Karlsruher Institut für Technologie

Automatic Generation of Optimized Process Models from Declarative Specifications

Richard Mrasek, Jutta Mülle, and Klemens Böhm

2014

Fakultät für **Informatik**

Automatic Generation of Optimized Process Models from Declarative Specifications

Richard Mrasek, Jutta Mülle, and Klemens Böhm

Karlsruhe Institute of Technology (KIT)
Institute for Program Structures and Data Organization
76131 Karlsruhe, Germany
{ richard.mrasek | jutta.muelle | klemens.boehm }@kit.edu

Abstract. Process models often are generic, i. e., describe similar cases or contexts. For instance, a process model for commissioning can cover both vehicles with an automatic and with a manual transmission, by executing alternative tasks. A generic process model is not optimal compared to one tailored to a specific context. Given a declarative specification of the constraints and a specific context, we study how to automatically generate a good process model and propose a novel approach. We focus on the restricted case that there are not any repetitions of a task, as is the case in commissioning and elsewhere, e. g., manufacturing. Our approach uses a probabilistic search to find a good process model according to quality criteria. It can handle complex real-world specifications containing several hundred constraints and more than one hundred tasks. The process models generated with our scheme are superior (nearly twice as fast) to ones designed by professional modelers by hand.

1 Introduction

Scheduling tasks so that the overall execution is efficient and at the same time no constraints are violated continues to be a fundamentally important problem. Process models describe the possible arrangements of the tasks.

Example **1.** *Our application scenario is commissioning. Commissioning means configuring and testing the electronic components of a vehicle at the end of its production. Process models describe the arrangement of the configuration and testing tasks. For instance, a factory worker has to configure the transmission and to activate the anti-theft system. The transmission can either be manual, i. e., Task M does the configuration or automatic (Task A). Task T activates the anti-theft system. Before the activation, a central computer needs to generate a master key (Task G), and it opens the connection to the specific control unit (Task O). The connection has to be closed before the process finishes (Task C). The configuration of the transmission and the activation of the anti-theft system require a running engine; Task E turns it on. Figure 1(a) shows the tasks that may be part of the commissioning. The second column is the expected processing time of the tasks. Commissioning always has a context, i.e., the variation of*
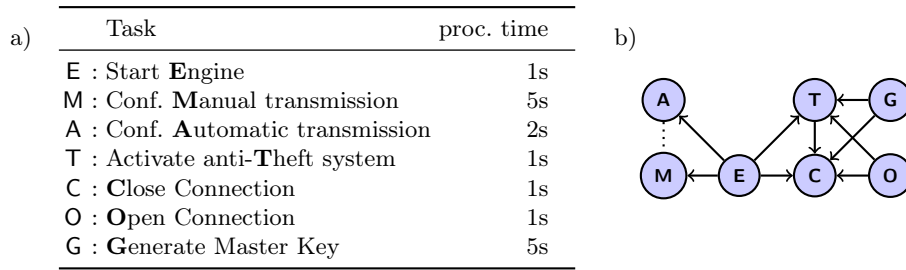
Fig. 1. The Tasks for the Commissioning Scenario (a) and the Ordering Relationship Graph (b).

the vehicle, its components, their relationships and the constraints the vehicle currently tested must fulfill. The variation determines which tasks have to be executed, e. g., a car with a manual transmission requires different tasks than a car with an automatic transmission.

A context c determines the tasks $\mathcal{T}_c$ required for a process. It is infeasible to model all processes for each possible set of required tasks by hand. This calls for generic process models for several contexts. With such generic models however, one optimal arrangement of tasks for any context does not exist.

Example 2. *The context characteristic* transmission *determines the required tasks as follows: If the vehicle has an automatic transmission, the commissioning requires execution of the tasks $\mathcal{T}_c = \{$E, A, T, C, O, G$\}$; for a manual transmission in turn the tasks are $\mathcal{T}_c = \{$E, M, T, C, O, G$\}$. Figure 1(b) shows the dependencies between the tasks as a graph. Directed edges represent ordering dependencies, while dashed lines represent exclusive dependencies. The graph is the declarative specification we generate the process model from. Section 3 shows how one can generate such a specification from input in other languages. Section 2 will introduce the notation behind that graph structure.*

Figure 2 shows two generic process models that comply with the dependency graph of Figure 1(b). Figure 2(a) has a shorter processing time if the transmission is automatic (7s to 10s). For a manual transmission in turn, the process model in Figure 2(b) has a shorter processing time (8s to 10s).


Fig. 2. Two distinct Process Models for the graph in Figure 1(b)

With at least one process model for each possible context, the number of such models increases exponentially with the number of context characteristics.

Example 3. *Say a vehicle has 10 different context characteristics, e. g., the kind of transmission, the navigation system, the safety system, etc. If each option can occur or not $2^{10} = 1024$ different contexts are possible.*

Next, several models typically are possible for a given context. The problem studied here is how to generate a good process model for a given context from a declarative specification. The model should be good according to predefined quality criteria, e. g., average throughput time of the process. Process models that comply with the dependencies can be very different with respect to quality and performance criteria. Section 6 will show that models generated with our approach are about 50% faster than ones designed by professionals with years of experience. We focus on the restricted case that there are none repetitions, i. e., we focus on process trees with inner nodes SEQ, AND, XOR, but not LOOP. There is a number of settings with this characteristic, for instance in manufacturing. In particular, loops are unnatural in commissioning processes, since a feature is tested only once. On the other hand, if a problem occurs and is fixed, a new commissioning process is started. Another assumption, which also holds for commissioning and elsewhere, is that context information together with experience from the past allows to reliably estimate the processing time of individual tasks.

The generation of a process model from a declarative specification bears several challenges. There often is a great variety of models that fulfill the specification, as mentioned before. To illustrate, the sequential arrangement of n nodes, in the absence of any constraint, give way to $n!$ different process models. For the largest process in our evaluation, 4.12×10^{340} models are possible. Generating all possible models is not possible. It is challenging to detect a good process model that does not violate any constraint.

Example 4. *There are four tasks A, B, C, and D. Suppose that the following constraints exist: B must always occur before A and C ($B \rightarrow A$, $B \rightarrow C$), and D always occurs before C ($D \rightarrow C$). It seems to be a good idea to put A and C in parallel, because this might reduce the throughput time. But putting A and C in parallel rules out having A and D in parallel.*

Related work in process synthesis is fully automatic only for processes that are fully specified by their dependencies [8][29]. In case of an under-specification, [8] requires a process modeler to manually make decisions, and [29] requires a manual clustering of the constraints. This is not practical, because of the daunting number of possible models. To this end, we propose a novel process synthesis algorithm whose output on the one hand complies with the dependencies and on the other hand is good according to predefined criteria. Our approach is as follows: First, it uses a modular decomposition of the dependencies to detect the fully specified regions of the process as well as the under-specified ones, so-called prime components. For each prime component, our approach partitions the corresponding ordering graph systematically, as follows. It selects a pivot

element and generates several smaller ordering graphs from the pivot partition. We reduce the problem in a divide and conquer fashion until it is small enough to explicitly generate all possible models. We repeat this for different pivot elements to have a better coverage according to our quality criterion the throughput time of the process. Other criteria such as overall energy consumption are possible as well. As we show in the evaluation with thousands of non-trivial process models, our approach is efficient, i. e., is able to test thousands of models in under a second, checking for complex constraints. On average, our approach nearly halves the processing time compared to the reference processes, which already are the output of a careful intellectual design. Our approach can handle complex real-world specifications containing several hundred dependencies as well as more than one hundred tasks. In our evaluation, the process models generated contain between 98 and 185 tasks, and their arrangement typically is nontrivial.

Section 2 introduces some fundamentals. Section 4 describes our algorithm for the process generation. Section 6 features our evaluation. Section 7 discusses related work, and Section 8 concludes.

2 Fundamentals

A meaningful input for process synthesis is the declarative specification in the form of an ordering relation graph (ORG) [25]. Section 3 shows how an ORG can be computed from other specification languages, e. g., DECLARE. The modular decomposition of a graph yields its components and implies a hierarchical structure of components called the Modular Decomposition Tree (MDT)[20], see Subsection 2.3. The MDT separates the under-specified regions from the fully specified ones.

2.1 Ordering Relation Graph

In an ordering relation graph, each node represents a task. Each edge represents a dependency between tasks. The dependencies consist of ordering dependencies, i. e., in which order do the tasks occur, and exclusive dependencies, i. e., when do two tasks exclude each other.

Definition 1. *The ordering relation graph is a directed attributed graph $G = (V, E)$, with V being nodes and $E \subseteq V \times V$ the edges. Each node corresponds to a task. E consists of two subsets $E_\rightarrow$ and $E_\#$ such that $E = E_\rightarrow \cup E_\#$ and $E_\rightarrow \cap E_\# = \emptyset$. $E_\rightarrow$ defines the ordering relation, i. e., two tasks that should be in a specific order have an edge in $E_\rightarrow$. $E_\rightarrow$ is transitive and anti-symmetric:*

$$(transitive) \quad \forall (x,y), (y,z) \in E_\rightarrow : (x,z) \in E_\rightarrow$$
$$(antisymmetric) \quad \forall (x,y) \in E_\rightarrow : (y,x) \notin E_\rightarrow$$

$E_\#$ defines the exclusiveness relation, i. e., if two tasks exclude each other they share an edge in $E_\#$. $E_\#$ is symmetric, i. e., $\forall (x,y) \in E_\# : (y,x) \in E_\#$. We do not allow self-edges, i. e., $\forall v \in V : (v,v) \notin E$.

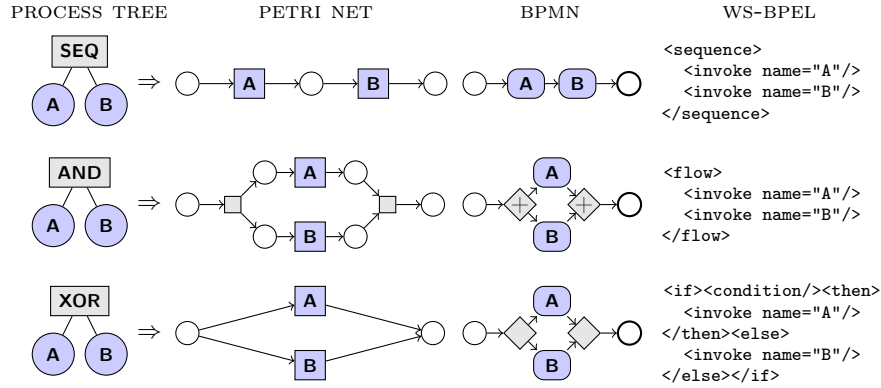| PROCESS TREE	PETRI NET	BPMN	WS-BPEL

Fig. 3. Transformation of a SEQ-, AND-, and XOR-Node of a Process Tree to Petri net, BPMN or WS-BPEL.

Note that $E_\rightarrow$ does not contain any cycle. For each task we determine the processing time. The average error of the estimated execution times of our tasks from our application scenario is less than 17%. We had calculated these times by analyzing the logs of existing traces.

Definition 2. *The neighborhoods* $N^{out}(v)$, $N^{in}(v)$ *of a node* v *are defined as:*

$$N^{out}(v) := \{w \mid w \in V \wedge (v,w) \in E_\rightarrow\} \quad N^{in}(v) := \{w \mid w \in V \wedge (w,v) \in E_\rightarrow\}$$

$N^{out}(v)$ *is the set of nodes with an incoming ordering edge from* v. $N^{in}(v)$ *is the set of nodes that have an outgoing ordering edge to* v. *For a set of nodes* V *the incoming and outgoing set are defined as* $N^{out}(V) := \bigcup_{v \in V} N^{out}(v)$ *and* $N^{in}(V) := \bigcup_{v \in V} N^{in}(v)$ *respectively.*

The ordering relation graph (ORG) describes the tasks in a process model as nodes and their relationships to each other by different types of edges. In contrast to an imperative process language like BPMN, ORG is a declarative description and not necessarily fully specified.

2.2 Process Tree

We want to generate the process model in the form of a process tree (PT). In contrast to a graph-based process model, the process tree has two important characteristics. First, it can be easily transformed into an executable process language, see Figure 3. Second, a process tree is sound by default [14]. This means the following: First, the process will terminate properly. Second, for each task there is at least one process instance containing it. Each *process tree* $PT = (\mathcal{V}, \mathcal{E})$ is an ordered tree, thus a rooted tree for which an ordering is specified for the children of each vertex. $\mathcal{V}$ consists of leaf nodes $\mathcal{V}_t$ and inner nodes $\mathcal{V}_c$, $\mathcal{V}_t \cup \mathcal{V}_c = \mathcal{V}$, $\mathcal{V}_t \cap \mathcal{V}_c = \emptyset$. Each leaf node corresponds to a task, and each inner node corresponds to a control structure. In this paper we consider three control

Fig. 4. An ordering relation graph (a), its modular decomposition (b) and the corresponding modular decomposition tree (c)

structures, namely sequence SEQ, parallel AND and exclusive XOR. These control structures correspond to the basic control workflow patterns [4]. This study focuses on the synthesis of process models without cycles. Hence, we do not define a loop operator. It is possible to model the commissioning processes using those control structures. Each control structure can be translated to another block-oriented language, e. g., WS-BPEL, OTX, or to a graph-oriented process language, e. g., Petri nets, BPMN. Figure 3 shows how each of the control structure of a PT can be easily transformed to three of the major process languages.

2.3 Modular Decomposition

We want to generate a process tree from the declarative specification, i. e., from the ORG. Let $G = (V, E)$ be such a graph. For any $W \subseteq V$ we say that $G_W(V^W, E^W)$ is the sub-graph induced by W, i. e., $V^W = W$ and $E^W = E \cap (W \times W)$. We call W a component iff $\forall v, v' \in W$, $N^{out}(v) \backslash W = N^{out}(v') \backslash W$ and $N^{in}(v) \backslash W = N^{in}(v') \backslash W$. Thus v and v' have identical neighborhoods outside of W. In other words, a component consists of tasks with the same dependencies regarding tasks outside of the component. In our use case, a component often consists of tasks operating on the same electronic control unit of the vehicle. W is a strong component if, for each component $W' \subseteq V$, one of the following holds: $W \cap W' = \emptyset$, $W \subseteq W'$, or $W' \subseteq W$. The decomposition of a graph into strong components is called Modular Decomposition, and the resulting hierarchical structure is called Modular Decomposition Tree (MDT). Figure 4 shows a simple ordering relation graph (a), its decomposition in four components (b) and the corresponding modular decomposition tree (c). [20] shows that a node W in a MDT with children $S_1, S_2, \ldots, S_k$ is of one of the following:

> Complete : $\forall I \subset \{1, \ldots, k\}$, with $1 < |I| < k : \bigcup_{i \in I} S_i$ is a component
> Prime   : $\forall I \subset \{1, \ldots, k\}$, with $1 < |I| < k : \bigcup_{i \in I} S_i$ is not a component

***Example* 5.** *The root node in Figure 4 (c) is a prime node. None of the subsets of the children with size 2 or 3, e. g., $\{X, R\}$ or $\{R, Y, Z\}$, do form a component.*

A complete component W with the induced graph $G_W(V^W, E^W)$ either does not contain any edges or is a clique in $E^W_\#$ or $E^W_\rightarrow$, see the proof of Lemma 1. A

complete component can easily be transformed to a process tree deterministically, see [25]. For a prime component our approach will use a heuristic optimization.

Lemma 1. *A strong complete component W is of exactly one of four types:*

trivial $\quad: |V^W| = 1$
serial $\quad\;\;: For\ every\ v, v' \in V^W \;: (v, v') \in E_{\to}^W \;\lor\; (v', v) \in E_{\to}^W.\ Recall\ that\ the$
$\qquad\qquad edges\ in\ E_{\to}^W\ are\ cycle\text{-}free.$
branch $: For\ every\ v, v' \in V^W \;: (v, v') \in E_{\#}^W$
parallel $: For\ every\ v, v' \in V^W \;: (v, v') \notin E^W$

Proof. For $|V^W| = 1$ the component is trivial. For $|V^W| = 2$ the pair (v, v') is either $(v, v') \in E_{\to}^W$ (serial), $(v, v') \in E_{\#}^W$ (branch), or $(v, v') \notin E^W$ (parallel). For $|V^W| \geq 3$ there exist two pairs of nodes (v_1, v_2), (v_3, v_4) with $(v_1, v_2) \neq (v_3, v_4)$. At least one element in the pair differs, with out loss of generality let $v_1 \neq v_3$. Each set of two elements is a component, and as shown earlier it is either a serial, branch, or parallel. We call this the *type* of the pair, namely $type_{v_1,v_2}$. Assume that the lemma is false. Then at least two pairs of nodes exist with $type_{v_1,v_2} \neq type_{v_3,v_4}$. W is a strong component and thus all combinations of child elements are strong components. The set of nodes $\{v_1, v_3\}$ is a strong component $\Rightarrow N(v_1) = N(v_3) \Rightarrow type_{v_2,v_3} = type_{v_1,v_2}$. If $v_2 = v_4$ this is a contradiction to the assumption, so let $v_2 \neq v_4$. The set of nodes $\{v_2, v_4\}$ is a strong component $\Rightarrow N(v_2) = N(v_4) \Rightarrow type_{v_3,v_4} = type_{v_2,v_3} \Rightarrow type_{v_3,v_4} = type_{v_1,v_2}$. This is a contradiction to the assumption. $\qquad\square$

[20] proves that the decomposition of a directed graph (V, E) can be done in $O(|V| + |E|)$, thus in time linear with the size of the graph – We use the MDT to transform the ORG into a process tree.

3 Generating an Ordering Relationship Graph

In this section, we explain how to generate an ordering relationship graph from a declarative specification, e. g., DECLARE, Compliance Rule Graphs or BPMN-Q. In our application domain, each task executes exactly once and for simplicity we consider here the core set of specification elements that are supported by the majority of graphical specification languages, see Table 3.

Figure 5 shows the core set of these elements and the representation for DECLARE. The study of [26] indicates a similar set of mostly used specification elements. Empirical studies [11] as well as our experience [22] show that this core set is sufficient for the majority of the specifications. Bounded existence and chain ordering only occur rarely (in 10 out of 555 models) [11].

We describe our algorithm for DECLARE, for Compliance Rule Graphs or BPMN-Q the algorithm would function equivalently. We start with a set of tasks $\mathcal{T}_c$ the commissioning process has to comprise for the context c. Next, we check if a task $t \in \mathcal{T}_c$ has an outgoing response, succession, responded existence or co-existence edge to a task t_2 not in $\mathcal{T}_c$. If so, we add t_2 to $\mathcal{T}_c$. Our algorithm repeats these

Table 1. Four different declarative specification language and the patterns they can represent directly (black mark) or by combination (gray mark).

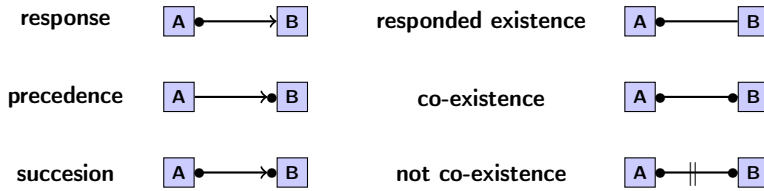|	Existence	Resp. Existence	Not co-existence	Bounded Existence	Response	Precedence	Succession	Chain Ordering	Data Constraints
BPMN-Q [7]	✓				✓	✓	✓		
Comp. Rule Graph [18]	✓	✓	✓		✓	✓	✓		✓
Declare [1]	✓	✓	✓	✓	✓	✓	✓	✓	
Prop. Spec. Pattern [10]	✓	✓	✓	✓	✓	✓	✓	✓	

Fig. 5. Core set of the DECLARE elements

steps until no further change happens (Step 1-3). For each task we generate a node in the ORG (Step 7-9). If an ordering edge (an edge with an arrow) exists in the DECLARE graph between nodes $v, w \in \mathcal{T}_c$ we add an edge $(v, w) \in E_\rightarrow$ to the ORG (Step 10-12). At last, the algorithm delivers the ordering relationship graph ORG (Step 13).

***Example* 6.** *We want to generate the ordering relationship graph for the DE-CLARE graph in Figure 6(a). The commissioning requires executing the tasks B and E, highlighted in dark red in Figure 6(a). E needs the occurrence of C and C the precedence of task D. The result is an extended set of tasks, see Figure 6(b)–*


Fig. 6. Generation of a ORG from a DECLARE graph and a list of tasks $\mathcal{T}_c$

Algorithm 1 generateORG (DECLARE graph DG, task set $\mathcal{T}_c$) : ORG
___
1: **while** $t \in \mathcal{T}_c$ has an activation edge to a task $t_2 \notin \mathcal{T}_c$ **do**
2:      $\mathcal{T}_c \leftarrow \mathcal{T}_c \cup \{t_2\}$
3: **end while**
4: **for all** $t \in \mathcal{T}_c$ **do**
5:      Add a state t to ORG
6: **end for**
7: **for all** response, precedence or succession edges between tasks $t_1, t_2 \in \mathcal{T}_c$ **do**
8:      Add an edge $(t_1, t_2) \in E_\rightarrow$ to ORG
9: **end for**
10: **return** The ordering relationship graph ORG
___

(c). Finally we transform the DECLARE *graph to an ordering relationship graph, see Figure 6(d).*

4   Generating a Process Tree

In this section we explain the conceptual design of our approach. Subsection 4.1 gives an overview, and Subsection 4.2 states how the algorithm handles under-specified regions.

___
Algorithm 2 synthesize(ORG G, context c): ProcessTree PT
___
1: Determine $\mathcal{T}_c$ from c
2: $G \leftarrow$ subgraph G_W of G with the nodes $W = \mathcal{T}_c$
3: $PT \leftarrow$ Modular Decomposition of G
4: **for all** prime nodes $\mathcal{P} \in PT$ **do**
5:      Process tree $\text{PT}_\mathcal{P} \leftarrow synPrime(\mathcal{P})$
6:      Replace $\mathcal{P}$ with $PT_\mathcal{P}$
7: **end for**
8: **for all** leaf nodes $l \in PT$ **do**
9:      **if** l is a partition leaf node **then**
10:          $G_l \leftarrow$ ORG of l
11:          Process tree $\text{PT}_l \leftarrow synthesize(G_l)$
12:          Replace l with PT_l
13:      **end if**
14: **end for**
15: **return** PT
___

4.1   Overview of the Automatic Generation

Our goal is to automatically generate a process model from a declarative description. Algorithm 2 synthesizes a process tree from an ORG and a context c. The context c determines the required tasks $\mathcal{T}_c$ (Line 1). We then reduce the ORG G to the subgraph G_W with the nodes $W = \mathcal{T}_c$ (Line 2). The algorithm then computes a modular decomposition of the ORG (Line 3 in Algorithm 2).

The resulting modular decomposition tree (MDT) may contain both complete and prime components. For complete components, a transformation to process fragments exists, cf. [25]. For a prime component in turn, several fragments are possible, see Figure 7. In other words, each prime component stands for an under-specified region. For each prime component $\mathcal{P}$, we use a probabilistic optimization to find a solution (Line 5). We replace $\mathcal{P}$ with the solution found (Line 6). *synPrime()* splits the ORG of the prime components into partitions. It generates a graph with one node for each of these partitions. The algorithm recursively calls itself, in order to replace each node with a subtree. Finally, our approach transforms the PT into a process language, e.g., BPMN, WS-BPEL.


Fig. 7. The Neighborhood Graph to Directly Generate a Process Tree (a), three Possible Process Trees (b)(c)(d) for the Graph (a).

4.2 Under-Specified Regions

Each prime component $\mathcal{P}$ induces a graph $G_{\mathcal{P}} = (V_{\mathcal{P}}, E_{\mathcal{P}})$. $V_{\mathcal{P}}$ denotes the set of strong components that belong to $\mathcal{P}$. Figure 4 shows that the graph $G_{\mathcal{P}}$ for the prime component $\mathcal{P}$ consists of $V_{\mathcal{P}} = \{X, R, Y, Z\}$ with $E_{\mathcal{P}} = \{(R \rightarrow X), (R \rightarrow Y), (Z \rightarrow Y)\}$. $\mathcal{P}$ is not fully specified and thus there does not exist a unique corresponding process tree. Due to the large number of possible process models for a prime graph $G_{\mathcal{P}}$ it is not feasible to construct every possible one.

The modular decomposition detects the fully specified and the under-specified regions of the process. Even for small prime components it is not possible to generate and test all possible process models. Our overall idea is to reduce the size of the graph induced by a prime component iteratively until the number of remaining solutions is low (< 100) so that we can solve the problem. See Figure 7. Our intuition for the reduction is to select a pivot node v and detect which nodes (V_1) must occur before v, and which nodes (V_2) can be scheduled in parallel to v. V_1 as well as V_2 imply two smaller ordering graphs. We repeat this with several different pivot nodes. Our approach randomly selects a node $v \in V_{\mathcal{P}}$ with $N^{out}(v) = \emptyset$ as pivot node. Lemma 4 will show why we need this characteristic. The ORG $G_{\mathcal{P}}$ is cycle-free, and thus a node v with $N^{out}(v) = \emptyset$ always exists.

Definition 3. *The zero neighborhood of a pivot node v is $N^{(0)}(v) := \{v\}$, $N^{(1)}(v) := N^{in}(v)$. For $i \in \mathbb{N}$, $i > 1$ we define the i-neighborhood as:*

$$N^{(i)}(v) := \begin{cases} \left(\bigcup_{v' \in N^{(i-1)}(v)} N^{out}(v')\right) \setminus N^{(i-2)}(v) & \text{if } i \in \{2, 4, 6, \dots\} \\ \left(\bigcup_{v' \in N^{(i-1)}(v)} N^{in}(v')\right) \setminus N^{(i-2)}(v) & \text{if } i \in \{3, 5, 7, \dots\} \end{cases}$$

Lemma 2. *The neighborhoods for a pivot node v of a prime component $G_{\mathcal{P}}(V, E)$ contains each node exactly once. This means that:*

1. $\bigcup_{i \in \mathbb{N}} N^{(i)}(v) = V$
2. $\forall i, j \in \mathbb{N}, i \neq j : N^{(i)}(v) \cap N^{(j)}(v) = \emptyset$

Proof.

1. Assume that $G_{\mathcal{P}}$ is not connected $\Rightarrow$ two sub-graphs exists G' with the nodes in the connected graph of v and $G'' := G_{\mathcal{P}} \setminus G'$. This is a contradiction to the assumption that $G_{\mathcal{P}}$ is a prime, because G' and G'' would form a component. Let $w \in V_{\mathcal{P}}$ be a node, with $w \neq v$. $G_{\mathcal{P}}$ is connected thus a shortest undirected path $p = (a_0, a_1, a_2, \dots, a_{n-1}, a_n)$ exists between $v = a_0$ and $w = a_n$. p is alternating thus $\forall i \in [0, n] : (a_i, a_{i+1}) \Leftrightarrow (a_{i+2}, a_{i+1})$, or a shorter path would exist. $a_i \in N^{(i)}(v)$ and thus $w \in N^{(n)}(v)$.
2. Assume $N^{(i)}(v) \cap N^{(j)}(v) \neq \emptyset$ then a node w exists with $w \in N^{(i)}(v) \wedge w \in N^{(j)}(v)$. This means that two shortest undirected paths would exist between w and v. This is a contraction, only one path can be shortest.

$\square$

We use the neighborhood information to partition the graph. Each partition $n^{(i)}$ is a subgraph of the ORG $G_{\mathcal{P}}$ with the nodes $N^{(i)}(v)$. In other words, the partitioning implies a graph G_v where each $n^{(i)}$ is a node. We refer to this graph as the neighborhood graph. Formally, given a pivot node v, the neighborhood graph $G_v = (V_v, E_v)$ is as follows

$$\begin{aligned} V_v &= \{n^i \mid N^{(i)}(v) \neq \emptyset\} \\ E_v &= \{(n^i, n^{i+1}) \mid i \in \{1, 3, \dots\} \wedge n^i, n^{i+1} \in V_v\} \cup \\ &\quad \{(n^{i+1}, n^i) \mid i \in \{0, 2, \dots\} \wedge n^i, n^{i+1} \in V_v\} \end{aligned}$$

The graph contains each non-empty neighborhood as a node.

Example 7. *For the graph in Figure 4(b) and the pivot Y the neighborhoods are: $N^{(0)}(Y) = \{Y\}$, $N^{(1)}(Y) = \{R, Z\}$, $N^{(2)}(Y) = \{X\}$, and for $i > 2$ $N^{(i)}(Y) = \emptyset$. The neighborhood graph $G_Y(V_Y, E_Y)$ for the pivot Y is:*

$$G_Y = (\ \{n^0, n^1, n^2\}\ ,\ \{\ (n^1, n^0)\ ,\ (n^1, n^2)\ \})$$

Example 8. *Figure 8(a) shows a more complex graph which is a prime component, i.e., there is no unique corresponding tree. The possible pivot nodes are in violet. The pivot node at the top of Figure 8(a) leads to the partitioning in Figure 8(b). Figure 8(c) shows the respective neighborhood graph.*

Fig. 8. A prime component (a), its partitioning (b), and the neighborhood graph (c)


Fig. 9. Structure of the Neighborhood Graph

Lemma 3. *The partitioning into the neighborhood graph for a pivot v preserves all order dependencies. In other words, for each edge $(v_1, v_2) \in E_{\mathcal{P}}$, one of the following holds:*

(a) $\exists i \in \mathbb{N}_0 : v_1, v_2 \in N^{(i)}(v)$
(b) $v_1 \in N^{(i)}(v), v_2 \in N^{(j)}(v), i \neq j \Rightarrow (n^i, n^j) \in E_v$

Proof. If $v_1, v_2 \in N^{(i)}(v)$ then the edge is in the subgraph for $N^{(i)}(v)$. We now focus on Case (b).
CASE $1 - i \in \{1, 3, 5, \dots\}$: If $j = i - 1$ then $(n^i, n^j) \in E_v$. Otherwise, if $j \neq i - 1$ then $v_2 \in N^{out}(v_1) \land v_2 \notin N^{(i-1)}(v) \Rightarrow v_2 \in N^{(i+1)}(v) \Rightarrow j = i + 1$, and $(n^i, n^j) \in E_v$.
CASE $2 - i \in \{0, 2, 4, \dots\}$: In this case, a node v_3 exists with $(v_3, v_1) \in E_{\mathcal{P}}$. The transitivity of $E_{\mathcal{P}}$ leads to $(v_3, v_2) \in E_{\mathcal{P}}$. $v_2, v_1 \in N^{out}(v_3) \Rightarrow v_1, v_2 \in N^{(i)}(v) \Rightarrow i = j$, i.e., there is a contradiction. $\square$

Lemma 3 states that our approach does not loose any dependencies. A symmetric solution would be to select pivots with $N^{in}(v) = \emptyset$ and change the definition of the neighborhood accordingly. However, a pivot v with $N^{out}(v) \neq \emptyset \land N^{in}(v) \neq \emptyset$ would loose a dependency, see Lemma 4.

Lemma 4. *The neighborhood graph G_v for a pivot node v with $N^{out}(v) \neq \emptyset \land N^{in}(v) \neq \emptyset$ does not preserve the order dependencies.*

Proof. Let v with $N^{out}(v) \neq \emptyset \land N^{in}(v) \neq \emptyset$ be the pivot node. Then a node $v^- \in N^{in}(v)$ and a node $v^+ \in N^{out}(v)$ exists. The ORG is cycle free, thus $v^+ \neq v^- \neq v$. The following dependencies, i.e., elements of $E_{\mathcal{P}}$, exist: (v^-, v), (v, v^+), and, because of the transitivity of the ORG, $(v^-, v^+) \in E_{\mathcal{P}}$. $v \in N^{(0)}(v)$, $v^+ \in N^{(1)}(v), v^- \in N^{(2)}(v)$ altogether imply that $(v^-, v) \in E_{\mathcal{P}}$ and $(n^2, n^0) \notin E_v$. This means that the dependencies between v and v^- have been lost. $\square$

Algorithm 3 synPrime (Neighborhood Graph $G(V,E)$) : ProcessTree PT

1: Pivot v ← randomly select a node $v \in G$ with $N^{out}(v) = \emptyset$
2: $G_v(V_v, E_v)$ ← calculate neighborhood of v
3: **if** $N^{(\lambda)} = \emptyset$ **then**
4:      **return** (select tree pattern randomly)
5: **else**
6:      **return** $synPrime(G_v)$
7: **end if**

Algorithm 3 generates a process tree for an under-specified region, i.e., a prime component. First, the algorithm randomly selects a pivot node v (Line 1) and calculates its neighborhood graph G_v (Line 2). The parameter $\lambda \in \mathbb{N}^+$ defines when the neighborhood graph is small enough to generate a process tree. If the neighborhood graph is too large, the algorithm calls *synPrime* again, and everything is repeated until the graph is processable. Figure 10 shows the reduction of a neighborhood graph. If our approach selects n^2 as the pivot element, it then builds the smaller graph on the right hand side.


Fig. 10. Reduction of a Neighborhood Graph with the Pivot n^2.

If the neighborhood graph is small enough ($N^{(\lambda)} = \emptyset$), Algorithm 3 randomly selects a tree pattern for it (Step 4). A tree pattern is a process tree for the neighborhood graph. The neighborhood graph in Figure 7(a) contains 5 nodes and 4 edges. For a graph with five nodes thousands of process trees are possible. For the graph in Figure 7(a) 53 trees are possible, given the constraints. For most of these 53 process trees, there is another tree with a lower overall processing time, for any processing times of the tasks. If we exclude these dominated trees, three trees remain. Figures 7(b) and (c) show two of them, randomly selected. Section 5 shows and explains all tree patterns for $\lambda \in [1,5]$. Figure 7(d) shows a process tree fulfilling the constraints in Figure 7(a), but the processing time of the tree in Figure 7(b) always is shorter.

For each ORG we have started out with, we calculate κ different process trees. The resulting trees differ depending on the probabilistic choices in Algorithm 3 (Line 1) and (Line 4). We select the best process tree found according to quality criteria, e.g., the processing time. We calculate a quality value of each tree as follows. The average processing time for each node in a process tree $PT(\mathcal{V}, \mathcal{E})$ is

calculated recursively with function $fit\colon \mathcal{V} \to \mathbb{R}$.

$$fit(n) := \begin{cases} runtime(n) & \text{if } type(n) = task \\ \max_{c \in child_n} fit(c) & \text{if } type(n) = \text{AND} \\ \sum_{c \in child_n} fit(c) & \text{if } type(n) = \text{SEQ} \\ \max_{c \in child_n} fit(c) & \text{if } type(n) = \text{XOR} \end{cases}$$

$type\colon \mathcal{V} \to \{task, \text{AND}, \text{SEQ}, \text{XOR}\}$ is a function to determine the type of the tree node. $child_n := \{c \mid (n, c) \in \mathcal{E}\}$ is the set of nodes in the process tree with parent node n. The fitness of a process tree $fit(PT)$ is the fitness of its root node. The algorithm returns the process tree with the highest fitness value.
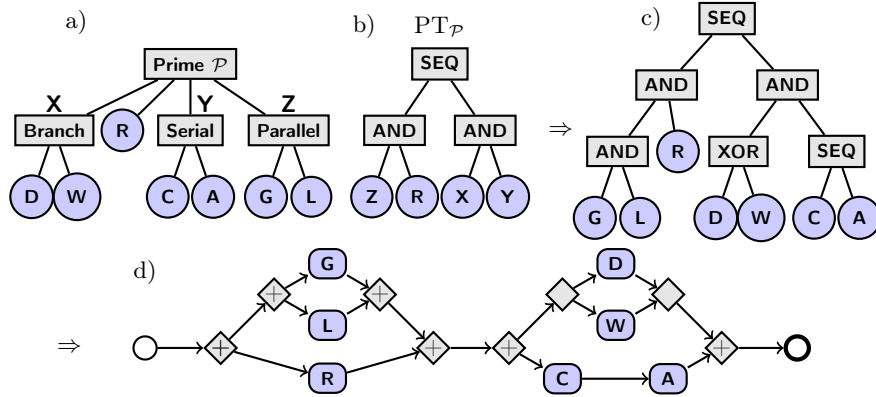

Fig. 11. The MDT for the ORG of Figure 4 (a), the PT for the prime node $\text{PT}_{\mathcal{P}}$ (b), the resulting process tree (c) and the process model in BPMN notation (d).

5 Tree Patterns

In this section we explain the tree patterns for a neighborhood graph G_v. For all possible runtimes of nodes in the neighborhood graph at least one of the tree patterns is optimal according to the fitness function fit. To minimize the fitness value we are trying to set tasks in parallel.

$\lambda = 2$ and $\lambda = 3$

For $\lambda = 2$ only one process tree is possible that fulfills the requirements. Figure 12(a) shows the ORG for $\lambda = 2$ and the only possible process tree. For $\lambda = 3$ it is only possible to execute n^0 and n^2 in parallel. This parallel arrangement leads to one possible process tree, Figure 12(b).

**Fig. 12.** Tree patterns for $\lambda = 2$ (a) and $\lambda = 3$ (b).

Fig. 13. The two Tree Patterns for $\lambda = 4$

$\lambda = 4$

For $\lambda = 4$ we can set in parallel three pairs of nodes (n^0, n^2), (n^0, n^3), (n^1, n^3). If we set (n^0, n^2) in parallel, it allows to set (n^1, n^3) in parallel and leads to the pattern in Figure 13(a). The dotted line shows the dependency added according to the tree pattern. If we set (n^0, n^3) in parallel, we have to add the two dependencies in Figure 13(b). This leads to the second tree pattern.

$\lambda = 5$

For $\lambda = 5$ we can set five pairs of nodes in parallel (n^0, n^2), (n^0, n^3), (n^0, n^4), (n^1, n^3), (n^1, n^4), and (n^2, n^4). If we set (n^0, n^2) in parallel, then we additionally can set in parallel either (n^1, n^3) or (n^1, n^4). If we choose (n^1, n^3) we can add n^4 to the parallel (n^0, n^2) resulting in the process tree of Figure 14(a). If we choose (n^1, n^4) no other parallelism is possible resulting in the process tree of Figure 14(b).

Fig. 14. The three Tree Patterns for $\lambda = 4$

6 Evaluation

Our evaluation uses 21 process models from a German car manufacturer that specify the testing and commissioning of middle-class vehicles. Professional process developers have designed these process models. The tasks to be executed depend on the components built into the vehicle to be tested. In cooperation with those domain experts we have built the specification for the 21 process models, i. e., the ordering relationship graphs, automatically using a knowledge base. See [23]. The process models contain up to 185 tasks and over 3000 dependencies. We choose $\lambda = 5$ for our evaluation. Figure 15(a) shows the full ordering relationship graph for one process model. Figure 15(b) shows a prime element that the graph contains.

Table 2 shows the results for commissioning process models A, B, and C. We have chosen A, B, and C because they are representative for the whole set, ranging from a relative small one (C) to one of the largest (B). For a summary of all models see Table 3. The first row in Table 2 lists the number of tasks that belong to the process model as well as the ordering relationship graph. The second row in Table 2 shows the processing time measured for the process model created by hand. Table 2 then lists the expected processing time of the process (PT) and the time our approach needs to generate the respective model (computation time CT) for 10 to 100,000 iterations. In all cases, the algorithm has been able to generate a process model in less than $100\,\mathrm{ms}$ that outperforms the reference process model. After 100,000 iterations (in less than 1.5 minutes) it could find process models
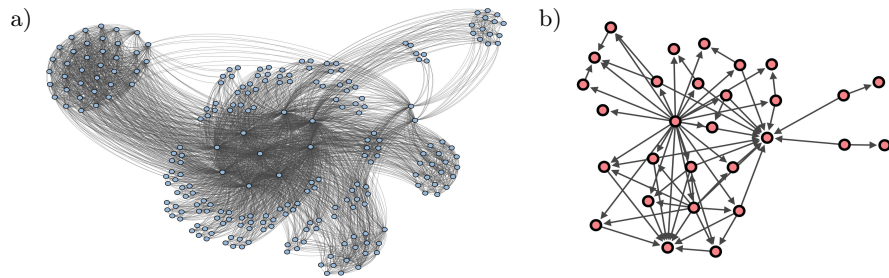
Fig. 15. The ORG for a Commissioning Process Model (a) and a Nontrivial Prime Graph Contained (b) in it.

Table 2. Computation time (CT) and processing time (PT) of our approach

	Process A		Process B		Process C	
No. of Tasks	171		185		116	
Ref. process time	171 780ms		169 606ms		148 014ms	
	CT in ms	PT in ms	CT in ms	PT in ms	CT in ms	PT in ms
10 Iterations	35	188,420	34	227,260	**32**	**132,998**
50 Iterations	**69**	**127,687**	**71**	**131,121**	66	103,234
100 Iterations	113	127,687	113	131,121	104	103,234
1 000 Iterations	823	127 687	964	116,155	788	97,264
10,000 Iterations	8,207	112,918	8,298	113,874	7,817	71,513
100,000 Iterations	78,409	112,624	86,594	106,216	77,335	65,892
PT reduction	34.437 %		37.375 %		50.456 %	

with processing times 34%, 37%, and 50% lower than their manually generated counterparts.

For all 21 process models, Table 3 shows the minimum, maximum, and the quartile for 7 values of the evaluation. The process models contain between 98 and 185 tasks, and need up to 178s to perform. Our approach requires $\approx 30s$ and $\approx 37\,000$ iterations on average to generate the best result found. For all instances our approach has identified a solution that is better than the manually generated reference point in less than 100 ms. Moreover, our approach needs less than 3 iterations to do so in most cases. On average, our approach nearly halves the processing time of the commissioning process models (47.47%) compared to the reference points.

Table 3. The Minimum, Maximum, and the Quartile for the Evaluation of 21 Commissioning Process Models

	Minimum $Q_{0.00}$	$Q_{0.25}$	Median $Q_{0.5}$	$Q_{0.75}$	Maximum $Q_{1.00}$
Nr. of Tasks	98	123	133	146	185
Ref. Process Time	144.232s	151.623s	157.513s	166.138s	178.606s
best found pt	64.643s	72.637s	84.529s	93.487s	108.496s
iterations (it)	5 090	15 523	37 733	76 035	94 271
calculation time (ct)	4.641s	12.356s	30.284s	63.392s	77.480s
pt reduction	33.39%	40.62%	47.54%	53.62%	58.03%

7 Related Work

[29] synthesizes a process model directly from its specification. The specifications are in propols [29], a high-level temporal constraint specification language. The specifications are transformed into finite state machines and then integrated into one machine. Next, each accepting path is generated from the state machine. An algorithm similar to the α-algorithm [3] for process mining is applied to synthesize a process model from its set of paths. [29] can only be applied if the specification, i. e., the number of state machines, is small (≈ 6). To this end, [29] divides the specification into small groups, synthesizes a process fragment for each group and manually combines the fragments. For our use case, this approach would require over a hundred state machines for each commissioning process model, and the manual combination would not be feasible. [8] has specifications with ltl as starting point. It generates a pseudo model from the specification. The pseudo model lists all paths that fulfill the ltl formula. [8] generates an ordering relation graph from the set of paths and uses it to synthesize a process tree. For our use case the generation of all paths would not be feasible. This is because, the number of paths grows exponentially with the size of the specification. Even for the smallest process model we have evaluated the calculation of all paths has not been possible.

Process discovery means finding a process model that can reproduce the behavior given in a log [2]. [16] rediscovers a process model in the process-tree notation. It generates a graph (directly-follows graph) from the log and tries to find different kinds of cuts in the graph. Each kind of cut refers to a control structure in the process tree (seq, and, xor, loop). The cuts partition the graph and allow to hierarchically find a process tree for the log. In contrast to an org, a directly-follows graph is not transitive, and if two nodes are in parallel they share a two-way edge (no edge in the org). It is not possible to find a cut for a prime component, thus the approach of [16] does not help in case the specification is under-specified. Put differently, the problem statement in [16] is different from ours; the neighborhood graph of the complete log of a process tree never contains a prime component. For an incomplete log, a prime component

can occur. [17] propose to use probabilistic activity relations in the case of an incomplete log. The cut with the highest probability is chosen. This means that their algorithm generalizes from the incomplete log and assumes relationships that are not present. An ORG is an upper bound of the possible behavior. Assuming an additional relation would result in a violation of a constraint.

An approach different from generating the process model from scratch is to extract information from process models already specified and to create a similar process. [9] uses a CBR-based method for the composition of workflows. The search is based on keywords that are annotations of the workflows. [13] guides the process designer with suggestions on how to complete data-oriented visualization models. The suggestions are generated from paths of existing visualization process models stored in a process repository. [13] does not allow building a process model with an AND-Split and therefore is not sufficient in our case. [15] predicts which activity pattern (generic process fragment) will follow the partly modeled process. The paths of existing process models are extracted and analyzed with association rule mining. [13][15] extend an existing process model, while our approach generates one from a declarative specification. [9] requires annotations of the existing process models. None of the approaches mentioned optimize the runtime or consider constraints.

[25] transforms an unstructured model without cycles into a behaviorally equivalent structured process model. 'structured' means that for each Split-Gateway there is a corresponding Join-Gateway. Structured processes allow an effective verification [22] and are easy to understand [27]. [25] determines relationships between the tasks of a process model and generates an ORG using these relationships. Next, [25] decomposes the ORG into a Modular Decomposition Tree. In contrast to our approach, [25] generates the ORG from the behavior of an existing process model and not from a set of compliance rules. The behavior is definite, the result therefore is a unique process model. In our approach in turn, the behavior is under-specified, and several process models are possible.

AI planning is the task of defining a set of actions that achieve a specified aim [12]. In a nutshell, it is the search for an applicable plan in the solution space. [28] uses a genetic algorithm to find a manufacturing plan. Some approaches that synthesize business processes are discussed next: [19] uses an AI planning approach to synthesize service compositions. Without calling it AI planning, [5] uses a similar approach for configuration-based workflow composition. [6] introduces a planning algorithm to compose data workflows. None of these studies focuses on optimizing the runtime of the process or considers requirements similar to ours. These approaches are not applicable to our problem statement.

In contrast to imperative process models, declarative workflows allow for any behavior fulfilling the declarative specification [21]. Thus, declarative workflows provide maximum flexibility not limited by a process model. In comparison, [29], [8] and our approach generate an imperative process model from the declarative specification. The enactment of declarative workflows is not trivial [24], and tool support by major vendors is missing. To our knowledge, there is no tool that executes declarative process models comparable to the commissioning of vehicles.

8   Conclusions

We have proposed a novel approach to generate a process model for a specific context automatically, given a set of constraints. We study the restricted case that there are not any repetitions of a task, as is the case in commissioning and elsewhere, e.g., manufacturing. We use a probabilistic search to find a good process model according to quality criteria that fulfills the constraints. Our approach can handle complex real-world specifications consisting of several hundred constraints and more than one hundred tasks. The process models generated with our scheme are superior (nearly twice as fast) to ones designed by professional process designers.

References

[1]   W. M. P. v. d. Aalst and M. Pesic. "DecSerFlow: Towards a Truly Declarative Service Flow Language". In: *Web Services and Formal Methods*. 2006.

[2]   W. M. P. van der Aalst. *Process Mining: Discovery, Conformance and Enhancement of Business Processes*. 2011.

[3]   W. M. P. van der Aalst, T. Weijters, and L. Maruster. "Workflow mining: discovering process models from event logs". In: *IEEE Transactions on Knowledge and Data Engineering* (2004).

[4]   W. M. P. van der Aalst et al. "Workflow Patterns". In: *Distributed and Parallel Databases* (2003).

[5]   P. Albert, L. Henocque, and M. Kleiner. "Configuration based workflow composition". In: *IEEE International Conference on Web Services*. 2005.

[6]   J. L. Ambite and D. Kapoor. "Automatically Composing Data Workflows with Relational Descriptions and Shim Services". In: *The Semantic Web*. 2007.

[7]   A. Awad, G. Decker, and M. Weske. "Efficient Compliance Checking Using BPMN-Q and Temporal Logic". In: *Business Process Management*. 2008.

[8]   A. Awad et al. "An Iterative Approach for Business Process Template Synthesis from Compliance Rules". In: *Advanced Information Systems Engineering*. 2011.

[9]   E. Chinthaka et al. "CBR Based Workflow Composition Assistant". In: *IEEE World Conference on Services*. 2009.

[10]   M. B. Dwyer, G. S. Avrunin, and J. C. Corbett. "Property Specification Patterns for Finite-state Verification". In: *Proceedings of the Second Workshop on Formal Methods in Software Practice*. 1998.

[11]   M. Dwyer, G. Avrunin, and J. Corbett. "Patterns in property specifications for finite-state verification". In: *Proceedings of the 1999 International Conference on Software Engineering, 1999*. 1999.

[12]   J. Hendler, A. Tate, and M. Drummond. *AI Planning: Systems and Techniques*. Tech. rep. University of Maryland at College Park, 1990.

[13]   D. Koop et al. "VisComplete: Automating Suggestions for Visualization Pipelines". In: *IEEE Transactions on Visualization and Computer Graphics* (2008).

[14] O. Kopp et al. "The Difference Between Graph-Based and Block-Structured Business Process Modelling Languages". In: *Enterprise Modelling and Information Systems Architecture* (2009).

[15] J. M. Lau et al. "Discovery and Analysis of Activity Pattern Cooccurrences in Business Process Models". In: *Int'l Conf. on Enterprise Information Systems*. 2009.

[16] S. J. J. Leemans, D. Fahland, and W. M. P. van der Aalst. "Discovering Block-Structured Process Models from Event Logs - A Constructive Approach". In: *Application and Theory of Petri Nets and Concurrency*. 2013.

[17] S. J. J. Leemans, D. Fahland, and W. M. P. van der Aalst. "Discovering Block-Structured Process Models from Incomplete Event Logs". In: *Application and Theory of Petri Nets and Concurrency*. 2014.

[18] L. T. Ly et al. "SeaFlows Toolset – Compliance Verification Made Easy for Process-Aware Information Systems". In: *Information Systems Evolution*. 2011.

[19] M. Matskin and J. Rao. "Value-Added Web Services Composition Using Automatic Program Synthesis". In: *Web Services, E-Business, and the Semantic Web*. 2002.

[20] R. M. McConnell and F. de Montgolfier. "Linear-time modular decomposition of directed graphs". In: *Discrete Applied Mathematics* (2005).

[21] M. Montali et al. "Declarative Specification and Verification of Service Choreographiess". In: *ACM Trans. Web* (2010).

[22] R. Mrasek, J. Mülle, and K. Böhm. "A new verification technique for large processes based on identification of relevant tasks". In: *Information Systems* (2014).

[23] R. Mrasek et al. "User-Friendly Property Specification and Process Verification - a Case Study with Vehicle-Commissioning Processes". In: *Business Process Management*. 2014.

[24] M. Pešić, D. Bošnački, and W. M. P. van der Aalst. "Enacting Declarative Languages Using LTL: Avoiding Errors and Improving Performance". In: *Model Checking Software*. 2010.

[25] A. Polyvyanyy, L. García-Bañuelos, and M. Dumas. "Structuring acyclic process models". In: *Information Systems* (2012).

[26] E. Ramezani, D. Fahland, and W. M. P. van der Aalst. "Where Did I Misbehave? Diagnostic Information in Compliance Checking". In: *Business Process Management*. 2012.

[27] H. Reijers and J. Mendling. "A Study Into the Factors That Influence the Understandability of Business Process Models". In: *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans* (2011).

[28] J. Váncza and A. Márkus. "Genetic algorithms in process planning". In: *Computers in Industry* (1991).

[29] J. Yu et al. "Synthesizing Service Composition Models on the Basis of Temporal Business Rules". In: *Journal of Computer Science and Technology* (2008).