

Product-Based Workflow Support

Irene Vanderfeesten ^{a,*} Hajo A. Reijers ^a
Wil M.P. van der Aalst ^a

^a*Technische Universiteit Eindhoven, Department of Industrial Engineering and Innovation Sciences, PO Box 513, NL-5600 MB Eindhoven, The Netherlands*

Abstract

Despite the industrial need for the improvement of information-intensive business processes, few scientifically grounded approaches exist to support such initiatives. In this paper, we propose a new approach that builds on concepts that are part of a product-oriented view on process optimization. Essentially, this approach allows end users to flexibly decide on the best possible way to create an informational product within the limits that are imposed by regulations and logical dependencies. We argue that this provides various benefits in comparison to earlier work. To support the end user in making sensible decisions, we describe two alternative approaches to provide her with recommendations to this end. We formalize these alternatives and discuss their relative strengths and weaknesses. The feasibility of the overall approach, which we refer to as Product-Based Workflow Support, is demonstrated by a workflow system realized using ProM and DECLARE.

Key words: Business Process Modelling, Workflow Management, Product Data Model.

1 Introduction

Contemporary management concepts like “operational excellence”, “lean management”, and “business process redesign” all stress the importance of smoothly running business processes. It seems a natural angle to consider *processes* – complete chains of operations that are needed to produce certain products or

* This research is supported by the Technology Foundation STW, applied science division of NWO and the technology programme of the Dutch Ministry of Economic Affairs.

* Corresponding author. Tel. 0031 40 247 4366, Fax. 0031 40 243 2612
Email address: i.t.p.vanderfeesten@tue.nl (Irene Vanderfeesten).

services – to make organizations perform better. Unsurprisingly, market analyses consistently identify the improvement of business processes as the top business priority for CIO’s [20,21,22].

Given the importance of business processes and their tight relation to organizational performance, it may come as a surprise that few scientific approaches are available that address the issue of how to actually *design* a process or, as in many processes are already in place, how to *redesign* one. The best-known references are situated in the domain of the popular management literature, e.g. [11,14,17]. Understandably, it is often said that process design is “more art than science” [42,43].

One of the notable exceptions is Product-Based Workflow Design (PBWD) [35]. PBWD has been developed in close cooperation between academic and industrial parties to arrive at a method for process redesign that is repeatable, objective, and effective. The focus is on the design of processes that deliver informational products, the so-called *workflow processes*. Since its conception, PBWD has been adopted by various consultancy and service companies to improve the performance of various business processes in the services domain [33,34].

Highly characteristic for PBWD is that it aims first and foremost at developing a deep understanding of the characteristics of the *informational product* that is to be delivered, e.g. a particular type of decision, proposal, permit, etc., which is laid down in a *product data model*. This is subsequently used by the designer to determine the best process structure to create and deliver that product. Given that, in general, there are alternative ways to produce an informational product, PBWD discloses the various opportunities to produce a product a desirable way.

At this stage, considerable experience has been gained with the application of PBWD in practice. Aside from the tangible business benefits that PBWD has delivered, it has become apparent that the “product” notion is an extremely viable concept to reason about workflow processes. The shift of attention to what is the desired *outcome* of a workflow process without directly discussing *how* this is achieved leads to an interaction with stakeholders that quickly converges. This sharply contrasts with the problems that are often associated with process improvement projects, such as the confusion about what actions in the current process are really necessary and which ones are merely motivated by tradition [14,39]. Interestingly, IBM’s recent artifact-centric approach takes a similar indirect route, by first considering the objects that are manipulated in a process before the focus moves to the actual process design [6].

At the same time, it must be acknowledged that the translation of the product data model to a favorable workflow is a critical step. In the first applications of

PBWD this derivation was done manually [34, p.256-273]. Since this is time-consuming and error-prone, we have been developing IT tools to support the administration of a product data model, as well as algorithms that automatically generate workflow designs on the basis of a product data model [44]. Still, business users find it difficult to consider and compare all the options that are available for the final workflow design. After all, in general, there are many of such options. Some options may work well for some cases and not for others. Moreover, many potential workflow variants may only differ in subtle ways.

This paper presents an entirely new outlook on the use of the product data model. Instead of aiming at the derivation of a workflow design that is in general the best possible way to generate an informational product, the product data model *itself* is proposed as the vehicle to steer a workflow's execution. In other words, the need to translate a product data model into a workflow design disappears. Instead, a business user determines on a case-by-case basis the best possible way to create an informational product in accordance with the relevant product data model. This approach addresses the difficulty for business users to compare many alternative workflow designs, while it still relies on the product data model with its attractive properties. In addition, this approach allows for a highly dynamic and case-specific execution of workflows, as will be illustrated in the remainder of this paper.

The proposed approach builds on two pillars. On the one hand, we exploit the wide industrial proliferation of “process-aware” or “process-oriented” information systems [13,24]. We will assume the existence of this type of system to support the proposed approach and, along the way, show the feasibility of this idea in the form of a prototype workflow system. On the other hand, the offered solution rests on the idea that it is easier for a business user to determine what is the best possible action in the context of processing a *single* case than to do that for the *general* case. To guide the business user in this respect, we present two alternative approaches to provide her with recommendations: one that is optimal in relation to a dominant performance criterion but rather computing-intensive, and another that is lightweight and based on heuristics.

Our contribution can be summarized as follows. We present a rigorous approach for business process improvement, which addresses the need for guidance in this respect from practice. The innovative aspect is that we do not aim at the design of an underlying, generic process but instead provide a business user with direct support for delivering individual informational products. To do so, we build on the successful notions from Product-Based Workflow Design (PBWD), in particular the product data model, to arrive at a method that we coin “Product-Based Workflow Support” (PBWS). While the product data model specifies the elements to assemble a particular product, a process-aware information system recommends a business user on how to use these to deliver

the product in the best possible way.

The structure of this paper is as follows. Section 2 contains background information, a running example and the motivation for Product-Based Workflow Support. Next, Sections 3 and 4 present the two main, alternative realizations of the envisioned support. A comparison of these different approaches is given in Section 5, followed by a description of a workflow system (based on ProM and DECLARE) to support the overall approach in Section 6. Finally, the paper ends with related work and conclusions.

2 Background and Motivation

This section provides information that is essential as background for the remainder of the paper. In particular, the product data model will be explained, which will be illustrated by an example. Also, we will motivate the idea of Product-Based Workflow Support.

2.1 Workflow Products

The product of a workflow process is an *informational* product, e.g. a decision on an insurance claim, the allocation of a subsidy, the approval of a loan. Based on the input data provided by the client or retrieved from other systems, the process constructs the end product step-by-step. In each step new information is produced based on the specific data present for the case.

In this paper, we use a financial workflow process as a running example. The workflow process deals with the calculation of the maximum amount of mortgage a bank is willing to loan to a client. The bank has three alternative ways to decide on the maximum mortgage. First of all, if the client has a negative registration in the central register for credits (e.g. the client has a history of non-payment), the bank may directly reject this person for a mortgage (leading to a maximum amount of zero). The central credit register keeps track of all loans a person has and helps providers of loans in their assessment of the creditworthiness of such a person.

Secondly, if the client has previously requested a mortgage offer and the term of validity of this offer is not expired yet, this may determine the amount of mortgage. Typically, the percentage of interest changes over time and a mortgage offer is valid for some months. In case the interest has increased since the previous offer, the valid offer might be better than a new one which is based on the higher interest percentage.

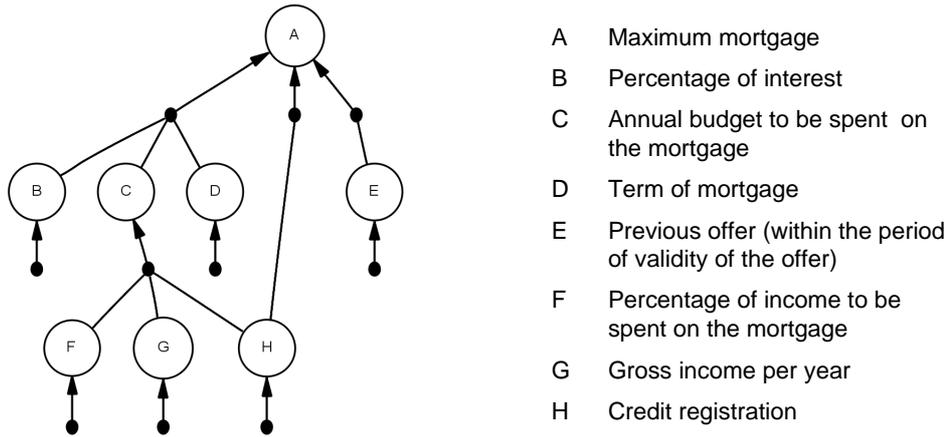


Figure 1. The PDM of the maximum mortgage calculation.

Finally, if the credit register shows a positive credit history, the bank needs more information on the client’s situation (e.g. gross income, type of mortgage) in order to decide on the maximum mortgage. To a certain extent, a bank defines its own internal rules and policies to decide on how much risk a mortgage applicant is to their business if they allow a mortgage. Therefore, each bank uses a percentage of the gross income of the client to calculate how much money the client is allowed to spend on the house. With this rule a bank ensures that the client can afford the cost of food and recurring expenses and that the probability is high that he will meet the monthly payment liabilities with the bank. However, this percentage is not fixed and can vary based on the bank’s current situation as well as the client’s.

In this example the maximum amount of mortgage is the end product of the workflow process, this is the piece of information that is “produced”. A graphical representation of the structure of this workflow product is given in Figure 1 and is explained below.

2.2 The Product Data Model

The structure of a workflow product is modeled by a *Product Data Model (PDM)*. The PDM is described by a tree-like structure similar to a Bill of Material [25]. The information that is processed in the workflow process is described by the so-called *data elements* of the PDM. For each specific case a data element can have a different *value*. Data elements are depicted as circles in the PDM, as can be seen in the example of Figure 1.

The actions that are taken on the data element values are called *operations* and are represented by arcs. Each operation can have zero or more *input*

data elements and produces exactly one *output data element*. The arcs are ‘knotted’ together when a value for all data elements is needed to execute the particular operation. Compare for instance the arcs from B , C , and D leading to A on the one hand, and the arc leading from E to A on the other hand in Figure 1. In the latter case, only one data element value is needed to determine the outcome of the process, while in the case of B , C , and D all three data element values are needed to produce A . An operation is said to be *executable* when a value for all of its input elements is available.

Several operations can have the *same* output element while having a different set of input elements. Such a situation represents *alternative ways* to produce a value for that output element. For example, a value for the end product A in Figure 1, can be determined in three alternative ways: (i) based on a value for E , (ii) based on a value for H , and (iii) based on values for B , C , and D .

The top element of the PDM, i.e. the end product, is called the *root* of the PDM. The *leaf elements* are the elements that are provided as inputs to the process. They are produced by operations with no input elements (e.g. the operations with output elements B , D , E , F , G , and H). The operations producing values for the leaf elements are denoted as *leaf operations* or input operations. An operation can be identified by a tuple consisting of the output element and a set of input elements, e.g. $(A, \{B, C, D\})$ for the operation producing a value for A based on data elements B , C , and D . Throughout this paper operations are also referred to by identifiers, such as $Op01$.

The construction of a PDM is a manual task. However, the information needed to construct a PDM can be obtained from e.g. rules and regulations, work instructions, forms, textual product descriptions, jurisprudence, information systems and the knowledge on the workflow product that is present with stakeholders.

Figure 1 expresses the PDM for the mortgage example. It shows that the maximum mortgage (element A in Figure 1) is dependent either on a previous mortgage offer (E), or on the registration in the central credit register (H), or on the combination of the percentage of interest (B), the annual budget to be spent on the mortgage (C), and the term of the mortgage (D). The annual budget (C) is determined from the gross income of the client per year (G), the credit registration (H), and the percentage of the income the client is allowed to spend on paying the mortgage (F).

2.3 Motivation

According to the PBWD method, the PDM is used as the basis for designing a process that on the one hand respects the dependencies between the various

data elements and on the other provides an attractive “walk through” along the various operations that need to take place [35]. It is important to stress that with PBWD such a process is proposed as a *generally* attractive way to deal with all thinkable instances of the product that the PDM relates to.

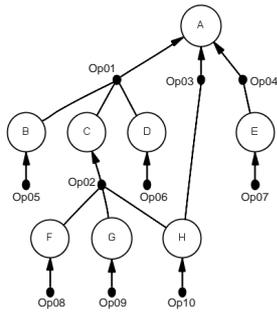
Both for the derivation of a general process (in accordance with PBWD) and for Product-Based Workflow Support (as proposed in this paper) the availability of various alternatives to achieve a particular outcome is exploited. Operations that create new values for data elements can take on different forms, e.g. an automatic calculation, an assessment by a human, or a rule-based decision, but all will consume time and money. In general, an operation can have a number of attributes associated to it that describe the characteristics of the operation in more detail:

- *Execution cost*: the cost associated with executing the operation (given by a probability distribution and its parameters).
- *Processing time*: the time that is needed to complete the operation (given by a probability distribution and its parameters).
- *Execution conditions*: conditions on the value of the input data elements that restrict the execution of the operation. If the condition is not satisfied, the operation is not executable even if a value for all of its input data elements is available.
- *Failure probability*: the probability that the operation is not performed successfully, i.e. the probability that the output data element is not produced.
- *Resource class*: the resource class or role that is required to perform the operation.

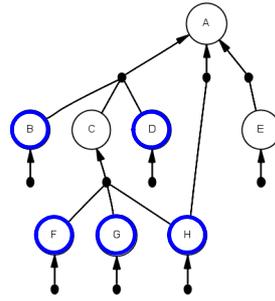
If we again consider our example, it could occur, for instance, that the cost of executing alternative operations differs. The involved bank has to pay for receiving a copy of the client’s registration in the credit register. Thus, the cost for this operation are higher than the cost for e.g. determining the gross income, since the information on income is provided by the client. The same holds for the processing time of these operations: Since an external party is consulted, it may take more time to retrieve the credit registration than to ask the client for his gross income.

To illustrate the notion of Product-Based Workflow Support, we will present with Figure 2 a step-by-step execution of the operations that occur in a PDM¹. Suppose that the values for the leaf elements B , D , F , G , and H (i.e. the interest percentage, term of mortgage, percentage of income to be paid in rent) are available at the start of the process *for one particular case*

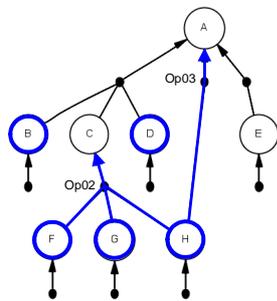
¹ For reasons of simplicity we abstract here from the execution conditions on the operations. Moreover, in Figure 2 we assume that all leaf operations have already executed in the initial state, i.e. a value is available for the output data elements of all successfully executed leaf operations.



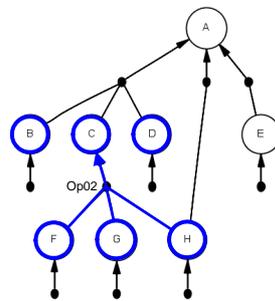
(a) The PDM for the mortgage example.



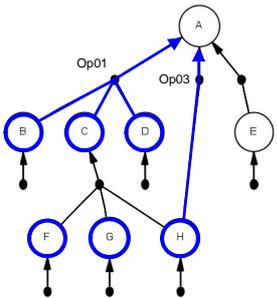
(b) The values for some of the leaf data elements (B, D, F, G, H) are available (indicated by bold circles). Throughout this paper, we refer to this situation as the *initial state*.



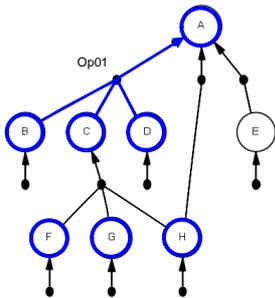
(c) Executable operations in the first step: $Op02$ and $Op03$.



(d) The value for data element C is produced by operation $Op02$.



(e) Executable operations in step two: $Op01$ and $Op03$.



(f) The value for the end product (A) is determined by executing operation $Op01$.

Figure 2. The step-by-step execution of the PDM for the mortgage example. Bold circles represent available data element values for the case under consideration; bold arrows indicate executable operations.

Table 1
Operations and their attributes for the mortgage example.

	Output	Input	Cost	Time	Prob.	Conditions
Op01	A	B, C, D	5.0	1.0	0.05	-
Op02	C	F, G, H	5.0	4.0	0.05	-
Op03	A	H	9.0	3.0	0.05	value(H) = 'negative'
Op04	A	E	2.0	2.0	0.00	-
Op05	B	-	0.0	0.0	0.00	-
Op06	D	-	0.0	0.0	0.00	-
Op07	E	-	1.0	1.0	0.50	-
Op08	F	-	0.0	0.0	0.00	-
Op09	G	-	0.0	2.0	0.00	-
Op10	H	-	3.0	10.0	0.15	-

(see Figure 2(b))². The operations that are now enabled for execution for this specific case are *Op02* and *Op03*, since a value for all of their input elements is available (Figure 2(c)). Operation *Op01* is not executable because a value for data element *C* is not available yet and *Op04* is not executable since there is no value for *E* present. Now, we have to choose between the two executable operations (*Op02*, *Op03*). Suppose we select *Op02*. Then, a value for data element *C* is produced (Figure 2(d)). The executable operations are calculated again, i.e. *Op01* and *Op03*, and one of these two operations is to be selected next. Suppose we select *Op01* (Figure 2(f)). Then, the value for the end product *A* is determined and the process ends.

The example illustrates that more than one operation may be executable at a certain point in time. For example, in the first step of this example we could have chosen for *Op03* instead of *Op02*. This would have led to the end product (*A*) immediately, but it generally also could have resulted in a different execution of the process with e.g. different total cost, throughput time, etc. For example, suppose the execution cost and processing time for the operations of the mortgage example are given in Table 1. If we focus on the cost of execution it seems best to choose *Op02* since the execution cost for *Op02* is lower than the execution cost for *Op03* (5.0 vs 9.0). However, if we consider the processing times for the operations, the selection of *Op03* as a next step would perhaps be a better decision (4.0 for *Op02* vs 3.0 for *Op03*). It

² Note that for this case not all leaf operations have been executed successfully. Because the execution of *Op07* has failed, there is no value for data element *E* available.

would be extremely valuable if at any point during the execution of operations these insights would be available.

The idea of Product-Based Workflow Support can now be described as follows. *An end user is supported on a case-by-case basis with recommendations on the most suitable way to carry out the available operations in a PDM.* A recommendation takes into account what type of performance is pursued for each case and how the various alternative executions differ from each other with respect to that performance criterion. Note that the general process model that is most suitable to deal with the average case, as prescribed by the PBWD method, becomes superfluous when using Product-Based Workflow Support. As an alternative, the end user is guided through the operations of the PDM in a way that is both flexible and performative.

In the remainder, we deal with the question how to determine proper recommendations, i.e. how to select in each possible situation from a set of executable operations the best operation to proceed with. It should be stressed that we define ‘best’ in the context of a single case, i.e. the performance goal (e.g. total cost, total processing time) of the case in isolation is optimized. Two main solution approaches for selecting candidate operations are presented in the next two sections.

3 Global Decision Strategies

The first solution approach aims at creating guaranteed optimal walkthroughs of the PDM by using global decision strategies. A global decision strategy takes into account the effect of the current decision on future decisions, i.e. the overall performance of the case. For example, in the first step of the mortgage case choosing *Op02* enables *Op01* in the next step and determines an alternative path to the path containing only *Op03* to produce the end product *A*. This effect is considered when determining a recommendation for the current situation. A global decision strategy takes the complete, alternative paths to the end product into account to optimize the overall performance of the case.

The technique we use to determine a walkthrough of the PDM using a global perspective is based on the theory of *Markov Decision Processes* (MDPs) [31,41]. An MDP extends the notion of a Markov chain with decisions. Markov chains are mathematical models of stochastic processes, i.e. systems that evolve over time [41]. They are widely used in e.g. operations research, biology, and computer science. In an MDP several decisions can be taken in each state of the Markov chain. The state transitions in the Markov chain are given by the transition matrix and are dependent on these decisions. Each decision has associated costs, which depend on the state of the system. The

goal of an MDP is to find a *strategy* that specifies which decision to take in each state, so as to minimize the overall cost [41]. Note that an important characteristic of a Markov chain is that it is *memoryless*, i.e. the probability of being in a certain state at time $n + 1$, is only dependent on the state at time n and not on earlier states ($< n$).

The execution of a PDM can be described as a memoryless stochastic process with decisions. We now present the formulation of an MDP based on the PDM to elaborate on this mapping.

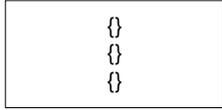
3.1 Formulation of the Markov Decision Process

An MDP is defined by a number of components: (i) the state space, (ii) the time space, (iii) the decision space, (iv) the transition function, and (v) the cost functions. To show that the decision problem in the execution of a PDM can be translated into an MDP, we describe for each component of the MDP the corresponding part in a PDM. A simple example, based on the mortgage process, is used to clarify this mapping.

State Space - The state space (S) describes the states the process can be in. The states of the execution of a PDM can be described by the operations that have been executed (either successfully or unsuccessfully) together with the data elements for which a value is available. A state in the state space is therefore represented by a tuple consisting of three sets: (i) the successfully executed operations, (ii) the unsuccessfully executed operations, and (iii) the data elements for which a value is available³. In Figure 3(b), for example, the state in the mortgage example in which a value for each of the leaf elements B , D , F , G , and H is available (cf. Figure 2(b)) is denoted by $(\{Op05, Op06, Op08, Op09, Op10\}, \{Op07\}, \{B, D, F, G, H\})$. Thus, $S \subseteq \mathcal{P}(O) \times \mathcal{P}(O) \times \mathcal{P}(D)$. The state space of the MDP is finite, since the number of operations is finite and we assume that each operation can only be executed once, either successfully or unsuccessfully.

Time Space - The time space describes the time points at which a decision is taken and at which a state transition occurs (also called *decision epochs*). The times used in our MDP are discrete and can be represented by the number of executed operations ($T = \{0, 1, 2, 3, \dots, |O|\}$), i.e., time is indicated by $t \in T$ indicating the number of executed operations. The times are not necessarily equidistant, but since there is a finite number of operations in a product data

³ Note that the set of available data elements in a state actually gives redundant information, since the set of available data elements can also be determined based on the set of successfully executed operations. However, this information is added for reasons of clarity and understandability of the examples.



(a) The state in which no data element values are available yet, cf. the initial state shown in Figure 2(b).



(b) The state in which the values for data elements B , D , F , G , and H are available, cf. Figure 2(b).

Figure 3. A state in the state space is described by three sets: (i) the operations that have been executed successfully so far, (ii) the operations that have been executed unsuccessfully, and (iii) the data elements for which a value is available.

model and we assume no concurrency, the time space is bounded by this number, i.e. the MDP problem has a finite time horizon.

Decision Space - In each state a number of decisions can be made. For the execution of a PDM these decisions are described by the set of operations that are executable in the current state of execution (i.e. those operations of which the input elements are available and that have not yet been executed)⁴. Moreover, if there are no executable operations for a certain state there is only one decision possible, i.e. to stop. Thus, the decision space A is equal to the set of operations plus the decision to stop, i.e. $A = (O \cup \{stop\})$. Furthermore, the decision space in a particular state A_i is a subset of $(O \cup \{stop\})$. Note that the decision space in a certain state is time-independent, i.e. given a state, at any point in time, the same decisions can be made.

Transition Probabilities - The transition probabilities are given by a matrix P that describes the probabilities that the system moves from the current state to any of the other states in the system. These transition probabilities are dependent on the decision that was made. For our application, a decision a in state i can lead to two new states: j_1 (for a successful execution of the operation given in decision a) and j_2 (for a failed execution of the operation given in decision a), each with their own probabilities. The two transition probabilities under decision a always add up to 1. For example, recall the execution of the mortgage example in Figure 2. If we start in the state with available values for data elements B , D , F , G , and H , there are two decisions which can be taken, i.e. $\{Op02, Op03\} \subseteq O$. Each of these two operations can either fail or be successfully executed. This leads to four new states (Figure 4). The transition probabilities correspond to the probabilities of failure or successful execution of an operation (See Table 1):

⁴ Throughout this paper we abstract from concrete values of data elements and conditions based on these values. Incorporating data is possible if the number of possible values is small. If many values are possible, the state space will be too large to allow for any form of analysis.

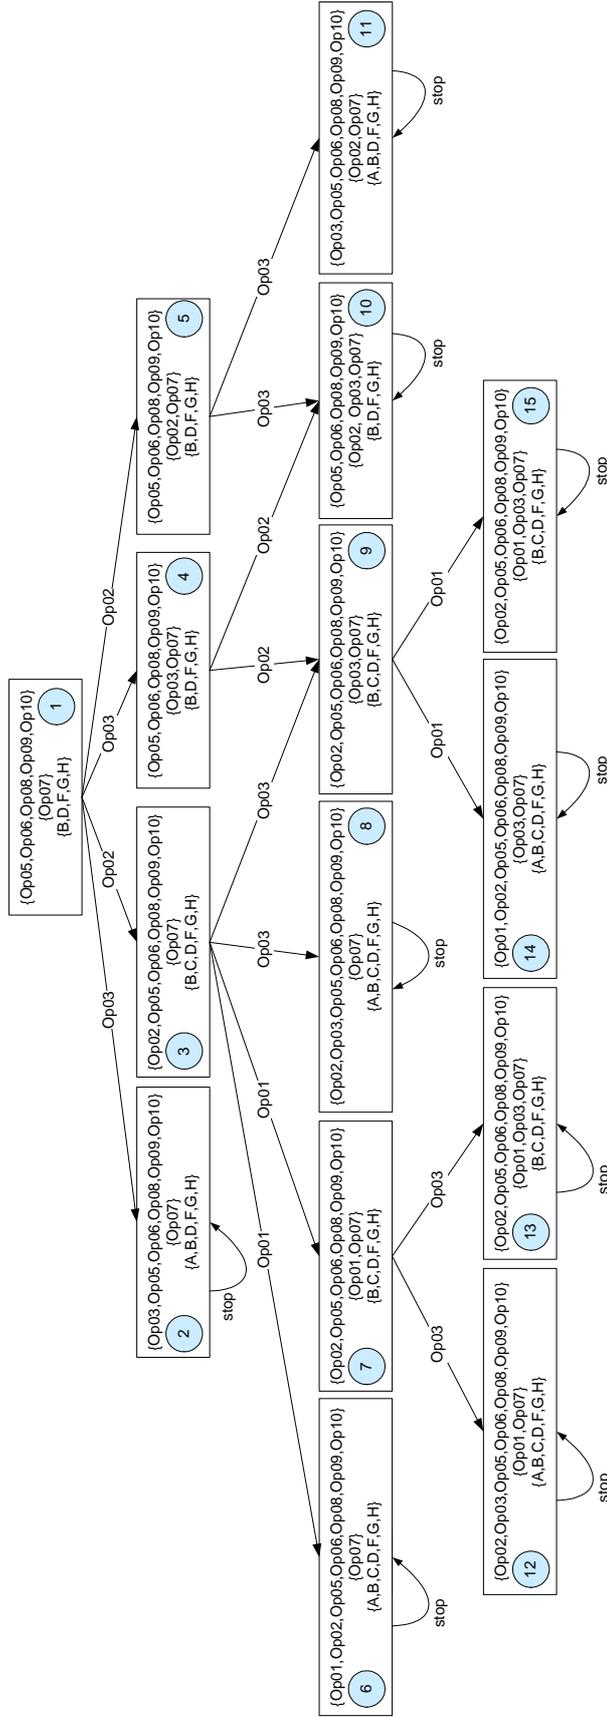


Figure 4. The state space of the mortgage example. Note that there are eight end states: {2, 6, 8, 10, 11, 12, 13, 14, 15}.

$$\begin{array}{ll}
p_{1,3}(Op02) = 0.95 & p_{1,2}(Op03) = 0.95 \\
p_{1,5}(Op02) = 0.05 & p_{1,4}(Op03) = 0.05
\end{array}$$

Thus, the probability of moving from state 1 to state 3 under decision $Op02$ is 0.95.

Immediate and Final Cost - The immediate cost of a transition from state i to state j under decision a are the cost of executing operation a . In our model there are no costs associated to residing in a state, i.e., costs are associated to decisions. Depending on the performance objective of the process, e.g. minimization of cost or processing time, the immediate cost are defined as the cost or the processing time of the operation. Suppose we want to minimize the execution cost of the process, then the immediate cost in state 1 are (See Table 1):

$$c^{Op02}(1) = 5.0 \qquad c^{Op03}(1) = 9.0$$

The final cost incurred at time $|O|$, when no decisions can be made anymore, in state i are zero, i.e. $q(i) = 0$.

By using the mapping presented above, the process of executing a PDM can be translated to a time-homogenous MDP with a finite horizon. Note that this mapping of a PDM to an MDP results in a Markov chain with some special properties. First of all, the state space is *finite*, i.e. there is a finite number of states. Secondly, each operation can be executed at most once. Thus, each decision can also be taken at most once during any execution. Moreover, at each point in time an operation is selected until there are no executable operations anymore. The only decision that can be made then is to stop and the system stays in the same state. These end states are *absorbing* states. The Markov chain is a *transient* chain since each closed class has exactly one absorbing state. Besides the absorbing end states there are *no cycles* in the Markov chain, i.e. it is not possible to return to a state that was previously visited.

We have used a small part of the mortgage example to illustrate the mapping from a PDM to an MDP (see Figure 4). Below, this example is used to demonstrate how a stationary Markov decision rule (f) is determined for the complete mortgage example.

3.2 The Mortgage Example as an MDP

In this section, the MDP problem for the mortgage example is explained as an illustration. The initial situation (see Figure 2(b)) in which there is already a

state	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	0.0	0.95	0.95	0.05	0.05	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
2	0.0	1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
3	0.0	0.0	0.0	0.0	0.0	0.95	0.05	0.95	0.05	0.0	0.0	0.0	0.0	0.0	0.0
4	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.95	0.05	0.0	0.0	0.0	0.0	0.0
5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.95	0.0	0.0	0.0	0.0
6	0.0	0.0	0.0	0.0	0.0	1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.95	0.05	0.0	0.0
8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
9	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.95	0.05
10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0	0.0	0.0	0.0	0.0
11	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0	0.0	0.0	0.0
12	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0	0.0	0.0
13	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0	0.0
14	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0
15	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.0

Table 2. Transition probabilities for the state space of the mortgage example in Figure 4. The transition probabilities correspond to the failure probabilities in Table 1.

value for data elements B , D , F , G , and H is used again because in this case the decision space is rather small. Therefore, the state space for this MDP problem stays small and readable. The state space is depicted in Figure 4 and describes all possible execution steps from the initial state. The decision space per state can also be derived from the figure, e.g. in state 3 two alternatives exist: $Op01$ and $Op03$. Because the execution of both operations can either be successful or unsuccessful, state 3 has four outgoing arcs to four new states. If operation $Op01$ is executed successfully, the system moves from state 3 to state 6. Confronted with an unsuccessful execution of $Op01$, the system moves to state 7. Similarly, upon execution of $Op03$ the system moves to state 8 or 9. The probabilities for these transitions are summarized in Table 2. For this example we focus on minimizing the total cost for a case. The cost for executing an operation can be found in Table 1.

The optimal decision strategy for this MDP can be calculated using a value iteration algorithm [41]. The algorithm is defined as follows:

Definition 3.1 Let $V_0(i) := q(i)$, and $V_t(i), t = 1, 2, \dots$ be recursively given by

$$V_t(i) := \min_{a \in A_i} [c^a(i) + \sum_{j \in S} p_{i,j}(a) \cdot V_{t-1}(j)].$$

Then, any decision rule f_t determined by

$$f_t(i) = \arg \min_{a \in A_i} [c^a(i) + \sum_{j \in S} p_{i,j}(a) \cdot V_{t-1}(j)]$$

gives the optimal decision for each state in the state space and achieves the minimal cost over t periods to the time horizon.

The value iteration algorithm uses backward induction to step-by-step derive the decision rules. First, the cost at the time horizon (see Figure 5) are determined. The time horizon is the point in time at which the process ends and no decisions can be made anymore. Then, a decision rule (f_1) is calculated for the last decision step in the process, i.e. the situation in which one decision can be made ($t = 1$). This calculation is based on the expected cost at $t = 0$. Subsequently, the decision rules for $t = 2, 3, \dots$ can be determined. Below, the first steps of this algorithm are elaborated for our MDP. $V_t(i)$ denotes the total expected cost at decision epoch t and state i . The cost at decision epoch $t = 0$ are equal to the final cost, i.e.

$$V_0(i) = 0, \forall i \in S.$$

Next, the values for $t = 1$ can be calculated based on all $V_0(i)$ values. For instance, two decisions can be made in state 3: $Op01$ and $Op03$. The expected cost for choosing $Op01$ are dependent on the immediate cost for choosing $Op01$ in state 3, and the expected cost and transition probabilities for the states the

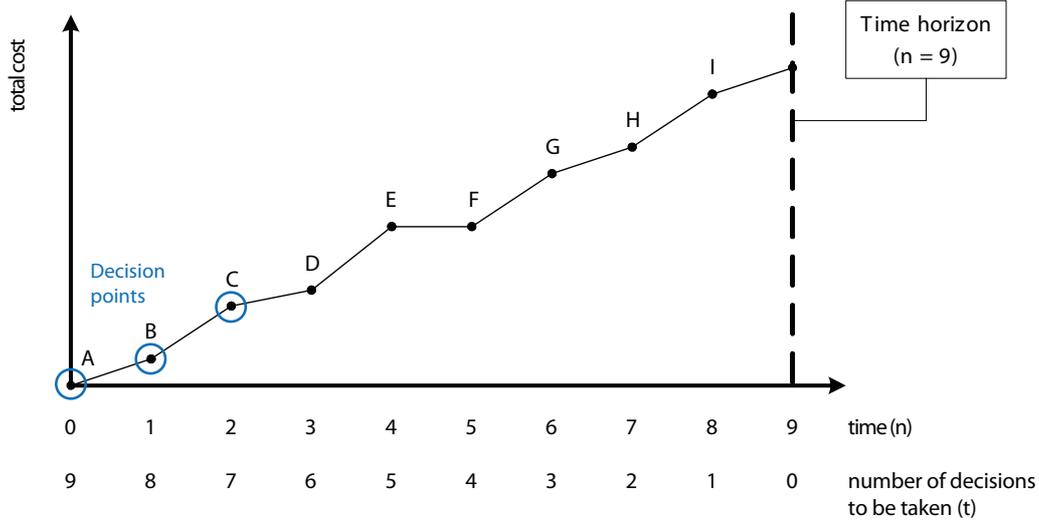


Figure 5. This figure describes a time line of an MDP with a finite time horizon. The process ends at time $n = 9$. This point is either represented by the time, i.e. $n = 9$, or by the number of decision epochs to the end of the process, i.e. $t = 0$.

process can move to under decision $Op01$:

$$c^{Op01} + p_{3,6}(Op01) \cdot V_0(6) + p_{3,7}(Op01) \cdot V_0(7) = 5.0 + 0.95 \cdot 0.0 + 0.05 \cdot 0.0 = 5.0.$$

Similarly, the expected cost in state 3 for decision $Op03$ can be calculated:

$$c^{Op03} + p_{1,2}(Op03) \cdot V_0(2) + p_{1,4}(Op03) \cdot V_0(4) = 9.0 + 0.95 \cdot 0.0 + 0.05 \cdot 0.0 = 9.0.$$

Since decision $Op01$ has the lowest expected total cost it is the best decision. Thus, if the process is in state 3 and there is only one decision to be taken until the end of the process, the best decision is $Op01$, i.e. $f_1(3) = Op01$, with expected total cost of 5.0, i.e. $V_1(3) = 5.0$. The same way the expected cost and best decisions for all other states at time $t = 1$ can be calculated. Tables 3 and 4 show that if there is only one time period left, it is best to choose $Op02$ in state 1 since $V_1(1)$ has the minimal value of 5.0 (vs. 9.0) for its argument $Op02$. Also, $Op02$ should be chosen in state 4 ($V_1(4) = 5.0$ for $Op02$), etc. In states 2, 6, 8, 10, 11, 12, 13, 14, and 15 the only decision that can be made is to stop since these states are end states in the state space. Thus, the decision rule f_1 (for time $t = 1$) is:

$$f_1 = \{(1, Op02), (2, stop), (3, Op01), (4, Op02), (5, Op03), (6, stop), (7, Op03), (8, stop), (9, Op01), (10, stop), (11, stop), (12, stop), (13, stop), (14, stop), (15, stop)\}.$$

In a similar way, the values and decision rules for all time points can be calculated. The details can be found in tables 3 and 4. The tables show that the decision rules become stationary for sufficiently large t (i.e. $t \geq 3$ for this case), since the Markov chain does not contain any cycles other than the absorbing end states. This results in the following global decision strategy:

$$f = \{(1, Op03), (2, stop), (3, Op01), (4, Op02), (5, Op03), (6, stop), (7, Op03), (8, stop), (9, Op01), (10, stop), (11, stop), (12, stop), (13, stop), (14, stop), (15, stop)\}.$$

Note that the best decision for the first state (describing the initial situation) has changed from $Op02$ at $t = 1$ to $Op03$ at $t \geq 2$. Thus, by looking at the complete state space and decision space, the algorithm calculates a global strategy that gives the best decision for each state in the MDP.

3.3 Computability of the decision strategy

In the previous section we have elaborated on the computation of a decision strategy for a small example with a small state space. As we mentioned before, the state space of our MDP problem is finite. Nevertheless, the state space can become extremely large. For instance, if we use an initial situation in the mortgage example in which none of the operations have been executed yet and no value is available for any of the data elements, i.e. the initial state with $(\{\}, \{\}, \{\})$ of Figure 6(a), the state space contains 2218 states and, as such, is much larger than in our example of Figure 4 with 15 states. This is because the six leaf operations can be executed in an arbitrary order. All possible combinations are then explicitly represented in the state space.

Generally, an upper bound for the size of the state space can be given in terms of the number of operations of the PDM. Each operation can be either (i) successfully executed, (ii) unsuccessfully executed (failed), or (iii) not yet executed. Thus, there are $3^{|O|}$ possible combinations of operations, which is an upper bound for the size of the state space. However, most state spaces will be smaller since not all combinations of executed operations are possible. For instance, the state space for the mortgage example will never contain the state $(\{Op02, Op07\}\{Op09\}\{C, G\})$, since $Op02$ can never be executed before both $Op08$ and $Op10$ are executed.

We have experimented calculating the state space for one of the industry cases in which PBWD was applied. The PDM of the UWV-case (described in [34]) contains 45 data elements and 50 operations, which is significantly larger than our running example but still reasonably small for a practical case. During

Table 3

The minimal cost values obtained by the value iteration algorithm, i.e. $V_t(i)$. Note that the values become stable, i.e. they do not change anymore, for $t \geq 3$.

t/i	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1	5.0	0.0	5.0	5.0	9.0	0.0	9.0	0.0	5.0	0.0	0.0	0.0	0.0	0.0	0.0
2	9.25	0.0	9.25	9.75	9.0	0.0	9.0	0.0	5.0	0.0	0.0	0.0	0.0	0.0	0.0
3	9.49	0.0	9.25	9.75	9.0	0.0	9.0	0.0	5.0	0.0	0.0	0.0	0.0	0.0	0.0
4	9.49	0.0	9.25	9.75	9.0	0.0	9.0	0.0	5.0	0.0	0.0	0.0	0.0	0.0	0.0
5	9.49	0.0	9.25	9.75	9.0	0.0	9.0	0.0	5.0	0.0	0.0	0.0	0.0	0.0	0.0
6	9.49	0.0	9.25	9.75	9.0	0.0	9.0	0.0	5.0	0.0	0.0	0.0	0.0	0.0	0.0
7	9.49	0.0	9.25	9.75	9.0	0.0	9.0	0.0	5.0	0.0	0.0	0.0	0.0	0.0	0.0
8	9.49	0.0	9.25	9.75	9.0	0.0	9.0	0.0	5.0	0.0	0.0	0.0	0.0	0.0	0.0
9	9.49	0.0	9.25	9.75	9.0	0.0	9.0	0.0	5.0	0.0	0.0	0.0	0.0	0.0	0.0

Table 4

The rows in this table denote the decision strategy (f_t) for each point in time (t).

t/i	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	Op02	-	Op01	Op02	Op03	-	Op03	-	Op01	-	-	-	-	-	-
2	Op03	-	Op01	Op02	Op03	-	Op03	-	Op01	-	-	-	-	-	-
3	Op03	-	Op01	Op02	Op03	-	Op03	-	Op01	-	-	-	-	-	-
4	Op03	-	Op01	Op02	Op03	-	Op03	-	Op01	-	-	-	-	-	-
5	Op03	-	Op01	Op02	Op03	-	Op03	-	Op01	-	-	-	-	-	-
6	Op03	-	Op01	Op02	Op03	-	Op03	-	Op01	-	-	-	-	-	-
7	Op03	-	Op01	Op02	Op03	-	Op03	-	Op01	-	-	-	-	-	-
8	Op03	-	Op01	Op02	Op03	-	Op03	-	Op01	-	-	-	-	-	-
9	Op03	-	Op01	Op02	Op03	-	Op03	-	Op01	-	-	-	-	-	-
10	Op03	-	Op01	Op02	Op03	-	Op03	-	Op01	-	-	-	-	-	-

our experiments we experienced that a normal PC⁵ was not able to process the complete state space of this case. Therefore, we tried to limit the amount of states in the state space based on a number of restrictions described in [44]. One of the main design objectives for the redesign project of the case was a minimum number of contacts with the client (see [34]). This requirement was realized in the design by requesting all input information from the client at once at the start of the case. Therefore, the restriction of using an initial state in which all leaf operations have been (successfully) executed is not unrealistic. Calculating the first 300.000 states of the state space with only this restriction took more than 60 days and we failed to gain the complete state

⁵ A normal PC at the time of writing this paper is e.g. an Intel Pentium 4 with a 3.8GHz processor and 3GB of RAM.

space. Furthermore, it also seemed realistic to assume that all operations are executed successfully, since the failure probabilities for the operations in the PDM are typically very low. With these two restrictions, the state space for the social benefits case was calculated in 2.5 days. The state space contained 66.581 states. Based on this partial state space a decision strategy can be computed. However, this decision strategy may not be optimal because of the assumptions that were made to limit the state space.

The experiments show that it is theoretically possible to find globally optimal recommendations for executing the PDM, but for real-life applications it may be intractable to find the optimal solution. Typically, PDMs may be even larger than the one for the UWV-case, e.g. containing hundreds of data elements. Therefore, the MDP may become unsolvable in practice. Note that the selection of the next operation needs to be done at real time and repeated after every step. Hence, a more pragmatical way of determining recommendations would be desirable. As a second solution approach, we now describe a number of *local decision strategies* that focus on the selection of the candidate operation for execution within the set of executable operations.

4 Local Decision Strategies

The issue of unsolvable MDPs due to large state spaces also exists in other application areas of MDP theory. In the field of production planning, simpler heuristics have been developed that perform well and approximate the optimal solution of the decision problem. These heuristics aim for a local instead of a global optimization. They do not consider the effects of a decision on the next steps in the execution, but merely look at the currently available set of possible decisions and take the best one out of it. For instance, in the initial situation of Figure 6(b), two operations are executable. When deciding on the next operation to be executed only those two operations are considered without looking at the further, possible effects of the choice for one of the two. Based on the characteristics of the two operations the best performing one with respect to the selected criterion is selected.

Based on production planning heuristics we have developed some *local decision strategies* for the direct execution of a PDM. We have drawn inspiration from *sequencing and scheduling rules* in the field of logistics and production planning [26,40]. Short-range production planning is a method to decide, a priori, which resource is going to perform which jobs in which order for production systems with a shop structure (i.e. flow shop, job shop, or open shop) [40]. Production planning problems are usually too large to be solved mathematically and researchers therefore have developed pragmatic ways (e.g. heuristics) that approximate the desired solution to a scheduling problem. Many different

strategies or rules exist to schedule jobs that have to be performed by a machine [26,40]. Well-known strategies are, for instance, First Come First Served (FCFS) or Earliest Due Date (EDD) [40]. The following decision strategies can be applied to our selection problem of finding the next operation in the execution of a PDM:

- *Random* - The operation is randomly selected from the set of executable operations (cf. [26]).
- *Lowest cost* - The operation with the lowest cost is selected.
- *Shortest processing time* - The operation with the shortest duration is chosen (cf. SPT defined in [26]).
- *Lowest failure probability* - The operation with the lowest probability of not being performed successfully is chosen.
- *Shortest distance to root element* - The operation with the shortest distance to the root element (measured in the total number of operations) is selected. The distance of an operation to the root element is the ‘shortest path’ from this operation to an operation that produces the root element. The distance to the root element can be measured as the total number of operations to the root element (cf. FOPNR defined in [26]).
- *Shortest remaining processing time* - The operation with the shortest remaining processing time is chosen. The shortest remaining processing time is another form of the distance to the root element. In this case the processing times of the operations on the path to the root element are added up (cf. SR defined in [26]).

If, for example, a lowest cost strategy was used for decision making on the initial situation of the mortgage example in Figure 6(b) then *Op02* would have been selected since the cost for *Op02* are lower than the cost for *Op03* (i.e. 5.0 vs. 9.0). If a shortest processing time strategy was used *Op03* would have been selected and with a shortest distance to root element strategy also *Op03* would have been chosen.

The above strategies present ways to select the next step based on a single strategy, i.e. only one performance goal is considered. A more advanced way to use these simple selection strategies is by combining them. With a single strategy there may be several candidates available for selection with the same value for the performance goal. To distinguish between the operations that all have an optimal value based on the first strategy, one can use another strategy to complement the first one. For instance, if the lowest cost strategy leads to three operations with the same minimal execution cost, another strategy (e.g. lowest failure probability) may be used to rank these three operations and find the best one among the three. Additionally, weighted criteria can be used. For instance, if the processing time and the execution cost are equally important for the selection of an operation, the selection strategy can be based on an equal weight (0.5 vs 0.5) for both performance criteria.

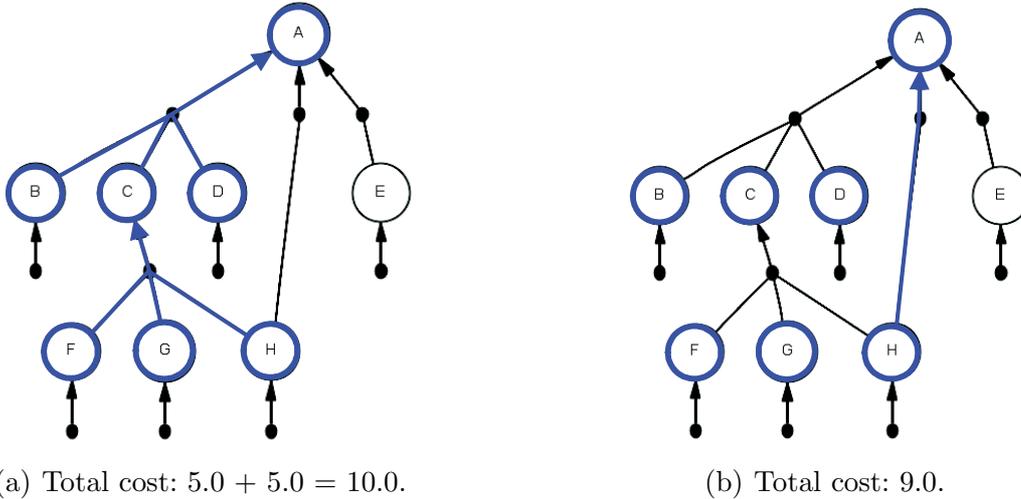


Figure 6. Two different execution paths leading to the end product of the PDM. The cost of executing an operation are presented in Table 1

Using the strategies presented above, the selection of the next candidate for execution is optimized locally (i.e. within the set of currently executable operations); the effect of the selected operation on future steps is not taken into account⁶. Such a local optimization may lead to a less desirable situation if we consider the overall performance of the case. Consider, for instance, the execution steps for the mortgage example with respect to the total execution cost. The first execution sequence contains *Op02* followed by *Op01*. The total cost of this execution are: $5.0 + 5.0 = 10.0$. The total cost of the alternative execution path containing only *Op03* are 9.0. So, the cost of the second execution path are lower. Selecting the best candidate based on the lowest cost selection strategy in this case does not lead to the best overall decision for the case considering the total cost of execution (see Figure 6). However, many of these pragmatic heuristics are proven to reach a near-optimal solution [26,32]. To assess our approaches, a comparison between some local strategies and the global decision strategy is made in the next section.

5 Evaluation

In this section, we evaluate the two proposed approaches based on the mortgage example. In order to assess the performance of the global and local decision strategies, we compare their results using two performance criteria. We

⁶ Note that this is comparable to a situation in the MDP in which there is only one decision left to the end of the process, i.e. $t = 1$

consider a situation in which the goal is to minimize cost as well as a situation in which the processing time is minimized.

For the global decision strategy, an expected value of each of these performance criteria can be calculated based on the MDP (see also Definition 3.1). Since the mortgage example is small enough we are able to compute a complete state space⁷ and decision strategy for the MDP derived from the PDM. To do so, we used the prototype workflow system presented in Section 6.

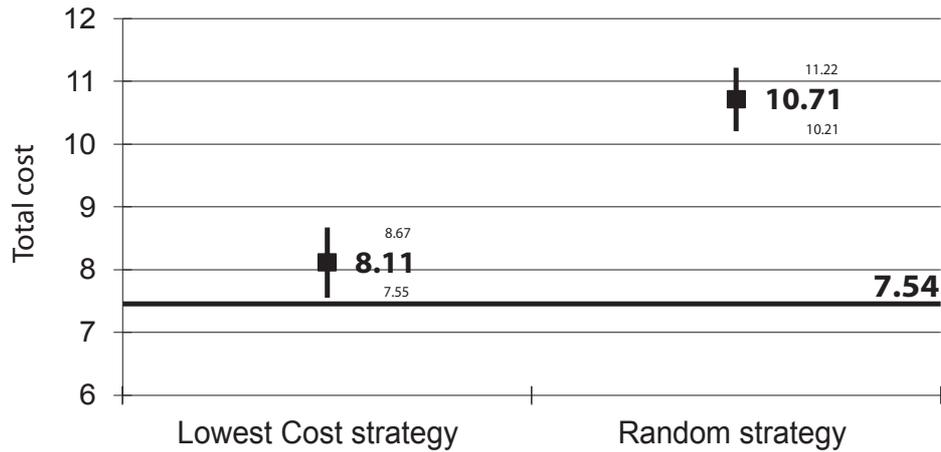
The performance of the local decision strategies is measured through simulation. In this simulation study, we have executed several simulation runs per local decision strategy. A simulation run corresponds to one complete execution of the PDM of the mortgage example. For each execution the total execution cost and total processing time were determined. Per strategy we collected 300 samples (i.e. $n = 300$) and constructed 90%-confidence intervals [23]. In the next section, the simulation model is explained first. Next, the simulation results are compared to the results of the global selection strategies.

5.1 Simulation model

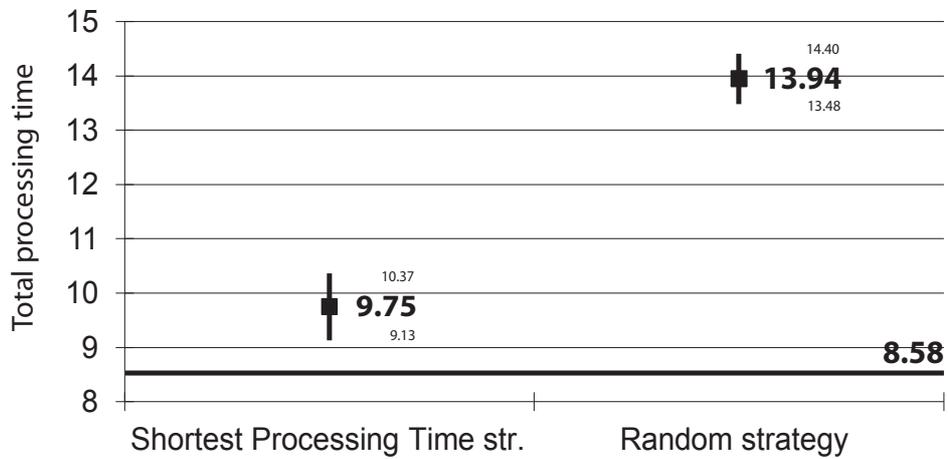
The simulation model for executing a PDM describes the functional design of a tool supporting the direct execution of a PDM based on a local decision strategy. The simulation model is represented by a Colored Petri Net (CPN, [16]) and has developed using CPN Tools [10]. The model takes care of a *correct* step-by-step execution of the PDM as was illustrated in Figure 2. This means that the functional design ensures that only executable operations are executed, that no duplicate data element values are produced, and that only one operation at a time is performed. It does not deal with performance and the selection of the best operation for the next step. The structure of the model is generic since the model does not need to be changed for a different PDM or a different selection strategy. The PDM is added to the model as a variable and the selection strategy as a function. This section describes the CPN model on a high level. For more details we refer to [44,45].

The main level of the CPN model is depicted in Figure 7. It contains eight places and three transitions. All three transitions are sub processes in which the actions to be taken are described in detail. For a detailed description of these subprocesses we refer to [44,45]. The main execution stream in the model is indicated by thick lines: First the executable operations are determined. Then one of the executable operations is selected based on the local decision strategy that was chosen. The selected operation is executed. If the execution

⁷ We take the initial situation in which no input data element values are available, i.e. $(\{\}, \{\}, \{\})$.



(a) total cost



(b) total processing time

Figure 8. Some results of the simulation study.

mean of 8.11. The 90%-confidence interval has a lower bound 7.55 and an upper bound of 8.67. The total cost under a *random* strategy has a mean of 10.71 (between 10.21 and 11.22 according to the 90%-confidence interval). Since these confidence intervals do not overlap, we may conclude that a lowest cost strategy leads to lower total cost than a random selection strategy. Besides the comparison of the local strategies we can also compare the outcome of the local strategies to the optimal solution resulting from the MDP. The total expected cost for executing the PDM of the mortgage example are 7.54 (see horizontal line in Figure 8(a)). Figure 8 shows that a lowest cost strategy closely approaches the optimal solution when a minimal cost strategy is selected in the MDP.

The second diagram depicts the results for the processing time performance criterion. The confidence intervals for the total processing time under a shortest processing time strategy and under a random strategy are given. The total processing time under a shortest processing time strategy has a mean of 9.75, a lower limit of 9.13 and an upper limit of 10.37. The *random* strategy has a mean of 13.94, a lower limit of 13.48 and an upper limit of 14.40. Clearly, the shortest processing time strategy performs better than a random strategy to minimize the total processing time. The results of the local decision strategies can again be compared to the expected processing time achieved by the global decision strategy obtained by the MDP: 8.58 (see horizontal line in Figure 8(b)). The shortest processing time strategy approaches the optimal solution compared to a minimal processing time decision strategy in the MDP.

We conclude from this simulation study that (i) using a heuristic (e.g. a lowest cost strategy to minimize cost) leads to better results than random selection of operations, and (ii) the outcomes of the heuristics closely approximate the optimal solution calculated by an MDP.

6 System Support

In this section we describe the workflow system that we have developed for Product-Based Workflow Support. It allows for the direct execution of a PDM by guiding the user through all steps to the end product. The system presents *execution recommendations* to the user based on the selected decision strategy and the current state of a case. Section 6.1 explains the architecture of our workflow system. Next, it is explained how the global and local strategies are supported by this system.

6.1 Architecture

Our approach builds on ProM and DECLARE.⁸ ProM [2,12] was initiated as a framework for *process mining*, but in recent years it has evolved into a broad and powerful *process analysis* tool supporting all kinds of analyses related to business processes [12]. ProM as a plug-able architecture. ProM 5.2 has 286 plug-ins, each realizing a particular type of functionality, e.g., the α miner plug-in is able to discover a process model from event logs and the conformance checker plug-in is able to measure the quality of a model in relation to an event log. DECLARE is a workflow management system [27,28,29]. DE-

⁸ ProM can be downloaded from www.processmining.org. DECLARE can be downloaded from <http://www.win.tue.nl/declare/>.

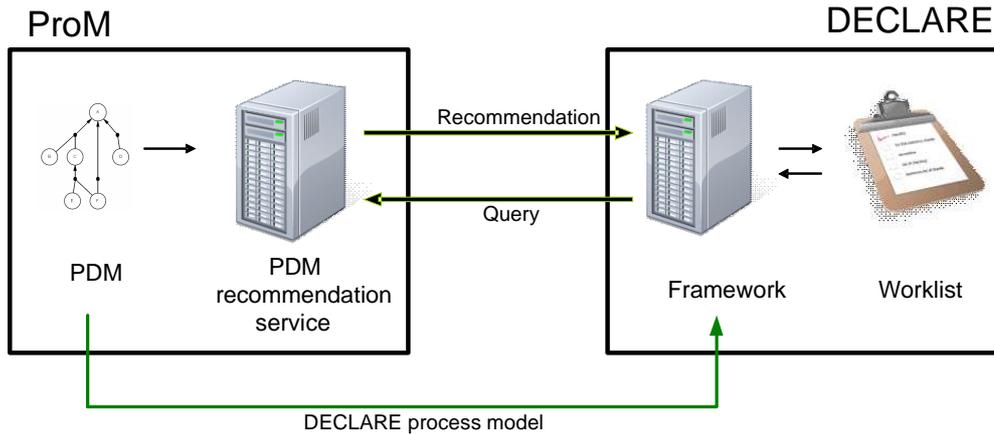


Figure 9. An overview of the architecture of the workflow system based on ProM and DECLARE.

CLARE is based on a *declarative* approach to business process modeling and execution and therefore provides more flexibility than conventional workflow management systems do. While the mainstream workflow management systems use a procedural approach where every activity needs to be triggered, DECLARE starts from the viewpoint that “anything is allowed unless explicitly forbidden”.

To realize a workflow system for Product-Based Workflow Support, we use DECLARE as-is and added several plug-ins to ProM. DECLARE is used to offer work to end users and could be replaced by another workflow system. However, we selected DECLARE because it is easy to specify workflow models that do not impose any control-flow constraints, i.e., from the viewpoint of the process any sequence of activities is possible (unless explicitly forbidden). Another reason for selecting DECLARE is the connection to ProM which allows to combine ProM’s analysis and reasoning capabilities with the flexibility of DECLARE. DECLARE does not know about data and dependencies between data elements, i.e., the PDM is only known to ProM. ProM uses a so-called *recommendation service*, the PDM, and knowledge about the history of cases, to inform DECLARE about recommended operations. Such a recommendation service (that is driven by a PDM and the availability of data) replaces the conventional control-flow oriented workflow engine. Two recommendations services have been realized using ProM’s plug-able architecture: one using the global MDP approach and one using local decision strategies.

An overview of the architecture of our system is shown in Figure 9. As a first step in the use of the PDM-based workflow system, a PDM is loaded into the ProM framework and exported to a DECLARE model. This DECLARE model solely contains the operations of the PDM as activities and their input and output data elements. Thus, no control flow is added to the DECLARE

model, i.e. the DECLARE activities are not related to each other in any way. Then, the ProM PDM recommendation service is started and a strategy is selected by the user. Actual cases can be handled after the DECLARE model has been loaded into the DECLARE framework and the DECLARE worklist has been started. The DECLARE framework communicates with the ProM recommendation service by sending *queries*. Such a query contains the current state of the case in terms of executed activities and available data element values. Based on this information, the PDM recommendation service calculates which operation is the best candidate to be executed for the case with respect to the selected decision strategy. The result of this calculation is sent back to the DECLARE framework as a *recommendation*. Finally, this recommendation is shown to the user in the worklist of DECLARE. The user may then follow up on the recommendation and execute the recommended operation. When the operation is executed a new query is sent to the ProM recommendation service.

Two different implementations of the PDM recommendation service have been developed. The first one is based on the global decision strategies introduced in Section 3. The second implementation uses the local decision strategies as described in Section 4 to generate recommendations for a query. Both implementations are described in more detail in the next sections.

6.2 Global Decision Strategies

In case a global decision strategy based on an MDP is chosen for generating recommendations, the user first has to calculate the state space and the optimal decision strategy in ProM (see Figure 10). Next, the recommendation service in ProM can be started. Also, the DECLARE framework and the DECLARE worklist are started. The user fills out the values for the leaf data elements. Suppose we take the same initial situation for the mortgage example as before (i.e. values for B , D , F , G , and H are available and E is not available), then a query containing the set of executable operations (i.e. $Op02$, $Op03$) and a set of available data elements and their values is sent to ProM. ProM determines a recommendation based on the calculated global decision strategy (i.e. $Op03$ under a minimize cost strategy) and sends it to DECLARE. In the DECLARE worklist, the recommendations are presented (see Figure 11). The first one is the recommended activity according to the calculated strategy. The others are second best options. Each recommendation also has a field called ‘rationale’ in which the total expected cost or processing time to finish the process is shown.

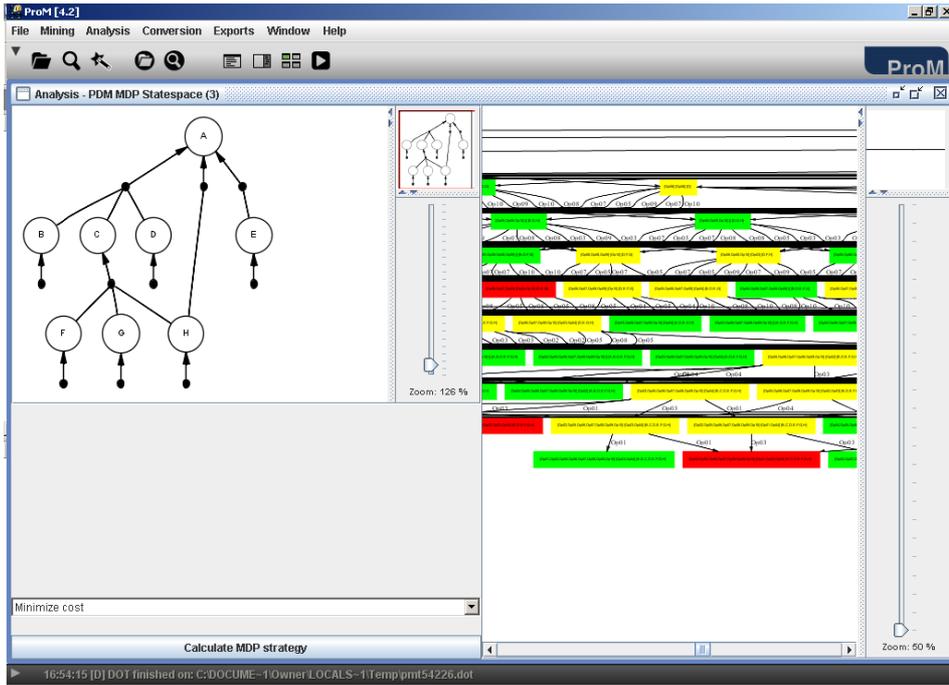


Figure 10. On the right hand side, this screenshot of ProM shows the complete state space that was computed for the mortgage example. In the lower left corner the user can select the performance objective to be optimized (e.g. minimize cost, or minimize processing time) via a drop-down menu. Next, the decision strategy is calculated.

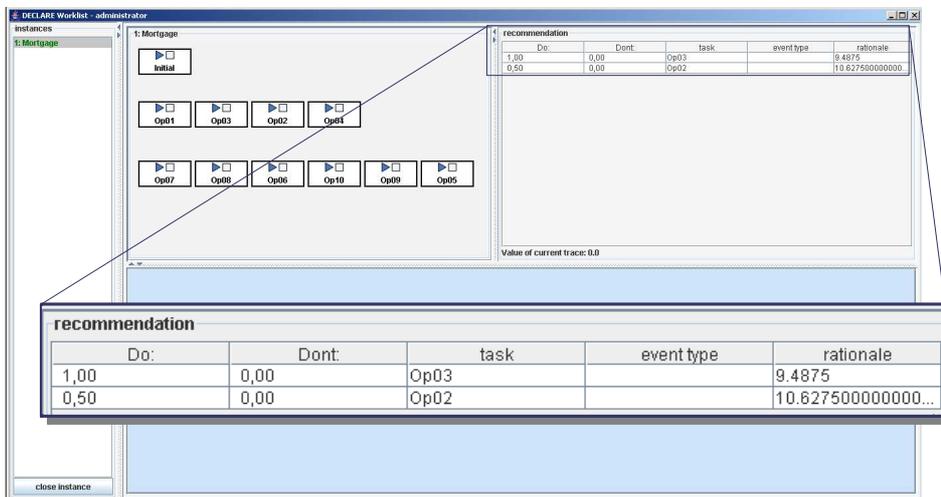


Figure 11. Two recommended activities are shown for the process which is in the state where all leaf operations have been executed and the values for data elements B , D , F , G , and H are available. Based on the global decision strategy calculated by using the MDP, $Op03$ is recommended as the next operation to be executed (cf. $V_n(3)$ and $f_n(3)$, for $n \geq 2$, in Table 3 and 4). Note that all executable operations are ranked with respect to the performance objective: the second best next step is $Op02$. Other operations are not executable.

6.3 Local Decision Strategies

In the PDM recommendation service based on local decision strategies, the user can immediately select one of the strategies introduced in Section 4 via a drop-down menu. Once the strategy has been selected, the recommendation service can be started. The execution of the process based on these recommendations is similar to the execution based on the global strategies described above. Given the state of the case, the recommendation service determines which operations are executable and which one of these is the best choice considering the selected local decision strategy. A list of all executable operations is presented with the best next step listed first. The user may decide whether or not to follow the recommendation and execute a next step. In any case, the recommendation service will again calculate a recommendation for the new state of the process until the end product is produced.

A detailed description of a PDM execution example using the prototype with a local decision strategy can be found in [44]. Note that it is possible to change a local decision strategy during the execution of a case. This is useful, for example, if a company decides to switch from a low cost strategy to a strategy focusing on a short throughput time to be able to deal with extra case arrivals as quickly as possible. In contrast to this, changing a global strategy during the execution of a case is not allowed, since this would force the system to recalculate the complete strategy. Recalculation may take some time and therefore is not acceptable at run-time.

7 Related Work

Most of the workflow approaches described in literature and supported by commercial systems and academic prototypes assume a procedural approach [3]. Although there are huge differences in the expressiveness and the suitability of procedural languages ranging from YAWL and BPMN to BPEL and ADEPT, the starting point of all these languages is the modeling of control-flow dependencies among tasks. Few alternative approaches have been proposed. For example, the DECLARE system supports a more declarative style of modeling grounded in LTL [27,28,29]. However, also DECLARE focuses on control-flow dependencies among tasks. This paper proposes to start from data and dependencies between data elements. The PDM has a similar role as the Bill-of-Material (BPM) in production planning. This relation has been explored in the nineties [1,30]. However, no concrete workflow support was provided so far. In our earlier work [1,34,35,44] we showed that in some cases the PDM can be converted into procedural model. However, as indicated in this paper this has the drawback that logical dependencies and performance

related concerns get mixed and it becomes impossible to dynamically change the strategy.

Also related is the work on case handling as it is supported by systems such as BPM[one [5]. In the case handling paradigm the availability of data also plays a dominant role. However, hard-coded control-flow dependencies need to be added to steer the user. The case handling paradigm is one of many approaches aiming at providing flexibility in workflows. Clearly, the direct execution of a PDM also offers flexibility. However, most of the literature focusing on workflow flexibility assumes a procedural language and changes of the underlying model. See for [46,37] for related work and typical flexibility mechanisms.

Recently, there has been revival of approaches that balance data and process aspects, cf. the artifact-centric approach [7] and the proclets approach [4]. However, these approaches do not use the data dependencies as a way to replace control-flow.

Most related to our work is the work by Kress and Seese on Executable Product Models [18,19]. This approach is based on our definition of a PDM [1,34,35] and the authors also propose to directly execute the PDM. Kress and Seese combine this with multi-agent systems, reinforcement learning, and genetic algorithms. The core idea is to learn how to best execute the PDM to reach particular goals.

The workflow solution presented in this paper combines DECLARE [28] and ProM [2,12]. The core idea is that a recommendation service in ProM provides guidance to users of the DECLARE system. Such a recommendation-based approach was proposed in [38] and our current system is using this idea. However, the recommendation approach presented in [38] is not using a PDM. Instead a model is learned and it is assumed that all logical dependencies and other constraints are handled outside of the recommendation service. Hence, this approach is not data driven and serves as an add-on rather than as a full-fledged workflow engine. The recommendation approach presented in [38] is closely linked to *recommender systems* [8,9,36]. Recommender systems are widely used in other domains such as information filtering. They are used to recommend items (films, television, video on demand, music, books, news, images, web pages, products, etc.) that are likely to be of interest to the user. These systems typically do not consider a process, i.e., the creation of a product using a sequence operations. The broader area of Web mining, i.e., the application of data mining techniques to discover patterns by observing users of the Web [15], has the same characteristic. Therefore, these techniques cannot be applied to workflow management. Moreover, these approaches aim at analysis rather than execution support involving data and processes.

8 Conclusion

This paper presented a new approach to workflow execution on the basis of a product description, named the Product Data Model (PDM). Instead of aiming at the derivation of a workflow design that is in general the best possible way to generate an informational product, the PDM *itself* is proposed as the vehicle to steer a workflow's execution. Based on data element values readily available for a specific case on the one hand and a selected strategy on the other hand, this approach recommends the next step that should be performed for the case. There is a clear separation of concerns at work here: the product data model is based on functional requirements while the selected strategy focuses on performance (e.g. minimizing costs or time). Therefore, the execution of one case can be still dynamic and flexible, i.e. changing strategies along the execution.

We have introduced two types of strategies to calculate the recommendations: (i) a global decision strategy, which takes into account the effect of the current decision that is made on future decisions, and (ii) a number of local decision strategies, which only look at the set of directly available, executable steps. The first approach is computationally demanding and may be infeasible for large decision problems, but presents the overall best decision for a case. The second approach is much faster and more flexible, but may result in sub-optimal solutions. We have assessed the performance of several strategies by a simulation study, showing that the local decision strategies indeed reach a near-optimal solution to the global decision strategy. Finally, the feasibility of the presented ideas is shown by the description of a fully operational prototype that supports the approach presented in this paper.

A limitation to this work is that the different decision strategies presented in this paper all focus on improving performance while considering a case in isolation. The best decisions are defined in terms of the performance objective on the case level, e.g. the minimization of throughput time or cost for a single case. Optimization on the process level with respect to e.g. the utilization of the process, or the optimal distribution of work among resources, is not considered yet. Kress et al. [18,19] propose an approach to optimize the whole process with an agent based system. However, they do not consider the performance of a single case. Moreover, the presented decision strategies are not yet able to deal with parallel execution of operations. Concurrency would lead to difficulties with different values for the same data element. In general, data updates are an issue for the use of PDMs and other data-driven approaches. Note that current data-driven workflow management systems, e.g. the case-handling system FLOWer, prevent parallel executions for the same case by blocking the whole case as soon as someone is working on it.

For future work we plan to cooperate with our industrial partners to incorporate the direct execution of a PDM and the decision strategies in a commercial tool. The availability of such a tool will certainly enhance the practical applicability of the ideas, which will be useful to further evaluate the approach in practice.

Acknowledgements

We would like to thank Maja Pesic for her work on DECLARE. Moreover, we thank the many people working on ProM. Their work enabled us to realize the workflow solution presented in this paper.

References

- [1] W.M.P. van der Aalst. On the Automatic Generation of Workflow Processes based on Product Structures. *Computers in Industry*, 39:97–111, 1999.
- [2] 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.
- [3] W.M.P. van der Aalst, A.H.M. ter Hofstede, B. Kiepuszewski, and A.P. Barros. Workflow Patterns. *Distributed and Parallel Databases*, 14(1):5–51, 2003.
- [4] W.M.P. van der Aalst, R.S. Mans, and N.C. Russell. Workflow Support Using Procllets: Divide, Interact, and Conquer. *IEEE Bulletin of the Technical Committee on Data Engineering*, 32(3):16–22, 2009.
- [5] 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.
- [6] K. Bhattacharya, C. Gerede, R. Hull, R. Liu, and J. Su. Towards formal analysis of artifact-centric business process models. *Lecture Notes in Computer Science*, 4714:288, 2007.
- [7] K. Bhattacharya, C. Gerede, R. Hull, R. Liu, and J. Su. Towards Formal Analysis of Artifact-Centric Business Process Models. 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 288–304. Springer-Verlag, Berlin, 2007.

- [8] R. Burke. The Wasabi Personal Shopper: A Case-Based Recommender System. In *Proceedings of the 11th Conference on Innovative Applications of Artificial Intelligence*, pages 844 – 849. American Association for Artificial Intelligence, 1999.
- [9] R. Burke. A Case-Based Reasoning Approach to Collaborative Filtering. In E. Blanzieri and L. Portinale, editors, *Proceedings of the 5th European Workshop on Advances in Case-Based Reasoning (EWCBR 2000)*, volume 1898 of *Lecture Notes in Computer Science*, pages 370–379. Springer-Verlag, Berlin, 2000.
- [10] CPN Tools Home Page. <http://wiki.daimi.au.dk/cpntools/cpntools.wiki>.
- [11] T.H. Davenport. *Process Innovation : Reengineering Work through Information Technology*. Harvard Business School Press, Boston, 1993.
- [12] B.F. van Dongen, A.K. Alves de Medeiros, H.M.W. Verbeek, A.J.M.M. Weijters, and W.M.P. van der Aalst. The ProM framework: A New Era in Process Mining Tool Support. In G. Ciardo and P. Darondeau, editors, *Application and Theory of Petri Nets 2005*, volume 3536 of *Lecture Notes in Computer Science*, pages 444–454. Springer-Verlag, Berlin, 2005.
- [13] 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.
- [14] M. Hammer and J. Champy. *Reengineering the Corporation*. Nicolas Brealey Publishing, London, 1993.
- [15] P. Hofgesang. Online mining of web usage data: An overview. In *Web Mining Applications in E-commerce and E-services*, volume 172 of *Studies in Computational Intelligence*, pages 1–24. Springer-Verlag, Berlin, 2009.
- [16] K. Jensen. *Coloured Petri Nets. Basic Concepts, Analysis Methods and Practical Use. Volume 1, Basic Concepts*. Monographs in Theoretical Computer Science. Springer-Verlag, Berlin, 1997.
- [17] H.J. Johansson, P. McHugh, A.J. Pendlebury, and W.A. Wheeler. *Business Process Reengineering: Breakpoint Strategies for Market Dominance*. Wiley & Sons, 1993.
- [18] M. Kress, J. Melcher, and D. Seese. Introducing Executable Product Models for the Service Industry. In *Proceedings of the 40th Hawaii International Conference on System Sciences (HICSS '07)*, pages 1–10. IEEE Computer Society, 2007.
- [19] M. Kress and D. Seese. Executable Product Models - The Intelligent Way. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics (ISIC '07)*, pages 1987 – 1992, 2007.
- [20] M. McDonald, M. Blosch, T. Jaffarian, L. Mok, and S. Stevens. Growing IT's Contribution: The 2006 CIO Agenda. *Gartner January*, 2006.
- [21] M. McDonald and T. Nunno. Creating Enterprise Leverage: The 2007 CIO Agenda. *Gartner January*, 2007.

- [22] M. McDonald, T. Nunno, and D. Aron. Making the Difference: The 2008 CIO Agenda. *Gartner January*, 2008.
- [23] D.C. Montgomery and G.C. Runger. *Applied Statistics and Probability for Engineers*. John Wiley and Sons, New York, 1999.
- [24] B. Mutschler, M. Reichert, and J. Bumiller. Unleashing the effectiveness of process-oriented information systems: Problem analysis, critical success factors and implications. *IEEE Transactions on Systems, Man, and Cybernetics (Part C)*, 38(3):280–291, 2008.
- [25] J.A. Orlicky. Structuring the Bill of Materials for MRP. *Production and Inventory Management*, pages 19–42, Dec 1972.
- [26] S.S. Panwalkar and W. Iskander. A Survey of Scheduling Rules. *Operations Research*, 25:45–61, 1977.
- [27] M. Pesic and W.M.P. van der Aalst. A Declarative Approach for Flexible Business Processes. In J. Eder and S. Dustdar, editors, *Business Process Management Workshops, Workshop on Dynamic Process Management (DPM 2006)*, volume 4103 of *Lecture Notes in Computer Science*, pages 169–180, Berlin, 2006. Springer-Verlag.
- [28] M. Pesic, M.H. Schonenberg, and W.M.P. van der Aalst. DECLARE: Full Support for Loosely-Structured Processes. In *Proceedings of the 11th IEEE International Enterprise Distributed Object Computing Conference (EDOC 2007)*, pages 287–300. IEEE Computer Society, 2007.
- [29] M. Pesic, M.H. Schonenberg, N. Sidorova, and W.M.P. van der Aalst. Constraint-Based Workflow Models: Change Made Easy. In R. Meersman and Z. Tari, editor, *On the Move to Meaningful Internet Systems 2007: OTM 2007 Confederated International Conferences, Part I.*, volume 4803 of *Lecture Notes in Computer Science*, pages 77–94, 2007.
- [30] E.A.H. Platier. *A Logistical View on Business Processes: BPR and WFM concepts*. PhD thesis, Eindhoven University of Technology, Eindhoven, 1996. (in Dutch).
- [31] M.L. Puterman. *Markov Decision Processes*. Wiley, New York, 1994.
- [32] R. Ramasesh. Dynamic Job Shop Scheduling: A Survey of Simulation Research. *OMEGA International Journal of Management Science*, 18(1):43–57, 1990.
- [33] H.A. Reijers. Product-based design of business processes applied within the financial services. *Journal of Research and Practice in Information Technology*, 34(2):110–122, 2002.
- [34] H.A. Reijers. *Design and Control of Workflow Processes: Business Process Management for the Service Industry*, volume 2617 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 2003.
- [35] H.A. Reijers, S. Limam Mansar, and W.M.P. van der Aalst. Product-Based Workflow Design. *Journal of Management Information systems*, 20(1):229–262, 2003.

- [36] P. Resnick and H.R. Varian. Recommender Systems. *Communication of the ACM*, 40(3):56–58, 1997.
- [37] 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.
- [38] H. Schonenberg, B. Weber, B.F. van Dongen, and W.M.P. van der Aalst. Supporting Flexible Processes Through Log-Based Recommendations. In M. Dumas, M. Reichert, and M.-C. Shan, editors, *Proceedings of the 6th International Conference on Business Process Management (BPM 2008)*, volume 5240 of *Lecture Notes in Computer Science*, pages 51–66. Springer-Verlag, Berlin, 2008.
- [39] A. Sharp and P. McDermott. *Workflow Modeling: Tools for Process Improvement and Application Development*. Artech House, 2nd edition, 2009.
- [40] E.A. Silver, D.F. Pyke, and R. Peterson. *Inventory Management and Production Planning and Scheduling*. John Wiley and Sons, Hoboken, NJ, 1998.
- [41] H.C. Tijms. *A First Course in Stochastic Models*. Wiley, Chichester, England, 2003.
- [42] P. Trkman. The Critical Success Factors of Business Process Management. *International Journal of Information Management*, 2009.
- [43] M. Vakola and Y. Rezgui. Critique of existing business process re-engineering methodologies. *Business Process Management Journal*, 6(3):238–50, 2000.
- [44] I. Vanderfeesten. *Product-Based Design and Support of Workflow Processes*. PhD thesis, Eindhoven University of Technology, Eindhoven, the Netherlands, 2009.
- [45] I. Vanderfeesten, W.M.P. van der Aalst, and H.A. Reijers. Modeling a Product-Based Workflow System in CPN tools. In K. Jensen, editor, *Proceedings of the Sixth Workshop on the Practical Use of Coloured Petri Nets and CPN Tools (CPN 2005)*, volume 576 of *DAIMI*, pages 99–118, Aarhus, Denmark, October 2005. University of Aarhus.
- [46] 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.