section 4 multiple automata are used to determine the state of an instance and states of all constraints. DECLARE relies on PROM for monitoring processes at a higher level [4]. In the future we also hope to use the work-item visualization of YAWL [29]. This visualization allows for the definition of several “maps”

on which work-items may be projected. We believe that the more flexible processes are, the more important it is to visualize “work” while using contextual information.

6.2.4 Learning from processes

DECLARE can export historic event information (e.g., starting and completing activities) by creating so-called MXML logs while executing instances and these can be loaded into PROM. PROM provides numerous plug-ins to extract knowledge from these DECLARE logs [4]. For example, PROM can extract process models based on the work that has actually been done or construct social networks based on the interactions between people. Moreover, DECLARE templates and constraints can be exported into the so-called LTL format readable by PROM. PROM’s LTL checker can then verify properties specified in LTL with respect to some event log [4].

A-posteriori analysis of instances of the model presented in Figure 5 can provide information that can help to improve the model. For example, it is possible to check which instances violated the optional constraint in Figure 5. Moreover, additional properties not specified in the model can be checked. Consider, for example, the verification of instances processed in the past against the following two properties: (1) activity *surgery* was executed in the instance, and (2) activity *rehabilitation* was executed after activity *surgery* in the instance. On the one hand, verification could show that the first property (1) holds in 80% of the instances processed in the previous year, while it holds in only 40% of the instances processed in the year before. This can be an indication that the hospital should hire more surgeons. On the other hand, verification could show that the second property (2) does not hold in 90% of the instances, i.e., the medical staff violated the optional constraint from the model shown in Figure 5 in 90% of the instances. This result may indicate that this constraint should either be removed from the model or made mandatory.

In [28] it is shown that DECLARE models can also be discovered using inductive logic programming. This makes it possible to discover the primary constraints by just observing the process. The discovered models can be loaded into DECLARE.

6.2.5 Enforcing correct changes

As shown in sections 4 and 5, it is possible to change processes. For an ad-hoc change only one instance needs to be migrated, while for an evolutionary change in principle all running instances need to be migrated. In Figure 9 it is shown that after making a choice the new model can be verified and that instances can be migrated. Verification of the new model is done in the same way as described in Section 6.1.1. This way it can be enforced that the new model is syntactically correct, i.e., free of deadlocks, etc. Moreover, using “replay” it is enforced

that only those instances are migrated that actually fit into the new model. Hence, DECLARE ensures syntactical correctness even when processes are changing.

7 Limitations

In this paper, we demonstrated that declarative workflow languages can assist in balancing between flexibility and support. As a proof of concept, the DECLARE system was presented. However, the current approach also has some shortcomings. These are characterized by the following three limitations.

The first limitation is that a constraint-based approach is not very suitable for processes that are of a strict procedural nature. This becomes clear when looking at Figure 8. If the desired control-flow is highly procedural, it is easier to simply describe what should happen rather than describing the constraints that should be satisfied. As shown in Figure 8, the starting point of our constraint-based approach is that anything is allowed unless not explicitly forbidden. In traditional approaches everything is forbidden unless explicitly specified. The later is more suitable if the desired behavior is highly restricted.

The second limitation is that declarative workflow specifications may be less readable if many (interacting) constraints are added. A well-known problem of rule-based systems is the cognitive load on the designer if rules interact in various ways. Similar problems may be encountered if our constraint-based approach is used. However, these problems are (partly) addressed by our graphical language and our verification techniques. Unlike rule-based systems, we provide a powerful graphical notation which immediately shows possible interactions between constraints. Moreover, our verification techniques can be used to easily detect conflicting constraints (cf. Section 6.1.1).

An alternative way to tackle some of the problems associated to the first two limitations is to use composition and mixtures of various languages. In [3; 33] we show that flexibility can be provided as a “service”, i.e., using service-orientation various languages can be mixed. For example, we have integrated YAWL, Worklets, and DECLARE using such an approach. Moreover, composition can be used to define a hierarchy of simpler models (“divide and conquer”).

The third limitation is the efficiency of the current implementation. The current DECLARE engine has problems dealing with large specifications. This is due to the complexity of the model-checking techniques used. So far we have followed a rather straightforward implementation strategy and a much faster implementation is possible, e.g., by a better encoding of the automaton and by caching results. Moreover, the underlying constraint language can be restricted to not use the full power of LTL. By limiting the expressiveness more efficient imple-

mentations come into reach. We expect that, based on focussed research and development efforts, highly performing implementations are possible.

8 Related work

Many researchers have been trying to provide ways of avoiding the apparent paradox where, on the one hand, there is the desire to control the process and to avoid incorrect or undesirable executions of the processes, and, on the other hand, users want lots of flexibility and to feel unconstrained in their actions [2; 5; 6; 8; 11; 12; 16; 22; 38; 40; 48].

See [5; 9; 23; 26; 33; 39; 42; 43; 47] for various taxonomies/classifications of workflow flexibility. The taxonomy in this paper is based on [43].

It is impossible to provide a complete overview of related work. Therefore, we refer to only some of the most related papers in this area. In [47] 18 “change patterns” (e.g. inset/delete process fragments) and 7 “change support features” (e.g., correctness enforcement and instance migration) are identified. The case handling concept is advocated as a way to avoid restricting users in their actions [8]. This is achieved by a range of mechanisms that allow for implicit deviations that are rather harmless. In [12] completely different techniques are used, but also the core idea is that implicit paths are generated to allow for more flexibility. In [9; 11] pockets of flexibility are identified that are specified/selected later in the process, i.e., there is some form of “late binding” at run time. A similar approach was proposed in [42] where the process of defining a change is integrated in the process itself. Many papers look at problems related to ad-hoc and/or evolutionary change [2; 16; 22; 38; 40; 48]. The problem of the dynamic change bug was introduced in [22]. In [2] this problem is addressed by calculating so-called change regions based on the structure of the process. A particular correctness property is described in [48] and the problem of instance migration is also investigated in [16]. In the context of the ADEPT system the problem of workflow change has been investigated in detail (including data analysis) [38; 39; 40]. In [41] a holistic approach is given which combines the ADEPT framework for process change with case-based reasoning technology to learn from change. Case based reasoning was also proposed in [32] to reuse past experiences with change. Here a suspension mechanism was proposed that allows the designer to modify suspended parts of the workflow while other parts continue to be executed.

Another popular stream of research is applying rule-based or constraint-based process modeling languages [25; 20; 46] that are able to offer multiple execution alternatives and, therefore, can enhance flexibility by design. In [25], Glance et al. use process grammars for definition of rules involving activities and documents. Process models are executed via execution of rules that trigger

each other. The Freeflow prototype presented in [20] uses constraints for building declarative process models. Freeflow constraints represent dependencies between states (e.g., inactive, active, disabled, enabled, etc.) of different activities, i.e., an activity can enter a specific state only if another activity is in a certain state. Some approaches consider process models based on dependencies between events involving activities [18; 46]. For example, the constraint-based language presented in [46] uses rules involving (1) preconditions that must hold before an activity can be executed, (2) postconditions that must hold after an activity is executed and (3) additional conditions that must hold in general before or after an activity is executed. A similar idea was presented in [27] by Joeris, who proposes flexible workflow enactment based on event-condition-action (ECA) rules. In [30], a temporal constraint network is proposed for business process execution. The authors use thirteen temporal intervals defined by Allen [13] (e.g., *before*, *meets*, *during*, *overlaps*, *starts*, *finishes*, *after*, etc.) to define selection constraints (which define activities in a process) and scheduling constraints (which define when these activities should be executed). Several approaches propose using *intertask dependencies* for specification of the process models. In [14; 15], Attie et al. propose using Computational Tree Logic (CTL) for the specification of intertask dependencies amongst different unique events (e.g., commit dependency, abort dependency, conditional existence dependency, etc.). Dependencies are transformed into automata, which are used by a central scheduler to decide if particular events are accepted, delayed or rejected. In [36; 37], Raposo et al. propose a larger set of basic interdependencies and propose modeling their coordination using Petri nets.

It is also interesting to mention some commercial WFMSs in this context. Historically, InConcert of Xerox and Ensemble of FileNet were systems among the first commercial systems to address the problem of change. Both supported ad-hoc changes in a rather restrictive setting. Several systems have been extended with some form of late binding. For example, the Staffware workflow system allows for the dynamic selection of subprocesses at run time. Probably the most flexible commercial system is FLOWer of Pallas Athena [8]; this system supports a variety of case handling mechanisms to enable flexibility at run time while avoiding changes of the model.

This paper is based on the earlier work on DECLARE and related languages such as ConDec and DecSerFlow [6; 33; 34; 45; 35]. In [3] it is explained how DECLARE can be combined with other (not constraint-based) approaches.

9 Conclusion

In this paper, we advocated the use of declarative language as a means to balance between flexibility and support. The ideas have been implemented which resulted in DECLARE, an open source workflow management system. DECLARE provides a wide range of flexibility mechanisms: *defer* (decide to decide later), *change* (decide to change model), and *deviate* (decide to ignore model). Both ad-hoc and evolutionary change are supported and models can be verified. As far as we know, there is no other workflow management system that supports such a wide range of flexibility mechanisms. In this paper, we also elaborated a bit on the link with process mining. The connection between PROM and DECLARE enables innovate means of analysis, e.g., users may get recommendations based on historic information.

Literatur

1. W.M.P. van der Aalst. The Application of Petri Nets to Workflow Management. *The Journal of Circuits, Systems and Computers*, 8(1):21–66, 1998.
2. W.M.P. van der Aalst. Exterminating the Dynamic Change Bug: A Concrete Approach to Support Workflow Change. *Information Systems Frontiers*, 3(3):297–317, 2001.
3. W.M.P. van der Aalst, M. Adams, A.H.M. ter Hofstede, M. Pesic, and H. Schonenberg. Flexibility as a Service. BPM Center Report BPM-08-09, BPMcenter.org, 2008.
4. W.M.P. van der Aalst, B.F. van Dongen, C.W. Günther, R.S. Mans, A.K. Alves de Medeiros, A. Rozinat, V. Rubin, M. Song, H.M.W. Verbeek, and A.J.M.M. Weijters. ProM 4.0: Comprehensive Support for Real Process Analysis. In J. Kleijn and A. Yakovlev, editors, *Application and Theory of Petri Nets and Other Models of Concurrency (ICATPN 2007)*, volume 4546 of *Lecture Notes in Computer Science*, pages 484–494. Springer-Verlag, Berlin, 2007.
5. W.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.
6. W.M.P. van der Aalst and M. Pesic. DecSerFlow: Towards a Truly Declarative Service Flow Language. In M. Bravetti, M. Nunez, and G. Zavattaro, editors, *International Conference on Web Services and Formal Methods (WS-FM 2006)*, volume 4184 of *Lecture Notes in Computer Science*, pages 1–23. Springer-Verlag, Berlin, 2006.
7. W.M.P. van der Aalst, H.A. Reijers, A.J.M.M. Weijters, B.F. van Dongen, A.K. Alves de Medeiros, M. Song, and H.M.W. Verbeek. Business Process Mining: An Industrial Application. *Information Systems*, 32(5):713–732, 2007.
8. W.M.P. van der Aalst, M. Weske, and D. Grünbauer. Case Handling: A New Paradigm for Business Process Support. *Data and Knowledge Engineering*, 53(2):129–162, 2005.
9. M. Adams. *Facilitating Dynamic Flexibility and Exception Handling for Workflows*. Phd thesis, Queensland University of Technology, 2007.
10. M. Adams, A.H.M. ter Hofstede, W.M.P. van der Aalst, and D. Edmond. Dynamic, Extensible and Context-Aware Exception Handling for Workflows. In F. Curbera, F. Leymann, and M. Weske, editors, *Proceedings of the OTM Conference on Cooperative Information Systems (CoopIS 2007)*, volume 4803 of *Lecture Notes in Computer Science*, pages 95–112. Springer-Verlag, Berlin, 2007.
11. M. Adams, A.H.M. ter Hofstede, D. Edmond, and W.M.P. van der Aalst. Worklets: A Service-Oriented Implementation of Dynamic Flexibility in Workflows. In R. Meersman and Z. Tari et al., editors, *On the Move to Meaningful Internet Systems 2006, OTM Confederated International Conferences, 14th International Conference on Cooperative Information Systems (CoopIS 2006)*, volume 4275 of *Lecture Notes in Computer Science*, pages 291–308. Springer-Verlag, Berlin, 2006.
12. A. Agostini and G. De Michelis. Improving Flexibility of Workflow Management Systems. In W.M.P. van der Aalst, J. Desel, and A. Oberweis, editors, *Business Process Management: Models, Techniques, and Empirical Studies*, volume 1806 of *Lecture Notes in Computer Science*, pages 218–234. Springer-Verlag, Berlin, 2000.
13. J.F. Allen. Maintaining Knowledge About Temporal Intervals. *Communications of the ACM*, 26(11):832–843, 1983.
14. P.C. Attie, M.P. Singh, E.A. Emerson, A. Sheth, and M. Rusinkiewicz. Scheduling Workflows by Enforcing Intertask Dependencies. *Distributed Systems Engineering Journal*, 3(4):222–238, December 1996.
15. P.C. Attie, M.P. Singh, A. Sheth, and M. Rusinkiewicz. Specifying and Enforcing Intertask Dependencies. In *19th International Conference on Very Large Data Bases (VLDB)*, pages 134–145, San Francisco, CA, USA, 1993. Morgan Kaufmann Publishers Inc.
16. F. Casati, S. Ceri, B. Pernici, and G. Pozzi. Workflow Evolution. *Data and Knowledge Engineering*, 24(3):211–238, 1998.
17. E.M. Clarke, O. Grumberg, and D.A. Peled. *Model Checking*. The MIT Press, Cambridge, Massachusetts and London, UK, 1999.
18. G. Decker, A. Grosskopf, and A. Barros. A Graphical Notation for Modeling Complex Events in Business Processes. In *Proceedings of the 11th IEEE International Enterprise Distributed Object Computing Conference (EDOC 2007)*, pages 27–36. IEEE Computer Society, 2007.
19. Declare, 2008. <http://declare.sf.net>.
20. P. Dourish, J. Holmes, A. MacLean, P. Marqvardsen, and A. Zbyslaw. Freeflow: Mediating Between Representation and Action in Workflow Systems. In *Proceedings of the CM Conference on Computer Supported Cooperative Work (CSCW '96)*, pages 190–198. ACM Press, New York, NY, USA, 1996.
21. M. Dumas, W.M.P. van der Aalst, and A.H.M. ter Hofstede. *Process-Aware Information Systems: Bridging People and Software through Process Technology*. Wiley & Sons, 2005.
22. C.A. Ellis, K. Keddara, and G. Rozenberg. Dynamic Change within Workflow Systems. In N. Comstock, C. Ellis, R. Kling, J. Mylopoulos, and S. Kaplan, editors, *Proceedings of the Conference on Organizational Computing Systems*, pages 10 – 21, Milpitas, California, August 1995. ACM SIGOIS, ACM Press, New York.
23. D. Georgakopoulos, M. Hornick, and A. Sheth. An Overview of Workflow Management: From Process Modeling to Workflow Automation Infrastructure. *Distributed and Parallel Databases*, 3:119–153, 1995.
24. D. Giannakopoulou and K. Havelund. Automata-Based Verification of Temporal Properties on Running Programs. In *ASE '01: Proceedings of the 16th IEEE international conference on Automated software engineering*, page 412, Washington, DC, USA, 2001. IEEE Computer

- Society.
25. N. Glance, D. Pagani, and R. Pareschi. Generalised Process Structure Grammars (GPSG) for Flexible Representations of Work. In *Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW96)*, pages 190–198. ACM Press, New York, NY, USA, 1996.
 26. P. Heinl, S. Horn, S. Jablonski, J. Neeb, K. Stein, and M. Teschke. A Comprehensive Approach to Flexibility in Workflow Management Systems. In G. Georgakopoulos, W. Prinz, and A.L. Wolf, editors, *Work Activities Coordination and Collaboration (WACC'99)*, pages 79–88, San Francisco, February 1999. ACM press.
 27. G. Joeris. Decentralized and Flexible Workflow Enactment Based on Task Coordination Agents. In B. Wangler and L. Bergman, editors, *Proceedings of the 12th International Conference on Advanced Information Systems Engineering (CAiSE00)*, volume 1789 of *Lecture Notes in Computer Science*, pages 41–62, Stockholm, Sweden, September 2000. Springer-Verlag, Berlin.
 28. E. Lamma, P. Mello, M. Montali, F. Riguzzi, and S. Storari. Inducing Declarative Logic-Based Models from Labeled Traces. In G. Alonso, P. Dadam, and M. Rosemann, editors, *International Conference on Business Process Management (BPM 2007)*, volume 4714 of *Lecture Notes in Computer Science*, pages 344–359. Springer-Verlag, Berlin, 2007.
 29. M. de Leoni, W.M.P. van der Aalst, and A.H.M. ter Hofstede. Visual Support for Work Assignment in Process-Aware Information Systems. In M. Dumas, M. Reichert, and M.C. Shan, editors, *International Conference on Business Process Management (BPM 2008)*, volume 5240 of *Lecture Notes in Computer Science*, pages 67–83. Springer-Verlag, Berlin, 2008.
 30. R. Lu, S. Sadiq, V. Padmanabhan, and G. Governatori. Using a temporal constraint network for business process execution. In *In proceedings of the 17th Australasian Database Conference (ADC '06)*, pages 157–166, Darlinghurst, Australia, 2006. Australian Computer Society, Inc.
 31. L.T. Ly, S. Rinderle, and P. Dadam. Semantic Correctness in Adaptive Process Management Systems. In *Business Process Management*, volume 4102 of *Lecture Notes in Computer Science*, pages 193–208. Springer-Verlag, Berlin, 2006.
 32. M. Minor, D. Schmalen, Koldehoff, and R. Bergmann. Structural Adaptation of Workflows Supported by a Suspension Mechanism and by Case-Based Reasoning. In *Proceedings of WETICE 2007*, pages 370–375, 2007.
 33. M. Pesic. *Constraint-based Workflow Management Systems: Shifting Control to Users*. Phd thesis, Eindhoven University of Technology, May 2008.
 34. M. Pesic, H. Schonenberg, and W.M.P. van der Aalst. DECLARE: Full Support for Loosely-Structured Processes. In M. Spies and M.B. Blake, editors, *Proceedings of the Eleventh IEEE International Enterprise Distributed Object Computing Conference (EDOC 2007)*, pages 287–298. IEEE Computer Society, 2007.
 35. M. Pesic, M. H. Schonenberg, N. Sidorova, and W.M.P. van der Aalst. Constraint-Based Workflow Models: Change Made Easy. In F. Curbera, F. Leymann, and M. Weske, editors, *Proceedings of the OTM Conference on Cooperative information Systems (CoopIS 2007)*, volume 4803 of *Lecture Notes in Computer Science*, pages 77–94. Springer-Verlag, Berlin, 2007.
 36. A.B. Raposo and H. Fuks. Defining Task Interdependencies and Coordination Mechanisms for Collaborative Systems. In M. Blay-Fornarino, A.M. Pinna-Dery, K. Schmidt, and P. Zaratè, editors, *Cooperative Systems Design*, volume 74 of *Frontiers in Artificial Intelligence and Applications*, pages 88–103, Amsterdam, The Netherlands, 2002. IOS Press.
 37. A.B. Raposo, L.P. Magalhaes, I.L.M. Ricarte, and H. Fuks. Coordination of Collaborative Activities: A Framework for the Definition of Tasks Interdependencies. In *Proceedings of the 7th International Workshop on Groupware (CRIWG)*, pages 170–179. IEEE Computer Society, 2001.
 38. M. Reichert and P. Dadam. ADEPTflex: Supporting Dynamic Changes of Workflow without Loosing Control. *Journal of Intelligent Information Systems*, 10(2):93–129, 1998.
 39. S. Rinderle, M. Reichert, and P. Dadam. Evaluation of Correctness Criteria for Dynamic Workflow Changes. In W.M.P. van der Aalst, A.H.M. ter Hofstede, and M. Weske, editors, *International Conference on Business Process Management (BPM 2003)*, volume 2678 of *Lecture Notes in Computer Science*, pages 41–57. Springer-Verlag, Berlin, 2003.
 40. S. Rinderle, M. Reichert, and P. Dadam. Correctness Criteria For Dynamic Changes in Workflow Systems: A Survey. *Data and Knowledge Engineering*, 50(1):9–34, 2004.
 41. S. Rinderle, B. Weber, M. Reichert, and W. Wild. Integrating Process Learning and Process Evolution: A Semantics Based Approach. In W.M.P. van der Aalst, A.H.M. ter Hofstede, and M. Weske, editors, *International Conference on Business Process Management (BPM 2005)*, volume 2678 of *Lecture Notes in Computer Science*, pages 252–267. Springer-Verlag, Berlin, 2005.
 42. S. Sadiq, W. Sadiq, and M. Orlowska. Pockets of Flexibility in Workflow Specification. In *Proceedings of the 20th International Conference on Conceptual Modeling (ER 2001)*, volume 2224 of *Lecture Notes in Computer Science*, pages 513–526. Springer-Verlag, Berlin, 2001.
 43. H. Schonenberg, R. Mans, N. Russell, N. Mulyar, and W.M.P. van der Aalst. Process Flexibility: A Survey of Contemporary Approaches. In J. Dietz, A. Albani, and J. Barjis, editors, *Advances in Enterprise Engineering I*, volume 10 of *Lecture Notes in Business Information Processing*, pages 16–30. Springer-Verlag, Berlin, 2008.
 44. H. Schonenberg, B. Weber, B.F. van Dongen, and W.M.P. van der Aalst. Supporting Flexible Processes Through Recommendations Based on History. In M. Dumas, M. Reichert, and M.C. Shan, editors, *International Conference on Business Process Management (BPM 2008)*, volume 5240 of *Lecture Notes in Computer Science*, pages 51–66. Springer-Verlag, Berlin, 2008.
 45. M.H. Schonenberg, R.S. Mans, N.C. Russell, N.A. Mulyar, and W.M.P. van der Aalst. Towards a Taxonomy of Process Flexibility (Extended Version). BPM Center Report BPM-07-11, BPMcenter.org, 2007.
 46. J. Wainer and F. de Lima Bezerra. Constraint-based flexible workflows. In *Proceedings of the 9th International Workshop on Groupware: Design, Implementation, and Use (CRIWG 2003)*, volume 2806, pages 151 – 158. Springer-Verlag, Berlin, 2003.
 47. B. Weber, M. Reichert, and S. Rinderle-Ma. Change Patterns and Change Support Features: Enhancing Flexibility in Process-Aware Information Systems. *Data and Knowledge Engineering*, 66(3):438–466, 2008.
 48. M. Weske. Formal Foundation and Conceptual Design of Dynamic Adaptations in a Workflow Management System. In R. Sprague, editor, *Proceedings of the Thirty-Fourth Annual Hawaii International Conference on System Science (HICSS-34)*. IEEE Computer Society Press, Los Alamitos, California, 2001.

Wil van der Aalst is a full professor of Information Systems at the Technische Universiteit Eindhoven and an adjunct professor at Queensland University of Technology. His research interests include workflow management, process mining, Petri nets, business process management, process modeling, and process analysis. Many of his papers are highly cited (he has an H-index of 55 according to Google Scholar) and his ideas have influenced researchers, software developers, and standardization committees working on process support. For more information about his work visit: www.workflowpatterns.com, www.workflowcourse.com, www.processmining.org, www.yawl-system.com, www.wvdaalst.com.

Maja Pesic studied information systems at University of Belgrade, Serbia. She received a Ph.D. degree in 2008 at Eindhoven University of Technology, The Netherlands. Her research in the field of workflow management systems aims at increasing the flexibility of these systems. Currently, she is working as a researcher (Post-doc) at the department of Mathematics and Computer Science at the Eindhoven University of Technology, The Netherlands.

Helen Schonenberg is currently working as Ph.D. student in the group of professor Wil van der Aalst at Eindhoven University of Technology, The Netherlands. Her research focus is on supporting the execution of flexible processes by means of recommendations and predictions, using process mining techniques. Schonenberg received her Master of Science in Computer Science (Ingenieur) in 2006 from the University of Twente, The Netherlands.