

Compliance by Design for Artifact-Centric Business Processes

Niels Lohmann

Universität Rostock, Institut für Informatik, 18051 Rostock, Germany
niels.lohmann@uni-rostock.de

Abstract. Compliance to legal regulations, internal policies, or best practices is becoming a more and more important aspect in business processes management. Compliance requirements are usually formulated in a set of rules that can be checked during or after the execution of the business process, called *compliance by detection*. If noncompliant behavior is detected, the business process needs to be redesigned. Alternatively, the rules can be already taken into account while modeling the business process to result in a business process that is *compliant by design*. This technique has the advantage that a subsequent verification of compliance is not required.

This paper focuses on compliance by design and employs an *artifact-centric* approach. In this school of thought, business processes are not described as a sequence of tasks to be performed (i. e., imperatively), but from the point of view of the artifacts that are manipulated during the process (i. e., declaratively). We extend the artifact-centric approach to model compliance rules and show how compliant business processes can be synthesized automatically.

1 Introduction

Business processes are the main asset of companies as they describe their value chain and fundamentally define the “way, businesses are done”. Beside fundamental correctness criteria such as soundness (i. e., every started case is eventually finished successfully), also nonfunctional requirements have to be met. Such requirements are often collected under the umbrella term *compliance*. They include legal regulations such as the often cited Sarbanes-Oxley Act to fight accounting frauds, internal policies to streamline the in-house processes, or industrial best practices to reduce complexity and costs as well as to facilitate collaborations. Finally, compliance can be seen as a means to validate business processes [5].

Compliance requirements are usually defined without a concrete business process in mind, for instance in legal texts such as “*The Commission shall [...] certify [...] that the signing officers have designed such internal controls to ensure that material information [...] is made known [...] during the period in which the periodic reports are being prepared*”.¹ Such informal descriptions then must

¹ Excerpt of Title 15 of the United States Code, § 7241(a)(4)(B).

be translated by domain experts into precise rules that unambiguously capture the essence of the requirement in a shape that it can be checked in concrete business process. The example above could yield a rule such as “*Information on financial reports must be sent to the press team by the signing officers at most two weeks after signing*”. The rules clarify who has to take which actions and when. That is, they specifically deal with the execution order of actions of the business process or the reachability of data values. This formalization of domain-specific knowledge is far from trivial and out of scope of this paper. In the remainder, we assume—similar to other approaches [18]—that such rules are already present.

Compliance rules are often *declarative* and describe what should be achieved rather than how to achieve it. Temporal logics such as CTL [4], LTL [21] or PLTL [20] are common ways to formalize such declarative rules. To make these logics approachable for nonexperts, also graphical notations have been proposed [1,2]. Given such rules, compliance of a business process model can be verified using model checking techniques [6]. These checks can be classified as *compliance by detection*, also called after the fact or retrospective checking [25]. Their main goal is to provide a rigorous proof of compliance. In case of noncompliance, diagnosis information may help to fix the business process toward compliance. This step can be very complicated, because the rules may affect various parts and agents of the business process (e.g., financial staff and the press team). Furthermore, the declarative nature of the rules does not provide recipes on how to fix the business process. To meet the previous example rule, an action “send information to press team” needs to be added to the process and must be executed at most two weeks after the execution of an action “sign financial report”. Compliance can be eventually reached after iteratively adjusting the business process model, cf. Fig. 1(a).

An alternative approach takes a business process model and the compliance rules as input and automatically generates a business process model that is *compliant by design* [25], cf. Fig. 1(b). This has several advantages: First, a subsequent proof and potential corrections are not required. This may speed up the modeling process. Second, the approach is flexible as the generation can be repeated when rules are added, removed, or changed. Third, the approach is complete in the sense that an unsuccessful model generation can be interpreted

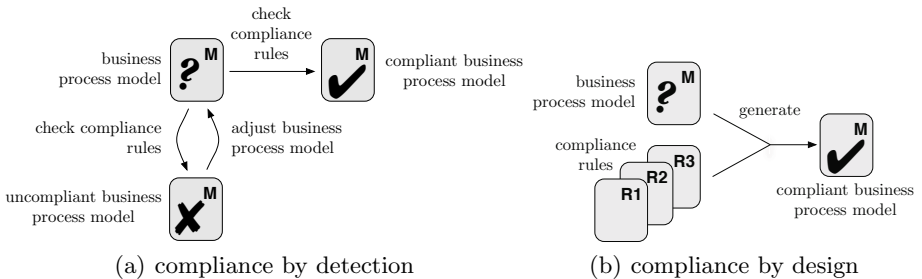


Fig. 1. Approaches to achieve compliance

as “the business process cannot be *made* compliant” rather than “the current model is not compliant”. Fourth, compliance is not only detected, but actually enforced. That is, noncompliant behavior becomes technically impossible.

This paper investigates the latter compliance-by-design approach. We employ a recent framework for *artifact-centric* business processes [17]. In this framework, a business process is specified by a description of the life cycles of its data objects (artifacts). From this declarative specification, which also specifies agents and locations of artifacts, a sound, operational, and interorganizational business process can be automatically generated.

Contribution. This paper makes two contributions: First, we extend the artifact-centric framework [17] to model a large family of compliance rules. Second, we use existing tools and techniques to achieve not only soundness, but also compliance by design. We also sketch the diagnosis of noncompliant models.

Organization. The next section introduces a small example we use throughout the paper to exemplify our approach and later extensions. Section 3 sets the stage for our later contributions. There, we introduce artifact-centric business processes and correctness by design. In Sect. 4, we demonstrate how a large family of compliance rules can be expressed with in our approach. Section 5 presents how compliance by design can be achieved. We also discuss the diagnosis of unrealizable compliance rules. Section 6 brings our approach in the context of related work, before Sect. 7 concludes the paper.

2 Running Example: Insurance Claim Handling

We use a simple insurance claim handling process (based on [24]) as running example for this paper. In this process, a customer submits a claim to an insurer who then prepares a fraud detection check offered by an external service. Based on the result of this check, the claim is either (1) assessed and the settlement estimated, (2) detected fraudulent and reported, or (3) deemed incomplete. In the last case, further information are requested from the customer before the claim is resubmitted to the fraud detection service. In this situation, the customer can alternatively decide to withdraw the claim. On successful assessment, a settlement case is processed by a financial clerk. The claim is settlement paid in several rates or all at once. A single complete payment further requires an authorization of the controlling officer. When the settlement is finally paid, the claim is archived.

One way to model this process is to explicitly order the actions to be taken and to give an operational business process model. An alternative to this *verb-centric* approach offers an artifact-centric framework which starts by identifying what is acted on (*noun-centric*) and to derive a business process from the life cycles of the involved artifacts. We shall discuss artifact-centric business processes in the next section.

3 Artifact-Centric Business Processes

In this section, we shall introduce artifact-centric business processes [17]. We first give an informal overview of all concepts involved. Then we present a Petri net formalization and discuss it in on the basis of the running example. Admittedly, this section takes a large part of this paper, but is required to discuss the contributions to compliance.

3.1 Informal Overview

Artifact-centric modeling promotes the data objects of a business process (called *artifacts*) and their life cycles to first-class citizens. In the running example, we consider two artifacts: an insurance claim file and a settlement case. Each life cycle describes how the state of an artifact may evolve over time. The actions that change states are executed by *agents*, for instance the customer, the insurer, the financial clerk, and the controlling officer. As multiple agents can participate in a business process, the artifact-centric approach is particularly suited to model interorganizational business processes.

Each artifact has at least one *final state* which models a successful processing of the artifact, for instance “claim archived”, or “settlement paid”. Artifact-centric business processes are inherently *declarative*: the control flow of the business process is not explicitly modeled, but follows from the life cycles of the artifacts. That is, any execution of actions that brings all artifacts to a final state can be seen as sound. However, not every sound execution makes sense. For instance, semantically ordered actions of different and independently modeled artifacts (e. g., “assess claim” and “create settlement”) may be executed in any order. In addition, not every combination of final states may be desirable, for instance “claim withdrawn” in combination with “settlement paid”. Therefore, the executions have to be constrained using *policies* and *goal states*. A policy is a way of expressing constraints between artifacts. For instance, a policy may constrain the order of state changes in different artifacts (e. g., always executing “assess claim” before “pay settlement”). Finally, goal states restrict final states by reducing those combinations of artifacts’ final states that should be considered successful.

In recent work [17], we presented an approach that takes artifacts, policies, and goal states as input and automatically synthesizes an interorganizational business process. This business process has two important properties: First, it is *operational*. It explicitly models which agent may perform which action in which state. In addition to the control flow, we can also derive the data flow and even the message flow, because artifacts may be sent between agents. Operational models can be translated into languages such as BPMN [19] and can be easily refined toward execution. Second, the business process is *weakly terminating*. Weak termination is a correctness criterion that ensures that a goal state is always reachable from every reachable state. Any actions that would lead to deadlocks or livelocks are removed. This means that the approach is *correct by design*. To summarize, the artifact-centric approach allows to model




Fig. 2. Artifact-centric business processes in a nutshell

artifacts and to restrict their manipulation by additional domain knowledge such as policies and goal states. From this declarative model we can then automatically generate a weakly terminating operational model. Figure 2 illustrates the overall approach.

3.2 Formalization

We model artifact-centric business processes with Petri nets [23]. Petri nets combine a simple graphical representation with a rigorous mathematical foundation. They can naturally express locality of actions in distributed systems. This allows us to model the life cycles of several artifacts independently, yielding a more compact model compared to explicit state machines.

Definition 1 (Petri net). A Petri net $N = [P, T, F, m_0]$ consists of two finite and disjoint sets P of places and T of transitions, a flow relation $F \subseteq (P \times T) \cup (T \times P)$, and an initial marking m_0 . A marking $m : P \rightarrow \mathbb{N}$ represents a state of the Petri net and is visualized as a distribution of tokens on the places. Transition t is enabled in marking m iff, for all $[p, t] \in F$, $m(p) > 0$. An enabled transition t can fire, transforming m into the new state m' with $m'(p) = m(p) - W([p, t]) + W([t, p])$ where $W([x, y]) = 1$ if $[x, y] \in F$, and $W([x, y]) = 0$, otherwise.

A Petri net shall describe the life cycle of an artifact. For our purposes, we have to extend this model with several concepts: Each transition is associated with an action from a fixed set $\mathcal{L} = \mathcal{L}_c \cup \mathcal{L}_u$ of *action labels*. This set is partitioned into a set \mathcal{L}_c of *controllable actions* that are executed by agents and a set \mathcal{L}_u of *uncontrollable actions* that are not controllable by any agent, but are under the influence of the environment. Such uncontrollable actions are suitable to model choices that are external to the business process model, such as the outcome of a service call, for instance to a fraud detection agency.

Definition 2 (Artifact [17]). An artifact $A = [N, \ell, \Omega]$ consists of (1) a Petri net $N = [P, T, F, m_0]$, (2) a transition labeling $\ell : T \rightarrow \mathcal{L}$ associating actions with Petri net transitions, and (3) a set Ω of final markings of N representing endpoints in the life cycle of the artifact.

Running example (cont.). Figure 3 depicts the claim and the settlement artifacts. Each transition is labeled by the agent that executes it (*insurer*, *customer*,

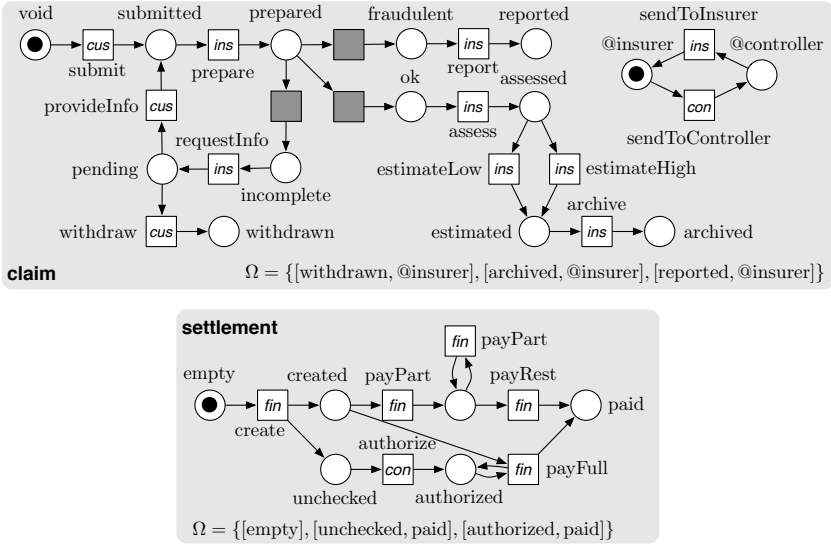


Fig. 3. Artifacts of the running example

controller, and financial clerk) or is shaded gray in case of uncontrollable actions. Two additional places (“@insurer” and “@controller”) model the *location* of the insurance claim file. The artifact-centric approach allows to distinguish physical objects (e. g., documents or goods) that need to be transferred between agents to execute certain actions and virtual objects (e. g., data bases or electronic documents). By taking the shape of the artifacts and their location into account, we can later derive explicit message transfer among the agents. In our example, we assume that after submitting the claim, a physical file is created by the insurer which can be sent to the controlling officer by executing the respective action “send to controller”. We further assume that the settlement case is a data base entry that can be remotely accessed by the insurer, the financial clerk, and the controlling officer.

We can now model each artifact of our business process independently with a Petri net model. Together, the artifacts implicitly specify a process of actions that may be performed by the agents or the environment according to the life cycles of the artifacts. This global model is formalized as the union of the artifact models, which is again an artifact.

Definition 3 (Artifact union [17]). Let A_1, \dots, A_n be artifacts with pairwise disjoint Petri nets N_1, \dots, N_n . Define the artifact union $\bigcup_{i=1}^n A_i = [N, \ell, \Omega]$ to be the artifact consisting of (1) $N = [\bigcup_{i=1}^n P_i, \bigcup_{i=1}^n T_i, \bigcup_{i=1}^n F_i, m_{0_1} \oplus \dots \oplus m_{0_n}]$, (2) $\ell(t) = \ell_i(t)$ iff $t \in T_i$ ($i \in \{1, \dots, n\}$), and (3) $\Omega = \{m_1 \oplus \dots \oplus m_n \mid m_i \in \Omega_i \wedge 1 \leq i \leq n\}$. Thereby, \oplus denotes the composition of markings: $(m_1 \oplus \dots \oplus m_n)(p) = m_i(p)$ iff $p \in P_i$.

The previous definition is of rather technical nature. The only noteworthy property is that the set of final markings of the union consists of all combinations of final markings of the respective artifacts. Conceptually, the union of the artifacts has several downsides as we discussed in Sect. 3.1. First, it may contain sequences of actions that reach deadlocks or livelocks. That is, the model is not necessarily weakly terminating. Second, the artifacts may evolve independently which may result in implausible execution orders or undesired final states. As stated earlier, the latter problems can be ruled out by defining *policies*, which restrict interartifact behavior, and *goal states*, which restrict the final states of the union. Before discussing this, we shall first cope with the first problem.

The transitions of the artifacts are labeled with actions that can be executed by agents. Hence, the agents have control about the evolution of the overall business process. To avoid undesired situations such as deadlocks or livelocks, their behavior needs to be coordinated. That is, it must be constrained such that every execution can be continued to a final state. This coordination can be seen as a *controller synthesis* problem [22]: given an artifact A , we are interested in a controller C (in fact, also modeled as an artifact) such that their interplay is weakly terminating. The interplay of two artifacts is formalized by their *composition*, cf. Fig. 4.

Definition 4 (Artifact composition [17]). Let A_1 and A_2 be artifacts. Define their shared labels as $S = \{l \mid \exists t_1 \in T_1, \exists t_2 \in T_2 : \ell(t_1) = \ell(t_2) = l\}$. The composition of A_1 and A_2 is the artifact $A_1 \oplus A_2 = [N, \ell, \Omega]$ consisting of:

- $N = [P, T, F, m_{0_1} \oplus m_{0_2}]$ with
 - $P = P_1 \cup P_2$,
 - $T = (T_1 \cup T_2 \cup \{[t_1, t_2] \in T_1 \times T_2 \mid \ell(t_1) = \ell(t_2)\}) \setminus (\{t \in T_1 \mid \ell_1(t) \in S\} \cup \{t \in T_2 \mid \ell_2(t) \in S\})$,
 - $F = ((F_1 \cup F_2) \cap ((P \times T) \cup (T \times P))) \cup \{[[t_1, t_2], p] \mid [t_1, p] \in F_1 \vee [t_2, p] \in F_2\} \cup \{[p, [t_1, t_2]] \mid [p, t_1] \in F_1 \vee [p, t_2] \in F_2\}$,
- for all $t \in T \cap T_1$: $\ell(t) = \ell_1(t)$, for all $t \in T \cap T_2$: $\ell(t) = \ell_2(t)$, and for all $[t_1, t_2] \in T \cap (T_1 \times T_2)$: $\ell([t_1, t_2]) = \ell_1(t_1)$, and
- $\Omega = \{m_1 \oplus m_2 \mid m_1 \in \Omega_1 \wedge m_2 \in \Omega_2\}$.

The composition $A_1 \oplus A_2$ is complete if for all $t \in T_i$ holds: if $\ell_i(t) \notin S$, then $\ell_i(t) \in \mathcal{L}_u$ ($i \in \{1, 2\}$).




Fig. 4. Composition of two artifacts

Given an artifact A , we call another artifact C a *controller* for A iff (1) their composition $A \oplus C$ is complete and (2) for each reachable markings of the composition, a final marking $m \oplus m'$ of $A \oplus C$ is reachable. The existence of controllers (also called *controllability* [27]) is a fundamental correctness criterion for communicating systems such as services. It can be decided constructively [27]: If a controller for an artifact exists, it can be constructed automatically [16]. Note that the requirement of a complete composition makes sure that the controller does not constrain the execution of uncontrollable actions.

Finally, we can define goal states and policies to constrain the interartifact behavior. *Goal states* are a set of markings of the artifact union and are used as final markings during controller synthesis. To model interdependencies between artifacts, we employ *policies*. We also model policies with artifacts (similar to behavioral constraints [15]); that is, labeled Petri nets with a set of final markings. These artifacts have no counterpart in reality and are only used to model dependencies between actions of different artifacts. The application of policies then boils down to the composition of the artifacts with these policies.

Running example (cont.). To rule out implausible behavior, we further define the following policies to constrain interartifact behavior and the location’s impact on actions:

- P1.** The claim may be archived only if it resides at the insurer and the settlement is paid.
- P2.** A settlement may only be created after the claim has been estimated.
- P3.** To authorize the complete payment of the settlement, the claim artifact must be at hand to the controlling officer.
- P4.** The claim artifact may only be sent to the controller if it has been estimated and the settlement has not been checked.
- P5.** The claim artifact may only be sent back to the insurer if the settlement has been authorized.

The policies address different aspects of the artifacts such as location (P1 and P3), execution order (P2), or data constraints (P4 and P5). The modeling of the policies as artifacts is straightforward and depicted in Fig. 5. Note that in policy P2, we use an unlabeled place to express the causality between the transitions.




Fig. 5. Policies for the running example

would be spurious as the respective choice could not be disabled in reality. Therefore, compliance rules must not restrict the behavior of the artifacts, but only restrict the final states of the model. This may classify behavior as undesired (viz. noncompliant), but this behavior remains reachable. Only if this behavior can be circumvented by the controller synthesis, we faithfully found a compliant business process which can be actually implemented. We formalize this nonrestricting nature as *monitor property* [15,27]. Intuitively, this property requires that in every reachable marking of an artifact, it holds that for each action label of that artifact a transition with that label is activated. This rules out situations in which the firing of a transition in a composition is inhibited by a compliance rule.

4.2 Expressiveness of Compliance Rules

Conceptually, we model compliance rules by artifacts with the monitor property. Again, adding a compliance rule to an artifact-centric model boils down to composition. The monitor property ensures that the compliance rule's transitions are synchronized with the other artifacts, but without restricting (i. e., disabling) actions. That is, the life cycle of a compliance rule model evolves together with the artifacts' life cycles, but may affect the final states of the composed model.

In a finite-state composition of artifacts, the set of runs reaching a final state forms a regular language. The terminating runs of a compliance rule (i. e., sequences of transitions that reach a final marking) describe compliant runs. This set again forms a regular language. In the composition of the artifacts and the compliance rules, these regular languages are synchronized—viz. intersected—yielding a subset of terminating runs. Regular languages allow to express a variety of relevant scenarios. In fact, we can express all patterns listed by Dwyer et al. [8], including:

- enforcement and existence of actions (e. g., “*Every compliant run must contain an action ‘archive claim’.*”),
- absence/exclusion of actions (e. g., “*The action ‘withdraw claim’ must not be executed.*”),
- ordering (precedence and response) of actions (e. g., “*The action ‘create settlement’ must be executed after ‘submit claim’, but before ‘archive claim’.*”), and
- numbering constraints/bounded existence of actions (e. g., “*The action “partially pay settlement” must not be executed more than three times.*”).

The explicit model of data states of the artifacts further allows to express rules concerning data flow, such as:

- enforcement/exclusion of data states (e. g., “*The claim’s state ‘fraud reported’ and the settlement’s state ‘paid’ must never coincide.*”), or
- data and control flow concurrence (e. g., “*The action ‘publish review’ may only be executed if the review artifact is in state ‘reviewers blinded’.*”).

On top of that, any combinations are possible, allowing to express complex compliance rules.

The presented approach is, however, not applicable to nonregular languages. For instance, a rule requiring that a compliant run must have an arbitrary large, but equal number of a and b actions or that a and b actions must be properly balanced (Dyck languages) cannot be expressed with a finite-state models. Similarly, rules that affect infinite runs (e.g., certain LTL formulae [21]) cannot be expressed. Infinite runs are predominantly used to reason about reactive systems. A business process, however, is usually designed to eventually reach a final state—this basically is the essence of the soundness property. Therefore, we shall focus on an interpretation of LTL which only considers finite runs, similar to a semantics described by Havelund and Roşu [11]. Just like Awad et al. [3], we also do not consider the X (next state) operator of CTL*, because we typically discuss distributed systems in which states are partially ordered. Finally, we do not use timed Petri nets and hence can make no statements on temporal properties of business processes. However, we can abstract the variation of time by events such as “time passes” or data states such as “expired” as in [10,18].

4.3 Formalization

As mentioned earlier, we again use artifacts (i.e., Petri nets with final markings and action labels) that satisfy the monitor property to model compliance rules. As an example, we consider the following compliance rules for our example insurance claim process:

- R1.** All insurance claims with an estimated high settlement must be authorized.
- R2.** Customers must not be allowed to withdraw insurance claims.
- R3.** Settlements should be paid in at most three parts.

Figure 7 shows the Petri net formalizations of these compliance rules. In rule R1, we exploited the fact that the actions “authorize” and “estimateHigh” are executed at most once. In rule R2 and R3, the monitor property is achieved by allowing “withdraw” and “payPart” to fire in any reachable state. Without

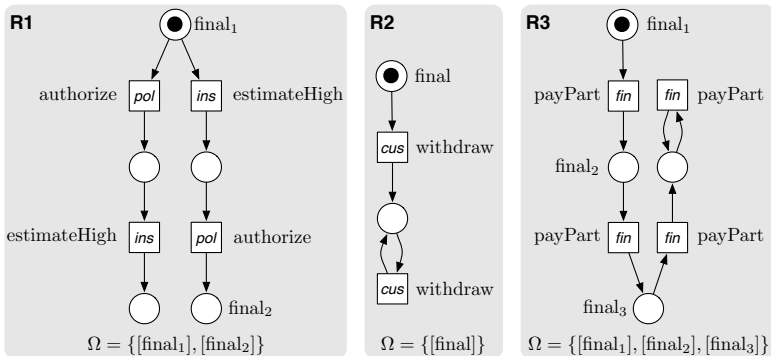


Fig. 7. Compliance rules modeled as Petri nets

restriction of the behavior, the final markings classify executions as compliant or not. For instance, executing “estimateHigh” in rule R1 without eventually executing “authorize” does not reach the final marking [final₂].

4.4 Discussion

We conclude this section by a discussion of the implications of using Petri nets to formalize compliance rules.

- *Single formalism.* We can model artifacts, policies, and compliance rules with the same formalism. Though we do not claim that Petri nets should be used by domain experts to model compliance regulations, using a single formalism still facilitates the modeling and verification process. Furthermore, each rule implicitly models compliant behavior which can be simulated. This is not possible if, for instance, arbitrary LTL formulae are considered.
- *Level of abstraction.* Rules can be expressed using minimal overhead. Each rule contains only those places and transitions that are affected by the rule and plus some additional places to model further causalities. In particular, no place-holder elements (e. g., anonymous activities in BPMN-Q [2]) are required.
- *Independent design.* The rules can be formulated independently of the artifact and policy models. That is, the modeler does not need to be confronted with the composite model. This modular approach is more likely to scale, because the rules can also be validated independently of the other rules.
- *Reusability.* The composition is defined in terms of action labels. Therefore, rules may be reused in different business process models as long as the labels match. This can be enforced using standard naming schemes or ontologies.
- *Runtime monitoring.* The monitor property ensures that the detection of noncompliant behavior is transparent to the process as no behavior is restricted. Therefore, the models of the compliance rules can be also used to check compliance during or after runtime, for instance by inspecting execution logs.
- *Rule generation.* Finally, the structure of the Petri nets modeling compliance rules is very generic. Therefore, it should be possible to automatically generate Petri nets for standard scenarios or to provide templates to which only the names of the constrained actions need to be filled. Also, the monitor property can be automatically enforced.

5 Compliance by Design

This section presents the second contribution of this paper: the construction of business process models that are compliant by design. Beside the construction, we also discuss the diagnosis of noncompliant business process models.

5.1 Constructing Compliant Models

None of the compliance rules discussed in the previous section hold in the example process depicted in Fig. 6. This noncompliance can be detected by standard

model checking tools. They usually provide a counterexample which describes how a noncompliant situation can be reached. For instance, the action sequence “1. submit, 2. prepare, 3. requestInfo, 4. withdraw” is a witness that the process does not comply with rule R2. To satisfy this requirement, the transition “withdraw” can be simply removed. However, implementing the other rules is more complicated, and each modification would require another compliance check.

We propose to *synthesize* a compliant model instead of verifying compliance. By composing the Petri net models of the artifacts (cf. Fig. 3), the policies (cf. Fig. 5), and the compliance rules (cf. Fig. 7) and by taking the goal states into account, we derive a Petri net that models the artifacts’ life cycles that are restricted by the policies and whose final states are constrained by the goal states and the compliance rules. *Compliant behavior is now reduced to weak termination*, and we can apply the same algorithm [27] and tool [16] to synthesize a controller. If such a controller exists, it provides an operational model that specifies the order in which the agents need to perform their actions. This model is *compliant by design* — a subsequent verification is not required. Beside weak termination (and hence, compliance), the synthesis algorithm further guarantees the resulting model is *most permissive* [27]. That is, exactly that behavior has been removed that would violate weak termination. Another important aspect of the approach is its flexibility to add further compliance rules. That is, we do not need to edit the existing model, but we can simply repeat the synthesis for the new rule set.

Running example (cont.). Figure 8 depicts the resulting business process model. It obviously contains no transition labeled with “withdraw”, but the implementation of the other rules yielded a whole different structure of the part modeling the settlement processing. *It is important to stress that the depicted business process model has been synthesized completely automatically using the partner synthesis tool Wendy [16] and the Petri net synthesis tool Petrify [7].* Admittedly, it is a rather complicated model, but any valid implementation of the compliance

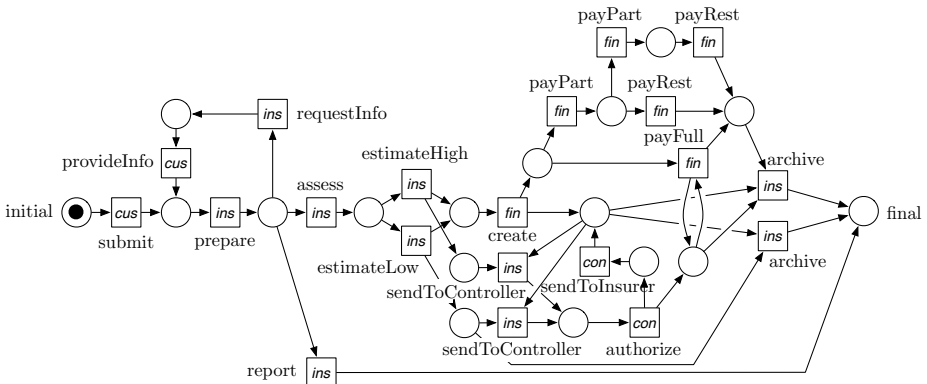


Fig. 8. Operational business process satisfying the compliance rules R1–R3

rules would yield the same behavior or a subset. Though our running example is clearly a toy example, experimental results [16] show that controller synthesis can be effectively applied to models with millions of states.

5.2 Diagnosing Noncompliant Models

So far, we only considered the case that the business process can be constructed such that it satisfies the compliance rules. Then, we can find a controller that guarantees weak termination. In case a compliance rule is not met by the business process, no such controller exists. That is, the business process cannot be made compliant. Intuitively, the intersection between the behavior of the business process and the compliance rule is empty. Just as a controller would be a witness for compliant behavior, such witness does not exist in case of noncompliance.

In previous work [14], we studied uncontrollable models and presented an algorithm that generates diagnosis information. This diagnosis information is presented as a graph that overapproximates the behavior of any controller. As no controller exists that can avoid states that violate weak termination, this graph contains paths to such deadlocks or livelocking states. These paths and the reason they cannot be avoided by a controller serve as a counterexample similar to those of standard model checking. Note that such a counterexample not only describes noncompliance of a concrete business process model, but for a whole family of operational models that can be derived from a set of artifacts.

Running example (cont.). To exemplify this diagnosis information, consider the following compliance rule:

R4. The customer should not be asked for further information more than once.

This rule is not realizable in the business process, because the outcome of the fraud detection service is not under the control of any agent. Furthermore, rule R2 excluded the possibility for the customer to withdraw the claim. Hence, we cannot exclude a noncompliant run in which the transition “requestInfo” is fired more than once. The diagnosis algorithm (which is also implemented in the tool Wendy [16]) generates a graph similar to that depicted in Fig. 9. From this graph, we pruned those paths that eventually reach a final state (represented as check marks) and replaced them by dashed arrows. As we can see, a final state cannot be reached after the second “requestInfo” action is executed. However,




Fig. 9. Counterexample for an unrealizable compliance rule

we cannot avoid this situation, because in the gray state, the external service decides whether or not the claim is fraudulent or whether further information are required. For the business process, this choice is external and uncontrollable and hence it must be correct for any outcome.

6 Related Work

Compliance has received a lot of attention in the business process management community. Contributions related to our approach can be classified as follows.

Compliance by detection. Awad et al. [3] investigate a pattern-based compliance check based on BPMN-Q [2]. They also cover the compliance rule classes defined by Dwyer et al. [8] and give a CTL formalization as well as an antipattern for each rule. These antipatterns are used to highlight the compliance violations in a BPMN model. Such a visualization is very valuable for the process designer and it would be interesting to see whether such antipatterns are also applicable to the artifact-centric approach. Sadiq et al. [26] use a declarative specification of compliance rules from which they derive compliance checks. These checks are then annotated to a business process and monitored during its execution. These checks are similar to the nonblocking compliance rule models that only monitor behavior rather than constraining it. Lu et al. [18] compare business processes with compliance rules and derive a compliance degree. This is an interesting approach, because it replaces yes/no answers by numeric values which could help to easier diagnose noncompliance. Knuplesch et al. [12] analyze data aspects of operational business process models. Similar to the artifact-centric approach, data values are abstracted into compact life cycles.

Compliance by design. Goedertier and Vanthienen [10] introduce the declarative language PENELOPE to specify compliance rules. From these rules, a state space and a BPMN model is generated which is compliant by design. This approach is limited to acyclic process models. Furthermore, the purpose of the generated model is rather the validation of the specified rules than the execution. Küster et al. [13] study the interplay between control flow models and object life cycles. The authors present an algorithm to automatically derive a sound process model from given object life cycles. The framework is, however, not designed to express dependencies between life cycles and therefore cannot specify complex policies or compliance rules.

To the best of knowledge, this paper presents the first approach that generates compliant and operational business process models from declarative specifications of artifact life cycles, policies, and compliance rules.

7 Conclusion

Summary. We presented an approach to automatically construct business process models that are compliant by design. The approach follows an artifact-centric modeling style in which business processes are specified from the point

of view of the involved data objects. We showed how artifact-centric business processes can be canonically extended to also take compliance rules into account. These rules can express constraints on the execution of actions, but can also take data and location information into account. By composing compliance rules to the artifact-centric model, we could reduce the check for compliant behavior to the reachability of final states. Consequently, we could use existing synthesis algorithms and tools to automatically generate compliant business process models. In case of noncompliance, we further sketched diagnosis information that can be used to visualize the reasons that make compliance rules unrealizable.

Lessons learnt. With Petri nets, we can use a single formalism to model artifact life cycles, interartifact dependencies, and compliance rules. Only this unified way of modeling enabled us to approach the compliance-by-design approach with only a few concepts (namely artifacts, composition, and partner synthesis). In addition, it is notable that the compliance rule models can also be used to check operational business process models for compliance: By composing compliance rules to existing Petri net models, we reduced the compliance check by a check for weak termination and allows to use standard verification tools [9].

Future work. We see numerous directions to continue the work in the area of compliance by design. We are currently working on an extension for BPMN to provide a graphical notation that is more accessible for domain experts to model artifacts, policies, and compliance rules. A canonic second step would then be the integration of the approach into a modeling tool and an empirical evaluation thereof. Another aspect that needs to be addressed is the expressiveness of the artifacts and compliance rules. Of great interest are *instances* and *agent roles*. To model more involved scenarios, it is crucial distinguish several instances of an artifact. We currently assume that for each artifact only a single instance exists. Furthermore, agent roles would allow a fine-grained description of access controls or concepts such as the four-eye principle.

References

1. van der Aalst, W.M.P., Pesic, M.: DecSerFlow: Towards a truly declarative service flow language. In: Bravetti, M., Núñez, M., Tennenholtz, M. (eds.) WS-FM 2006. LNCS, vol. 4184, pp. 1–23. Springer, Heidelberg (2006)
2. Awad, A.: BPMN-Q: a language to query business processes. In: EMISA 2007. LNI P-119, pp. 115–128. GI (2007)
3. Awad, A., Weidlich, M., Weske, M.: Visually specifying compliance rules and explaining their violations for business processes. *J. Vis. Lang. Comput.* 22(1), 30–55 (2011)
4. Ben-Ari, M., Manna, Z., Pnueli, A.: The temporal logic of branching time. In: POPL 1981, pp. 164–176. ACM, New York (1981)
5. Cannon, J.C., Byers, M.: Compliance deconstructed. *ACM Queue* 4(7), 30–37 (2006)
6. Clarke, E.M., Grumberg, O., Peled, D.A.: *Model Checking*. MIT Press, Cambridge (1999)
7. Cortadella, J., Kishinevsky, M., Kondratyev, A., Lavagno, L., Yakovlev, A.: Petrify: A tool for manipulating concurrent specifications and synthesis of asynchronous controllers. *Trans. Inf. and Syst.* E80-D(3), 315–325 (1997)
8. Dwyer, M.B., Avrunin, G.S., Corbett, J.C.: Patterns in property specifications for finite-state verification. In: ICSE 1999, pp. 411–420. IEEE, Los Alamitos (1999)

9. Fahland, D., Favre, C., Jobstmann, B., Koehler, J., Lohmann, N., Völzer, H., Wolf, K.: Instantaneous soundness checking of industrial business process models. In: Dayal, U., Eder, J., Koehler, J., Reijers, H.A. (eds.) BPM 2009. LNCS, vol. 5701, pp. 278–293. Springer, Heidelberg (2009)
10. Goedertier, S., Vanthienen, J.: Designing compliant business processes with obligations and permissions. In: Eder, J., Dustdar, S. (eds.) BPM Workshops 2006. LNCS, vol. 4103, pp. 5–14. Springer, Heidelberg (2006)
11. Havelund, K., Roşu, G.: Testing linear temporal logic formulae on finite execution traces. Technical Report 01.08, RIACS (2001)
12. Knuplesch, D., Ly, L.T., Rinderle-Ma, S., Pfeifer, H., Dadam, P.: On enabling data-aware compliance checking of business process models. In: Parsons, J., Saeki, M., Shoval, P., Woo, C., Wand, Y. (eds.) ER 2010. LNCS, vol. 6412, pp. 332–346. Springer, Heidelberg (2010)
13. Küster, J.M., Ryndina, K., Gall, H.: Generation of business process models for object life cycle compliance. In: Alonso, G., Dadam, P., Rosemann, M. (eds.) BPM 2007. LNCS, vol. 4714, pp. 165–181. Springer, Heidelberg (2007)
14. Lohmann, N.: Why does my service have no partners? In: Bruni, R., Wolf, K. (eds.) WS-FM 2008. LNCS, vol. 5387, pp. 191–206. Springer, Heidelberg (2009)
15. Lohmann, N., Massuthe, P., Wolf, K.: Behavioral constraints for services. In: Alonso, G., Dadam, P., Rosemann, M. (eds.) BPM 2007. LNCS, vol. 4714, pp. 271–287. Springer, Heidelberg (2007)
16. Lohmann, N., Weinberg, D.: Wendy: A tool to synthesize partners for services. In: Lilius, J., Penczek, W. (eds.) PETRI NETS 2010. LNCS, vol. 6128, pp. 297–307. Springer, Heidelberg (2010), <http://service-technology.org/wendy>
17. Lohmann, N., Wolf, K.: Artifact-centric choreographies. In: Maglio, P.P., Weske, M., Yang, J., Fantinato, M. (eds.) ICSSOC 2010. LNCS, vol. 6470, pp. 32–46. Springer, Heidelberg (2010)
18. Lu, R., Sadiq, S.W., Governatori, G.: Compliance aware business process design. In: ter Hofstede, A.H.M., Benatallah, B., Paik, H.-Y. (eds.) BPM Workshops 2007. LNCS, vol. 4928, pp. 120–131. Springer, Heidelberg (2008)
19. OMG: Business Process Model and Notation (BPMN). Version 2.0, Object Management Group (2011), <http://www.omg.org/spec/BPMN/2.0>
20. Pnueli, A.: In transition from global to modular temporal reasoning about programs. In: Logics and models of concurrent systems. NATO Advanced Summer Institutes, vol. F-13, pp. 123–144. Springer, Heidelberg (1985)
21. Pnueli, A.: The temporal logic of programs. In: FOCS 1977, pp. 46–57. IEEE, Los Alamitos (1977)
22. Ramadge, P., Wonham, W.: Supervisory control of a class of discrete event processes. SIAM J. Control Optim. 25(1), 206–230 (1987)
23. Reisig, W.: Petri Nets. EATCS Monographs on Theoretical Computer Science edn. Springer, Heidelberg (1985)
24. Ryndina, K., Küster, J.M., Gall, H.: Consistency of business process models and object life cycles. In: Kühne, T. (ed.) MoDELS 2006. LNCS, vol. 4364, pp. 80–90. Springer, Heidelberg (2007)
25. Sackmann, S., Kähler, M., Gilliot, M., Lowis, L.: A classification model for automating compliance. In: CEC/EEE 2008, pp. 79–86. IEEE, Los Alamitos (2008)
26. Sadiq, S.W., Governatori, G., Namiri, K.: Modeling control objectives for business process compliance. In: Alonso, G., Dadam, P., Rosemann, M. (eds.) BPM 2007. LNCS, vol. 4714, pp. 149–164. Springer, Heidelberg (2007)
27. Wolf, K.: Does my service have partners? In: Jensen, K., van der Aalst, W.M.P. (eds.) Transactions on Petri Nets and Other Models of Concurrency II. LNCS, vol. 5460, pp. 152–171. Springer, Heidelberg (2009)