


Journal Pre-proof

Verifying compliance in process choreographies: Foundations, algorithms, and implementation

Walid Fdhila, David Knuplesch, Stefanie Rinderle-Ma, Manfred Reichert



PII: S0306-4379(22)00002-3
DOI: <https://doi.org/10.1016/j.is.2022.101983>
Reference: IS 101983

To appear in: *Information Systems*

Received date : 1 April 2021
Revised date : 24 December 2021
Accepted date : 6 January 2022

Please cite this article as: W. Fdhila, D. Knuplesch, S. Rinderle-Ma et al., Verifying compliance in process choreographies: Foundations, algorithms, and implementation, *Information Systems* (2022), doi: <https://doi.org/10.1016/j.is.2022.101983>.

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2022 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Verifying Compliance in Process Choreographies: Foundations, Algorithms, and Implementation

Walid Fdhila^a, David Knuplesch^b, Stefanie Rinderle-Ma^c and Manfred Reichert^d

^aSecurity Business Austria (SBA-research), Vienna, Austria, wfdhila@sba-research.org

^balphaQuest GmbH, Ulm, Germany, d.knuplesch@alphaquest.de

^cTechnical University of Munich, Department of Informatics, Boltzmannstrasse 3, Garching, Germany, stefanie.rinderle-ma@tum.de

^dInstitute of Databases and Information Systems, Ulm University, Ulm, Germany, manfred.reichert@uni-ulm.de

ARTICLE INFO

Keywords:

Distributed business processes
Business rules
Compliance management
Process choreography
Business process compliance
Global compliance rules
Rule decomposition

ABSTRACT

The current push towards interoperability drives companies to collaborate through process choreographies. At the same time, they face a jungle of continuously changing regulations, e.g., due to the pandemic and developments such as the BREXIT, which strongly affect cross-organizational collaborations. Think of, for example, supply chains spanning several countries with different and maybe even conflicting COVID19 travelling restrictions. Hence, providing automatic compliance verification in process choreographies is crucial for any cross-organizational business process. A particular challenge concerns the restricted visibility of the partner processes at the presence of global compliance rules (GCR), i.e., rules that span across the process of several partners. This work deals with the question how to verify global compliance if affected tasks are not fully visible. Our idea is to decompose GCRs into so called assertions that can be checked by each affected partner whereby the decomposition is both correct and lossless. The algorithm exploits transitivity properties of the underlying rule specification, and its correctness and complexity are proven, considering advanced aspects such as loops. The algorithm is implemented in a proof-of-concept prototype, including a model checker for verifying compliance. The applicability of the approach is further demonstrated on a real-world manufacturing use case.

1. Introduction

Gartner regards interoperability as “strategic imperative”¹ for healthcare. Especially the global push by digitalization and the current pandemic require the collaboration and integration of (business) partners and organizations. Process technology serves as enabler for process-oriented collaborations between distributed business partners, realized and implemented through so-called *process choreographies*. Applications include healthcare (34), blockchain-based processes (48; 47), multi-modal logistics scenarios (7; 10), and supply chains (22).

Digitalization and ongoing changes due to, for example, the pandemic situation or the BREXIT flood enterprises and organizations with updated or even new regulations at a fast pace. For example, “bank regulations change every 12 minutes”.² Regulatory frameworks comprise application-independent frameworks such as the GDPR on “data processing boundaries of the personal data of European Union’s citizens” (66) and the ISO 27001 security standard³ as well as application-specific ones, e.g., the WHO regulations defined in the context of COVID19⁴. As a consequence, in our globalized world, regulations and their changes affect process collaborations (55) and lead to an increased need for *compliance verification* in process choreographies.

1.1. Problem Statement

What are the particular challenges with respect to compliance verification in process choreographies? Let us illustrate them by an example. Figure 1 depicts a BPMN choreography model of a supply chain scenario adapted from (22). It involves five process partners, i.e., *Bulk Buyer*, *Manufacturer*, *Middleman*, *Supplier*, and *Special Carrier* that interact

Corresponding author

✉ wfdhila@sba-research.org (W. Fdhila)

ORCID(s): 0000-0002-5245-6128 (W. Fdhila); 0000-0001-5656-6108 (S. Rinderle-Ma); 0000-0003-2536-4153 (M. Reichert)

¹<https://gtmr.it/3vGFB7f>

²<https://thefinanser.com/2017/01/bank-regulations-change-every-12-minutes.html/>

³<https://www.iso.org/iso/iec-27001-information-security.html>

⁴<https://www.who.int/teams/regulation-prequalification/procurement/covid-19>

Verifying Compliance in Process Choreographies

through message exchanges. First, the *Bulk Buyer* orders a set of products from the *Manufacturer* (e.g., an aircraft). The manufacturing of the product requires several sub-products (intermediates) to be provided by different suppliers. In this scenario, we assume that only one intermediate is required and provided by the *Supplier*. After processing the order, hence, the *Manufacturer* sends an order for the intermediate (e.g., the fuselage or engines) to the *Middleman*. The *Middleman* forwards the order of the intermediate to the *Supplier* and sends an order for a special transport to the *Special Carrier*. The *Special Carrier* requests the details on the transport from the *Supplier* and the *Supplier* provides them to the *Special Carrier*, followed by sending the waybill for the intermediate. The *Special Carrier* sends a notice on the arrival of the intermediate to the *Manufacturer*, which then delivers the product to the *Bulk Buyer*.

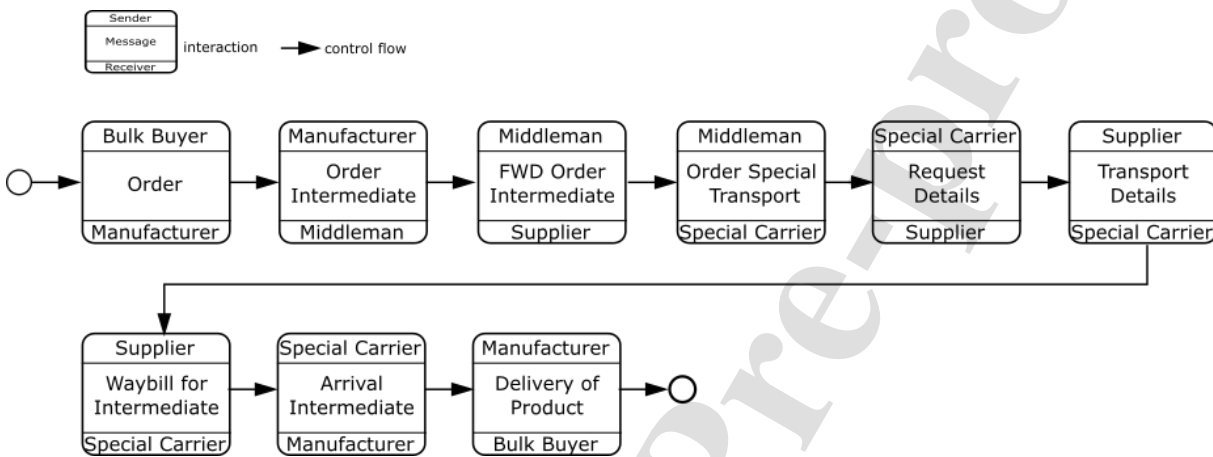


Figure 1: BPMN choreography model for a supply chain – running example with five process partners

Imagine now that the partners and the choreography are subject to the *Global Compliance Rules (GCR)* depicted in Fig. 2, which stem from legal regulations and standards such as GDPR or ISO 27001:

- C1 After *Production* a *Final test* must be performed.
- C2 *Pack Intermediate* is required before *Transport Intermediate*.
- C3 Each *Transport intermediate* requires *Permission of authority*. Further on, the transporter must pass a *Safety Check*.

Obviously, none of the GCR can be directly verified on basis of the choreography model in Fig. 1 as none of the public and message exchanging tasks corresponds to any of the tasks referred to in the GCR.

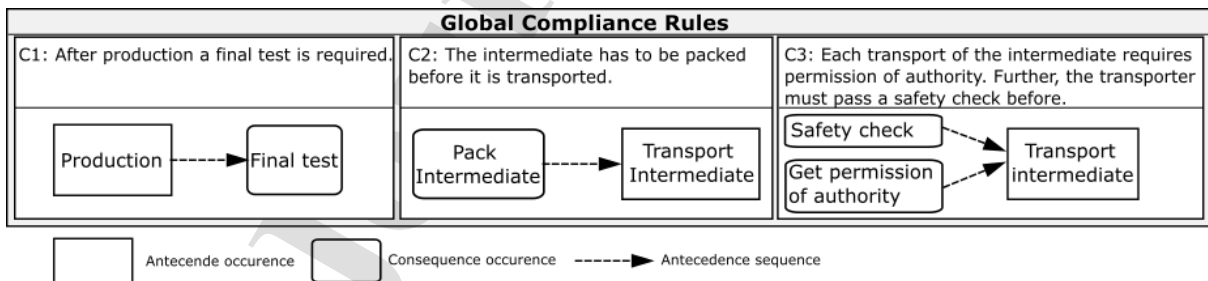


Figure 2: eCRG model of Global Compliance Rules imposed on supply chain choreography

Let us have a look at the *public processes* of the partners involved in the choreography as shown in Fig. 3. These public process models contain all public tasks that are visible to the other partners, including the tasks that exchange messages, but also other visible tasks such as *Production* at the *Manufacturer*. Based on the public process models, C1 and C2,

Verifying Compliance in Process Choreographies

as depicted in Fig. 2, can be verified: C1 refers to public tasks of the *Manufacturer* process, which obviously complies with C1, i.e., public task *Production* is followed by public task *Final Test*. C2 can be verified over the *Supplier* and *Special Carrier* processes. The order between tasks *Pack Intermediate* and *Transport Intermediate* is determined by the message exchange and receiving *Waybill Intermediate*. As opposed to C1 and C2, C3 cannot be verified based on the public processes of the partners as there are no public tasks for *Safety check* and *Get permission of authority*.

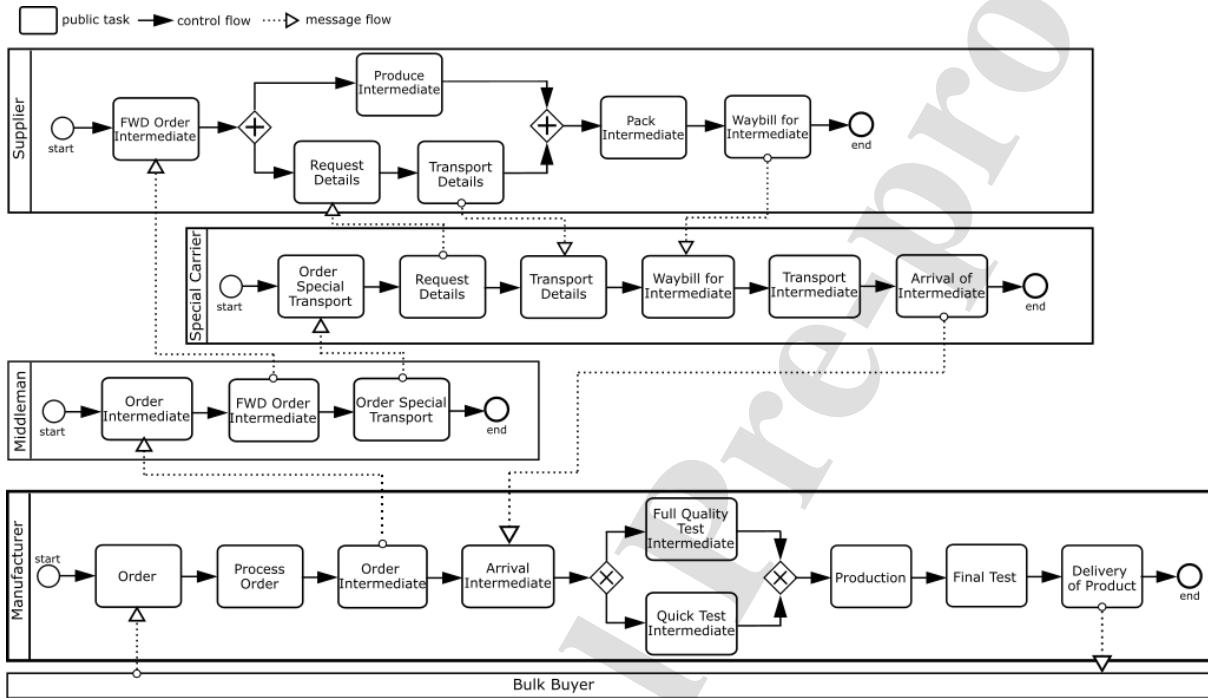


Figure 3: Public processes (collaboration model) – running example with five process partners cooperating in a supply chain (adapted from (22))

The presumption is that C3 also refers to *private tasks* of the partners, i.e., tasks that are only present in the *private process models* of the partners. In general, private process models of the partners implement and possibly extend the behavior of the corresponding public models. As opposed to public tasks, private tasks are not visible to the partners. Figure 4 shows the private process models of partners *Special Carrier* and *Middleman* where *private tasks* are highlighted in gray color. Although private tasks are usually hidden to other partners, restrictions over them might exist. In the supply chain, for example, C3 refers to private tasks *Safety Check* for partner *Special Carrier* and *Get permission of authority* for partner *Middleman*. If private tasks are affected by a GCR, no information about how and when these tasks are executed, or how they are connected to other nodes of the corresponding private process model, becomes visible to the other partners. Usually, this happens when a collaborating partner p1 imposes the execution of a specific task that must exist in its private process and comply with a given rule involving another partner p2. Partner p1 should then assure the existence of such task and that it follows the imposed rule.

As can be seen from the example depicted in Figs. 1–4, GCRs constrain actions of multiple partners and/or the interactions between them. Ensuring the compliance of process choreographies with a GCR is crucial and challenging (32) as a partner “only has the visibility of the portion of the process under its direct control” (49). Reconsider GCR C3 as an example. It asks for a *safety check* accomplished by a private task of the *Special Carrier*. To cope with this issue, *assertions* can be used. An assertion (A) corresponds to a commitment of a partner guaranteeing that its private/public process complies with the imposed rule (22). Figure 5 depicts the two assertions A1 and A2: the *Middleman* agrees to get the permission of the authority before ordering the special transport (A1). Moreover, the *Special Carrier* commits to perform a safety check before transporting the intermediate (A2). In combination, assertions A1 and A2 enable

Verifying Compliance in Process Choreographies

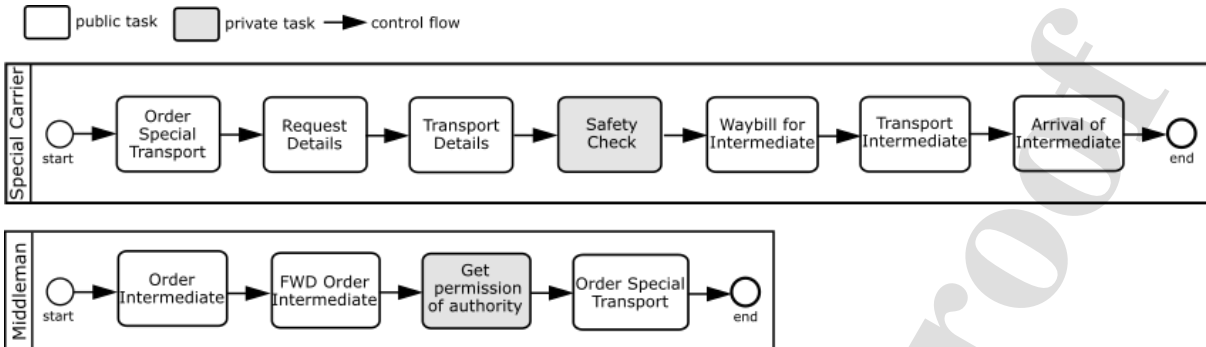


Figure 4: Private processes of partners *Special Carrier* and *Middleman*, omitting message exchanges (adapted from (22))

checking GCR C3.

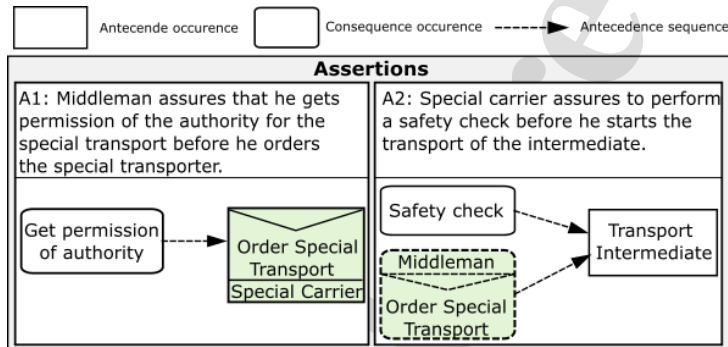


Figure 5: Assertions by partners *Middleman* and *Special Carrier*

1.2. Contribution

Overall, this leads to the overarching research question *RQ* tackled in this work:

RQ: How to verify GCRs in a decentralized setting of a process choreography where no central coordinator with complete knowledge on the private and public tasks of all partners exists?

In literature, there is a “*knowledge gap*” when it comes to compliance verification in process choreographies (32). (49) tackles the problem of checking a GCR on private tasks based on IoT-enabled artifacts. However, not all process choreography settings with compliance requirements feature IoT-enabled artifacts. Hence, this work aims at providing a formal approach that is independent of any technology or application. The central idea is to decompose the GCR into assertions in a lossless way, i.e., the verification of all assertions guarantees the one of the GCR. Note that this solves the problem as assertions may be checked separately by each of the partners. Hence, infringing the privacy of any partner is avoided.

The decomposition algorithm presented in this article exploits transitivity properties of the underlying GCR specification and was originally presented in (23). The decomposition relies on transitivity properties of the underlying GCR specification. The transitivity properties are shown by the example of a translation to first order predicate logic. Similarly, for example, (58) presents declarative patterns based on Linear Temporal Logic (LTL).

In our approach, GCRs are specified in a pattern-based and visual way using the eCRG formalism (40). This means that a GCR may be composed of so called *antecedence patterns* and *consequence patterns*. The patterns can be connected reflecting pre-/post-conditions of the respective GCR. C1 in Fig. 2, for example, connects antecedence

Verifying Compliance in Process Choreographies

pattern *Production* with consequence pattern *Final test*, demanding that after the production a final test is required. Note that antecedence and consequence patterns may require *occurrence* (i.e., something must happen) and *absence* (i.e., something must not happen). In (23), we relied on simple rules that consist of single antecedent and multiple occurrence patterns. Aside the decomposition algorithm itself, (23) provides basic proofs, simple GCR decomposition scenarios, and the embedding of the approach in the overall digitalized change and compliance management framework C³Pro.⁵ This article extends and elaborates the results presented in (23) in several directions:

- We allow for additional and more complex compliance rules with multiple antecedence patterns. This significantly increases the complexity of the theoretical considerations as well as the one of the provided GCR decomposition scenarios. As a result, we obtain new theorems and algorithms.
- The decomposition proofs are extended to cover the additional complexity of the GCR; in particular they now consider loops as well.
- The decomposition algorithm with extensions is prototypically implemented and integrated with the C³Pro framework, which deals with both change and compliance in process choreographies.
- A model checker for verifying decomposition correctness is provided.
- A manufacturing use case illustrates the applicability of the approach. Specifically, the use case demonstrates the applicability of the approach beyond regulatory compliance, i.e., it shows how decomposition can be used to lift implicit connections to explicit assertions.
- The related work section is significantly extended.

The remainder of the paper is structured as follows: Section 2 provides fundamentals required for understanding this work, and Section 3 introduces the foundations for GCR decomposition (including transitivity theorems). Section 4 then presents the decomposition algorithm for global compliance rules, whereas Section 5 deals with the automated verification of the resulting decompositions based on model checking. Sections 6 and 7 cover the evaluation of the approach, i.e., the implementation and application of the algorithms. Section 8 discusses related work. Section 9 concludes the paper with a summary and an outlook.

2. Fundamentals

This section presents definitions and formal backgrounds for global compliance rules (GCRs) to be obeyed by a process choreography. A *choreography* includes three types of overlapping models: (i) a *private model* representing the executable process and including both private and public activities (see Fig. 4 for examples of private process models), (ii) a *public model* (also called the *interface* of the process) that solely includes the public activities and the interactions of a given partner (see Fig. 3 for the public process models of our running example), and (iii) a *choreography model* providing a global view on the interactions between all partners (see Fig. 1 for the choreography model of our running example) (21). Compared to (21; 2), this paper assumes that public activities are not necessarily interactions with other partners, but may additionally represent tasks made visible by the corresponding partners. Therefore, both interactions and non-interaction public activities of a single partner are described in a public model. The latter serves as public (restricted) view on the private model of the partner, which “describes the internal logic of a partner including its private and public activities” (6). For a formal definition of process choreography, we refer to Definition 1.

Definition 1 (Choreography (21)). We define a choreography as a tuple $(\mathcal{P}, \mathcal{C}, \mathcal{M}, \mathcal{I}, \mathcal{R}, \mathcal{S})$ where

- \mathcal{P} is the set of all participating partners.
- \mathcal{C} is the choreography model representing the interactions between partners in \mathcal{P} (cf. Fig. 1).
- $\mathcal{M} = \{ \mathcal{M}_p \mid p \in \mathcal{P} \}$ is the set of all private models (cf. Fig. 4).
- $\mathcal{I} = \{ \mathcal{I}_p \mid p \in \mathcal{P} \}$ is the set of all public models (cf. Fig. 3).
- $\mathcal{R} = \{ \mathcal{R}_{p,q} \mid p, q \in \mathcal{P} \}$ is a partial mapping function between nodes of the public and private models.
- $\mathcal{S} : \mathcal{I} \rightarrow \mathcal{M}$ is a partial mapping function between nodes of different public models.
- $\mathcal{S} : \mathcal{C} \rightarrow \mathcal{I}$ is a partial mapping function between nodes of the choreography model and the public models.

⁵<http://www.wst.uni.wue.ac.at/communities/c3pro>

Verifying Compliance in Process Choreographies

As depicted in Figures 1, 3 and 4, choreography, public and private models are defined as graphs, where nodes are either activities (i.e., interaction, public or private activities) or gateways (e.g., sequence, exclusive or parallel), and arrows are the dependencies between them. As described above, each of these three models use specific type of activity nodes (e.g., interaction activities for choreography models). Because the focus of this paper is mainly on GCR decomposition, we abstract their respective formal definitions, but the reader may refer to (21) for more details.

While function γ maps the activities of the public models to those of the private models, function δ determines the dependencies between the interactions of different public models (e.g., $(\text{send}_p, \text{receive}_p) = \text{receive}_p$ (send_p)). Finally, function ϵ represents the dependencies between the activities in the choreography model and those of the public models (e.g., $(\text{send}_p, \text{receive}_p) = \{ (\text{send}_p, \text{receive}_p), (\text{receive}_p, \text{send}_p) \}$). Note that in the examples above, connected interaction activities (i.e., the send and the corresponding receive) of different public models have the same labels, while in practice, it is possible to have them different.

Based on functions γ and δ , certain soundness properties of choreography can be checked, including *structural and behavioral compatibility* between public models, and *consistency* between public and private models (15). Structural consistency requires that for each public activity of the public model of a partner p , there should be a matching element in the corresponding private model of p , but not vice versa (21). Structural compatibility states that for each interaction activity of the public model of a partner p , there should be a matching interaction activity in the public model of another partner. Note that this is a necessary, but not yet sufficient condition for ensuring compatibility and consistency—the models' behaviors (control flow dependencies) should also be compatible and consistent. In this paper, we assume that the choreography is sound.

In previous work (20), multiple formal languages employed for business process compliance modelling and checking (e.g., linear temporal logic LTL, event calculus EC, extended compliance rule graph eCRG) were compared. It was shown that most of these languages can deal with most qualitative time patterns, and can therefore be used to model the compliance constructs addressed in this paper. Similar results were proven in (30).

In order to specify these constructs as well as transitivity properties required for the GCR decomposition, this work utilizes the visual *eCRG (extended Compliance Rule Graph)* language (42; 38; 36). The eCRG language offers a visual rule notation for expressing compliance rules over process choreographies and is based on first order predicate logic (cf. Fig. 6). To distinguish between a precondition (i.e., antecedence) and corresponding postconditions (i.e., consequences), an eCRG can be partitioned into an *antecedence pattern* and a *consequence pattern*. The antecedence pattern specifies when the compliance rule is triggered (i.e., activated), whereas the consequence pattern specifies what the rule demands. As compliance rules may require the occurrence or absence of certain activities or interactions (i.e., message exchanges), the antecedence and consequence patterns are further sub-divided into occurrence and absence nodes. Sequence conditions between these events can be expressed using directed connectors between the respective nodes. We use the following notation: \square : Antecedence occurrence; \boxtimes : Antecedence absence; \circ : Consequence occurrence; \otimes : Consequence absence. Fig. 6 introduces the elements of the eCRG language. For a formal definition of eCRG, we refer to Def. 2.

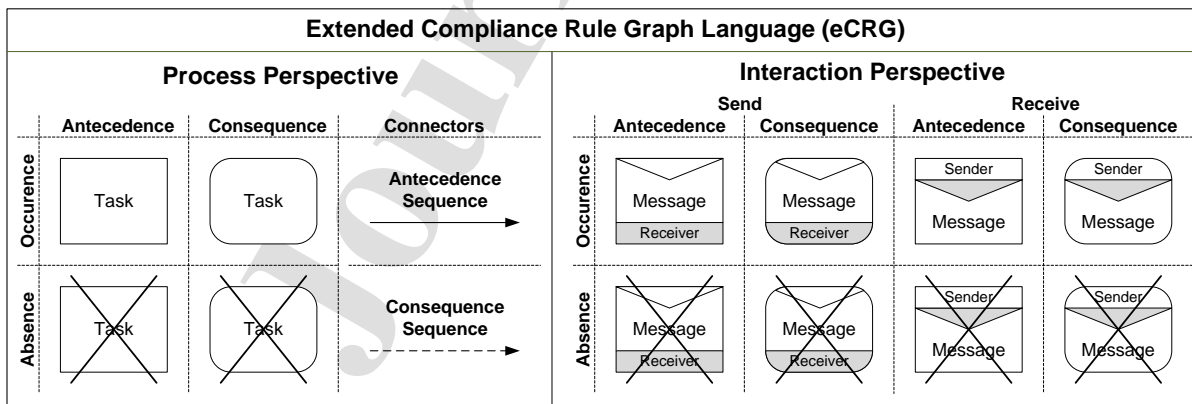


Figure 6: Elements of the eCRG language

Verifying Compliance in Process Choreographies

Definition 2 (Global Compliance Rule (GCR) structure). Given a process choreography $\mathcal{C} = (\mathcal{N}, \mathcal{A}, \mathcal{I}, \mathcal{S}, \mathcal{R}, \mathcal{C})$ (cf. Def. 1), let \mathcal{N}_p be the set of private and public non-interaction activities and \mathcal{N}_i be the set of interaction activities. Then: A GCR is defined as tuple $\mathcal{G} = (\mathcal{N}, \mathcal{A}, \mathcal{I}, \mathcal{S}, \mathcal{R}, \mathcal{C})$ with

- \mathcal{N} being the set of nodes,
- \mathcal{A} returning the partner responsible for a node.
- $\mathcal{I} \times \mathcal{N} \times \mathcal{N} \times \{ \rightarrow, \dashrightarrow \}$ returning the sequence flow connector between two nodes, i.e., consequence sequence and antecedence sequence connectors respectively.
- \mathcal{S} mapping each node to an activity or an interaction (i.e., message exchange).
- $\mathcal{R} \subseteq \{ \square, \boxtimes, \circ, \otimes \}$

Think of an eCRG as a graph of connected nodes, where each node is assigned to a particular partner (e.g., in C1, $(\text{node}) = (\text{partner})$). A node may either be a private, non-interaction public activity, or an interaction (see Figure 6). Given two nodes of an eCRG, function \mathcal{I} returns the sequence flow connector as depicted in Figure 6, where a dashed arrow (i.e., consequence connector) connects an antecedence pattern to set of consequence patterns (e.g., 1: After production a final test is required), and an antecedence connector expresses a relation between antecedence patterns solely (i.e., the pre-condition). For example, assume that we change \mathcal{C}_3 as follows: $\text{node}_1 \dashrightarrow \text{node}_2$. Then: if the pre-condition (i.e., execution of activity node_1) followed by the one of activity node_2) is met, the post condition (i.e., activity node_2) will be triggered. Finally, function \mathcal{R} evaluates whether a node is an antecedence or consequence, and whether or not it should occur.

3. Global Compliance Rule Decomposition Theorems

This section introduces the theoretical foundations for the decomposition of global compliance rules (including theorems and proofs) illustrated by a number of examples, which we have extracted from the application scenario introduced in Section 1. Section 3.1 first describes the basic idea of our approach (i.e., why do we need to decompose a GCR), followed by the presentation, proofs and illustrations of the theorems in Section 3.2.

3.1. Basic Idea

Our method for the decentralized checking of global compliance rules relies on the decomposition of the original GCR into a set of assertions that can be checked locally by each partner and collectively reproduce the behavior of the GCR (cf. Fig. 7). A communication between partners is only required in the setup phase to deduce the assertions. During runtime, however, no further compliance-related communication becomes necessary for checking the GCR unless a local assertion becomes violated. The decomposition of a GCR into a set of assertions is based on well-grounded theorems, which ensure that if a conjunction of hypotheses is true, the conclusion (GCR) is true as well.

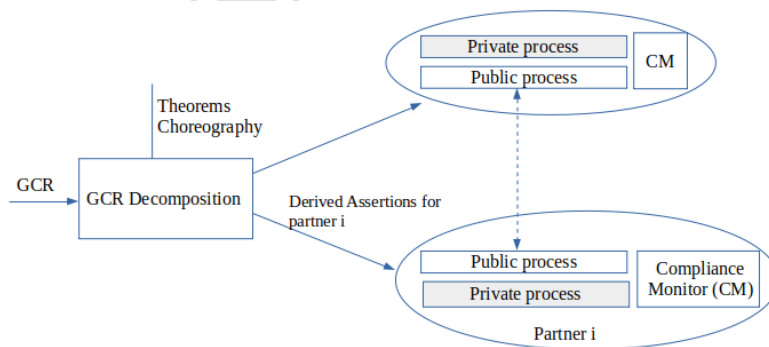


Figure 7: Method Overview

3.2. Theorems

In this section, we provide a decomposition method for selected global compliance patterns and show how they can be applied in a collaborative setting. In particular, we prove a set of theorems required for ensuring the correctness of our decomposition method. Each theorem represents a possible decomposition of a given compliance pattern.

We illustrate the translation of a GCR into a First Order Logic (FOL) expression using basic equivalences as in Def. 3.

Definition 3 (Basic Equivalences). *Based on (41), the following equivalences hold by definition. Predicate $(,)$ describes that at the point in time t an activity (message) of type act was executed (i.e., sent or received).*

- $\square \circ \quad (,) \quad ((,) <)$
 $(,) \quad ((,) <)$
- $\circ \square \quad (,) \quad ((,) <)$
 $(,) \quad ((,) <)$
- $\square \otimes \quad (,) \quad ((,) <)$
 $(,) \quad \neg((,) <)$
 $(,) \quad (\neg((,) <))$
- $\otimes \square \quad (,) \quad ((,) <)$
 $(,) \quad \neg((,) <)$
 $(,) \quad (\neg((,) <))$

For example, GCR $\boxed{\text{Production}}$ $\boxed{\text{Final test}}$ is translated into: $(, \text{Production}) \quad (, \text{Final test}) < .$ Thereby, relation $<$ expresses a temporal precedence between points in time t and t' . The decomposition algorithm presented in Section 4 exploits the transitivities for GRC as in Theorem 1. Specifically, by combining transitive relations, where each relation can be checked locally by a single partner, it becomes possible to reconstruct the original GCR behavior. Theorem 1 ensures that the behavior of the derived assertions reproduces the behavior of the GCR, but not vice versa.

Theorem 1 (Transitivities).

Let $act_1, act_2,$ and act_3 be three activity or message types. Then:

a. The rightwards transitivity holds:

$$\square \circ \square \circ \square \circ$$

b. The leftwards transitivity holds:

$$\circ \square \circ \square \circ \square$$

In the following, the correctness of Theorem 1 is proven applying Def. 3.

Proof 1 (Rightwards Transitivity).

Let $act_1, act_2,$ and act_3 be three activities or interactions. Then $\square \circ \square \circ$
 $(,) \quad ((,) <) \quad (,) \quad ((,) <)$
 $(,) \quad ((,) <) \quad (,) \quad ((,) <)$
 $(,) \quad ((,) <) \quad (,) \quad ((,) <)$
 $(,) \quad (,) \quad ((,) < <)$
 $(,) \quad ((,) <) \quad \square \circ \quad q.e.d.$

Verifying Compliance in Process Choreographies

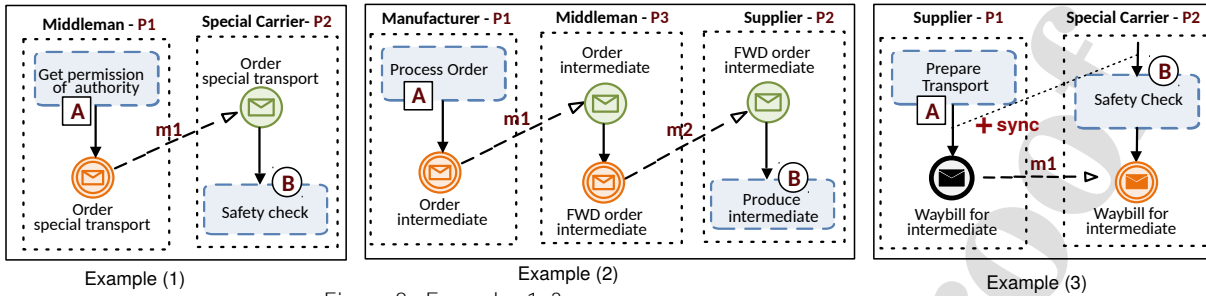


Figure 8: Examples 1-3

Leftwards transitivity can be proven similarly by replacing ' $<$ ' with ' $>$ '.

Corollary 1. Let α , β , and γ be activities or interactions. Then



In the following, we use Examples (1) - (3) (cf. Fig. 8), which we extracted from our running example (cf. Figs. 3 and 4), in order to illustrate how we use Theorem 1 for decomposing a simple compliance rule of type $\square \circ$ that involves two private tasks α and β of two different partners P_1 and P_2 respectively.

- Example (1): $\square \circ$. In this example, both activities are private, which would normally require α and β to share runtime information about the execution time of the respective activities. In turn, this would require an agreement on a time synchronization protocol that considers network failures and message transmission delays. This can be solved by identifying a transitive relation between the two private activities that include an interaction activity. According to Theorem 1, the interaction activity m_1 between α and β fulfills the conditions P_1 and P_2 :

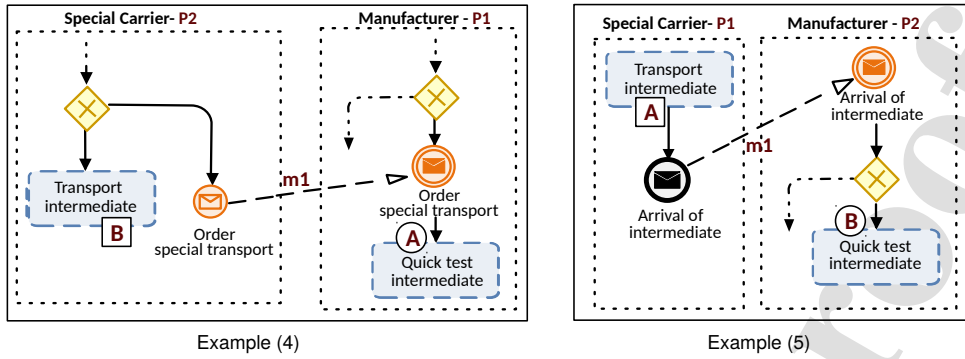


The behavioral and structural compatibility (cf. Section 2) between the partner processes ensures that message m_1 sent by P_1 shall be correctly received by P_2 . By locally checking α and β by P_1 and P_2 respectively, we can ensure that the original GCR is not violated as long as the assertions are not violated. If one assertion is violated, a communication between the two partners will become necessary. Note that this violation does not necessarily mean that the original GCR is violated. For example, assume that for a given process instance, assertion α evaluates to true, and β execute α activity before the message arrival. Although this would result in β being evaluated to false, it does not necessarily mean that β is executed before α .

- Example (2): $\square \circ$. In this example, α and β do not engage in any direct interaction. However, by looking at the public processes of the collaboration model from Fig. 3, it becomes possible to identify a double transitive relation through m_1 , which interacts with both partners. Therefore, using Corollary 1, the transitive relations (assertions): $\alpha \rightarrow m_1 \rightarrow \beta$, $\beta \rightarrow m_2 \rightarrow \alpha$, and $\alpha \rightarrow m_1 \rightarrow m_2 \rightarrow \alpha$ reproduce the behavior of $\alpha \rightarrow \beta$, which has initially not been involved in the GCR, becomes involved in the derived assertions. We call m_1 an *intermediary partner*.

- Example (3): $\square \circ$. In this example, it is not possible to identify any transitive relation between α and β that involve private activities α and β . The interaction activity m_1 connects both partners immediately after the activities in question, which discards any possibility of a transitive relation. In this case, it is not possible to satisfy Theorem 1 and, hence, additional message exchanges become necessary to inform about the execution state of the activities involved in the GCR. Message exchanges can be either synchronous or asynchronous. Asynchronous message exchange only allows for reactive GCR checking and, therefore, detecting violations after their occurrence. Synchronous message exchange, in turn, is proactive as it enforces the GCR with new

Verifying Compliance in Process Choreographies



Example (4)
Figure 9: Examples 4–5


restrictions to the process models, e.g., the execution of activity \square becomes enabled only after receiving a synchronization message (i.e., about whether or not \square is executed). \square shall also inform \square in case activity \square is not executed, as this does not prevent activity \square from being executed according to the GCR.

Rightwards transitivity (cf. Theorem 1.a) directly ensures the correctness of the assertions derived in the above examples. It should be clear that the correctness of the derived assertions in Example (2) can be directly concluded based on Corollary 1. The same examples can be also used to illustrate leftwards transitivity.

Theorem 2 (Zig zag Transitivity).

Let \square , \square , and \square be three activity or message types. Then:

a. The rightwards zig zag transitivity holds for the consequence absence:



b. The leftwards zig zag transitivity holds for the consequence absence:



Proof 2 (Rightwards Zig Zag Transitivity of Absence).

Let \square , \square , and \square be activities or interactions. Then:



Leftwards zig zag transitivity of absence can be proven similarly by replacing ' $<$ ' with ' $>$ ' and ' \square ' with ' \square '.

In the following, we use Examples (4) and (5) from Fig. 9 to illustrate and discuss how Theorem 2 can be used to decompose a GCR of type rightwards zigzag $\square \square$. Note that these two examples are adopted from the running example we introduced in Section 1 in order to fulfill the decomposition requirements.

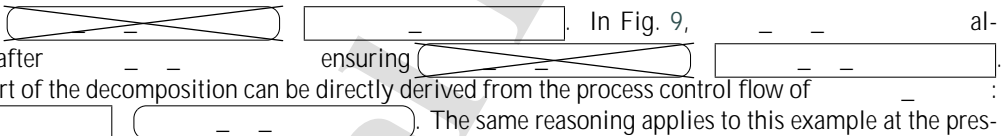
- Example (4): In this example, \square and \square belong to different XOR branches, which means that the

Verifying Compliance in Process Choreographies

execution of activity a implies the non-execution of activity b and vice versa (fulfilling assertion $\neg (a \wedge b)$). Additionally, in the interaction activity i and the private activity p belong to the same XOR branch, and fulfill assertion $\neg (i \wedge p)$. According to Theorem 2.a, the conjunction of $\neg (a \wedge b)$ and $\neg (i \wedge p)$ reproduces the behavior of the original GCR. Note that process compatibility ensures that whenever sending a occurs in P , receiving a should occur in Q as well. At the presence of loops that encapsulate the depicted process part of P , the XOR fragment can be executed multiple times possibly leading to an alternate execution of the corresponding branches. For example, if in the first loop iteration, a is executed and b is not executed, then, to this point both derived assertions are satisfied. Let us assume that a future iteration over the XOR fragment in the context of the same process instance triggers b execution, thus, violating the original GCR.

iteration 1: $P : \{ a \}$
 iteration 1: $Q : \{ \}$
 iteration 2: $P : \{ a, \text{order_special_transport} \}$
 iteration 2: $Q : \{ \text{quick_test_intermediate} \}$
 Combined trace: $\{ \text{transport_intermediate}, \text{quick_test_intermediate} \}$

By looking at the combined trace, it becomes clear that the GCR is violated. Unfortunately, this would require both partners to exchange the traces and use a common time stamping system to obtain the same chronological order of activities. Using the theorems, however, the process P can locally run its derived assertion against its own execution trace of the same process instance, and identify the violation. Note that the decomposition does not enforce the processes with new restrictions (except when no transitivity could be derived), but uses the existing control flow and interactions between partners to derive assertions that can be used for a decentralized checking of the original GCR.

- Example (5): . In Fig. 9, a always happens after b ensuring $\neg (a \wedge b)$. The second part of the decomposition can be directly derived from the process control flow of P : $P : \{ a, b, c, d, e, f, g, h, i, j, k, l, m, n, o, p, q, r, s, t, u, v, w, x, y, z \}$. The same reasoning applies to this example at the presence of loops. The correctness of Example (5) concludes from the leftwards zig zag transitivity (cf. Theorem 2.b), whereas Example (4) relies on the rightwards zig zag transitivity of the absence (cf. Theorem 2.a). The decomposition process is not limited to these scenarios and, as aforementioned, the decomposition cannot always be automated, but might require manual interaction and processing. Altogether, the decomposition eases global compliance rule checking, where each partner checks its corresponding derived assertions locally. A GCR is rechecked only if at least one of the derived assertions is not evaluated to true. Note that this does not necessarily imply that the GCR is violated.

Theorem 3 (Rightwards Chaining Transitivity).

Let a, b, c , and d be activities or interactions such as $\square, \square, \bigcirc, \bigcirc$: if a and b occur, and c shall occur afterwards. Let i_1, i_2 , and i_3 be three interactions such as:

- (1) $\bigcirc_1 \square$
- (2) $\square_1 \square \bigcirc_2$
- (3) $\square_2 \bigcirc \bigcirc_3$
- (4) $\square_3 \bigcirc$

Then: Whenever (1) (2) (3) (4) evaluates to true, $\square \square \bigcirc \bigcirc$ is true as well.


Proof 3 (Rightwards Chaining Transitivity).

- (1) (2)
- $$i_1 \quad (\cdot) \quad (i_1, i_1) \quad (i_1 <) \quad i_1, i_2 \quad ((i_1, i_1) \quad (\cdot) \quad i_1 <)$$

Verifying Compliance in Process Choreographies

$$\begin{aligned}
 & ((2, 2) < 2) \\
 & \quad 1 \quad 2 \quad (,) \quad (1, 1) \quad (1 <) \quad ((1, 1) \quad (,) \quad 1 <) \quad ((2, 2) < 2) \\
 & \quad 1 \quad 2 \quad (,) \quad (,) < \quad (1, 1) \quad (1 <) \quad ((2, 2) < 2) \\
 (1) \quad (2) \quad (3) \quad & \quad 2 \quad (,) \quad (,) < \quad (2, 2) < 2 \quad 2 \quad 3 \quad (2, 2) \quad (,) \quad (3, 3) \\
 & \quad 2 < \quad < 3 \\
 (1) \quad (2) \quad (3) \quad (4) \quad & \quad 3 \quad (,) \quad (,) < \quad (,) \quad (3, 3) < \quad < 3 \\
 & \quad 3 \quad (,) \quad (,) < \quad (,) \quad (3, 3) < \quad < 3 \quad 3 \quad (3, 3) \\
 & \quad (,) \quad 3 < \\
 & \quad (,) \quad (,) < \quad (,) \quad (,) < \quad <
 \end{aligned}$$

• Example (6): . In this example, all activities involved in the GCR are private and belong to separate partners. According to the process models shown in Figs. 3 and 4, each partner can separately derive the corresponding assertion based on Theorem 3 and involving the corresponding activity in the GCR.



In this example, partners will first engage in a setup phase, in which they agree on the interaction activities that satisfy all derived assertions following the assertions' templates of Theorem 3. For example, will start by identifying relations in its process of type , where must be a message exchange with that is the partner being responsible for the following antecedence occurrence . In this example, and have only one interaction that satisfies the derived assertion (1); however, it is also possible to identify several alternatives. The combination of the four derived assertions reproduce the behavior of the original GCR when all assertions are true.

The following theorem represents a generalization of Theorem 3 with antecedences' occurrences and consequences' occurrences. Note that the previous example also illustrates Theorem 4 with $n = 2$ and $m = 2$.

Theorem 4 (Generic Rightwards Chaining Transitivity).

Let a_1 and a_2 be + activities. ; if all occur such that $a_1 < a_2 < \dots < a_{n+1}$ holds, then all should occur afterwards such that $a_1 < a_2 < \dots < a_{n+1}$ holds.

Let a_k , where $1 < k < n+m-1$ be interactions such that:

- (1)
- (2) where $1 < k < n$
- (3)
- (4) where $1 < k < n+m$
- (5)

Verifying Compliance in Process Choreographies

Then: Whenever (1) (2) (3) (4) (5) evaluates to true, $\square_1 \dots \square_n \bigcirc_1 \dots \bigcirc_n$ is true as well.

Proof 4 (Generic Rightwards Chaining Transitivity).

- (1) $\square_1 (a, b) (c, d) (e < f)$
- (2) $\square_{1 < 2} \square_{-1} ((c, d) (e, f) (g < h)) ((a, b) (c < d))$
- (3) $\square_{-1} ((c, d) (e, f) (g < h)) ((a, b) (c < d))$
- (1) (2) (3): when evaluated to true, this ensures that all \square_{-1} were executed, and all these executions combined lead to $\square_{1 < 2}$ as a consequence. This includes the case where all \square_{-1} execute in ascending order. So, if we consider this particular order, the formula leading to $\square_{1 < 2}$ becomes true.
- (4) $\square_1 (a, b) \square_{+1} ((c, d) (e, f) (g < h)) ((i, j) (k < l)) \square_{+ < 2}$
- (5) $\square_{+ < 2} ((i, j) (k, l) (m < n)) \square_{+ < 1}$
- (4) (5): This formula transitively ensures that whenever $\square_{+ < 1}$ is executed, there is a least one execution in ascending order of all $\square_{+ < 2}$ becomes the link between all antecedence patterns and consequence patterns. Therefore, the conjunction of formulas (1) to (5) ensures that whenever an instance containing an ordered execution of $\square_{+ < 1}$ should lead to an ordered execution of execution of $\square_{+ < 2}$. Note that this conjunction represents a stronger constraint than the original GCR. However, as these formulas are deduced directly from the actual processes, they do not add new constraints.

Theorem 5 (Between Pattern 1).

Let a , b and c be three activities. $\square_1 \bigcirc_2 \square_3$: if a and b occur and c occurs after a , then b must occur in between.

- (1) $\square_1 \bigcirc_1 \bigcirc_2 \square_1$
- (2) $\bigcirc_2 \bigcirc_3 \square_1$
- (3) $\square_1 \bigcirc_3 \square_3$

Then: Whenever (1) (2) (3) evaluates to true, $\square_1 \bigcirc_2 \square_3$ is true as well.


Proof 5 (Between Pattern 1).

- (1) $\square_1 (a, b) \square_2 (c, d) \square_3 (e, f) (g < h) (i < j) (k < l)$
- (2) $\square_2 (c, d) \square_3 (e, f) \square_3 (g, h) (i < j) (k < l)$
- (3) $\square_1 \square_3 (e, f) (g, h) \square_3 (i, j) (k < l)$
- (1) (2) $\square_1 (a, b) (c, d) (e, f) (g, h) (i, j) (k < l) (m < n) (o < p) (q < r) (s < t)$
- (1) (2) (3) $\square_1 (a, b) (c, d) (e, f) (g, h) (i, j) (k < l) (m < n) (o < p) (q < r) (s < t) (u < v) (w < x) (y < z)$

- Example (7): $\square_1 \dots \square_n \bigcirc_1 \dots \bigcirc_n$. Again, in this example, we consider the worst case scenario where each activity referred to by the GCR belongs to a different process partner. In this example, \square_1 has one single alternative as it only has two interaction activities with \square_2 and \square_3 respectively, which follow the assertion template (1) of Theorem 5; i.e., each execution of \square_1 must be followed (not necessarily immediately) by \square_2 or \square_3 .

Verifying Compliance in Process Choreographies

which, in turn, should not be preceded by any $_$ execution. Similarly, $_$ and $_$ should identify assertions that follow rule templates (2) and (3) respectively. Note that the Between Pattern $\square \circ \square$ can be also checked using chaining transitivity $\square \circ \square$. However, this adds a stronger assumption on $_$ and $_$ that should follow whenever it occurs. For example, this holds in the running example (cf. Fig. 3), as $_$ transitively implies $_$, which, in turn, transitively implies $_$.



Theorem 6 (Between Pattern 2).

Let $_$, $_$ and $_$ be three activities. $\square \circ \square$: if $_$ and $_$ occur and $_$ occurs after $_$, then $_$ shall occur in between.

- (1) $\square_1 \square_2 \square_3 \square_4$
- (2) $\square_1 \square_2 \square_3 \square_4$
- (3) $\square_3 \square_5 \square_4$
- (4) $\square_2 \square_5$

Then: Whenever (1) (2) (3) (4) evaluates to true, $\square \circ \square$ is true as well.

Proof 6 (Between Pattern 2).

- (1) $_1 _2 _3 _4, (_1, _1) (_2, _2) (_3, _3) (_4, _4) (_1 < _2 < _3 < _4)$
- (2) $_1 _4, (_1, _1) (_4, _4) _2, (_2, _2) (_2 < _4)$
- (3) $_3 _5, (_3, _3) (_5, _5) (_3 < _5)$
- (4) $_2 _5, (_2, _2) (_5, _5) (_2 < _5)$

Using (1):

$$_1 _2 _3 _4, (_1, _1) (_2, _2) (_3, _3) (_4, _4) (_1 < _2 < _3 < _4) (_2 < _4) (_3 < _5) (_4 < _5) < (_1, _1) (_2, _2) (_3, _3) (_4, _4) ((_1 < _2 < _3 < _4) (_1 < _4)) ((_1 < _2 < _3 < _4) (_4 < _5))$$

Using (2), if $_4$ happens before $_2$ then $_3$ should not happen in between:

$$_1 _2 _3 _4, (_1, _1) (_2, _2) (_3, _3) (_4, _4) (_1 < _2 < _3 < _4) (_1 < _2 < _3 < _4) (_2 < _4) (_2, _2) (_3, _3) (_2 < _3 < _4)$$

Using (3), if $_3 < _5$, then there should be $_4$ in between:

$$_2 _3 _5, (_2, _2) (_3, _3) (_5, _5) (_2 < _3 < _5) (_2, _2) (_3, _3) (_5, _5) (_2 < _3 < _5 < _4) (_2, _2) (_5, _5) (_2 < _5 < _4)$$

Using (4), if there exist $_2$ and $_5$ such that $_2 < _5$, then there should be $_4$ in between:

$$_2 _5, (_2, _2) (_5, _5) (_2 < _5 < _4) (_2, _2) (_5, _5) (_2 < _5 < _4) (_2, _2) (_5, _5) (_2 < _5 < _4)$$

Verifying Compliance in Process Choreographies

$$1 < 2) \quad (2 < 3 <)$$

As there are no loops over 1 , and corresponding messages, then 2 can occur only once within one process instance. Consequently, $2 = 2$:

$$(1) \quad (2) \quad (\quad 1, \quad 2 \quad 3 \quad ((,) \quad ((,)) \quad ((1, 1) \quad (2, 2) \quad (3, 3) \quad (< 1 < 2 < 3 <))$$

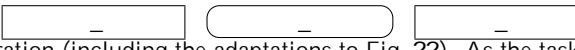
$$(\quad 1, \quad 3 \quad ((,) \quad ((,)) \quad ((1, 1) \quad (3, 3) \quad (< 1 < 3 <))$$

$$(3) \quad 1 \quad 3 \quad ((1, 1) \quad (3, 3) \quad ((1, 1) \quad (,) \quad (1 < < 3))$$

$$(1) \quad (2) \quad (3) \quad (\quad 1, \quad 3 \quad ((,) \quad ((,)) \quad ((1, 1) \quad (3, 3) \quad (< 1 < 3 <))$$

$$, \quad (,) \quad (< 1 < < 3 <)$$

$$(\quad ((,) \quad ((,)) \quad , \quad (,) \quad (< <)$$

- Example (9): . We use the same GCR as in previous theorem illustration (including the adaptations to Fig. ??). As the tasks involved in the GCR are not contained in any loop, Theorem 7 may be applied. The following assertions are the decomposition results:

$$(1) \quad \text{[Diagram: Three boxes with horizontal lines, connected by a line.]}$$

$$(2) \quad - \quad \text{[Diagram: Three boxes with horizontal lines, connected by a line, with a minus sign.]}$$

$$(4) \quad \text{[Diagram: Three boxes with horizontal lines, connected by a line.]}$$

Theorem 8 (Requires transitivity).

Let a and b be two activities or interactions such as $\square \quad \bigcirc$: if a occurs then b should occur (before or after, $(,)$, $(,)$):

Let a, b, c be three activities or interactions such that :

- (1) $\square \quad \bigcirc$
- (2) $\square \quad \bigcirc$.

If (1) (2) evaluates to true, then $\square \quad \bigcirc \quad \bigcirc \quad \square$ evaluates to true.

Proof 8 (Requires transitivity).

$$(1) \quad , \quad (,) \quad (,) \quad (<)$$

$$(2) \quad , \quad (,) \quad (,) \quad < \quad , \quad (,) \quad (,) \quad (<)$$

$$, \quad (,) \quad (,) \quad (<) \quad <)$$

$$(1) \quad (2) \quad , \quad (,) \quad (,) \quad (<) \quad , \quad (,) \quad (,) \quad (<) \quad <)$$

$$, \quad (,) \quad (,) \quad (<) \quad (,) \quad (<) \quad <)$$

$$, \quad (,) \quad (,) \quad (<) \quad (,) \quad (< <) \quad (< <) \quad (< <)$$

$$, \quad (,) \quad (,) \quad (<) \quad (<)$$

$$, \quad (,) \quad , \quad (,)$$

Note that this theorem also considers loops and multi-instance patterns. The illustration of Theorem 8 is similar to the rightwards and leftwards transitivity examples.

4. Algorithm for Decomposing Global Compliance Rules

At design time, checking a GCR that solely refers to interactions and/or public activities can be achieved by applying contemporary compliance checking techniques (cf. (29)) either on the choreography model or the public process models of the involved partners. By contrast, if a GCR refers to private activities of different partners, it becomes impossible to check its correctness as partners must not view the private process model parts of the other partners and,

Verifying Compliance in Process Choreographies

therefore, cannot identify the dependencies between the private activities involved in the GCR. To cope with this issue, we suggest decomposing the GCR into a set of assertions of which each can be checked locally by the corresponding partner. The decomposed rules then reproduce the behavior of the original GCR.

Decomposition in compliance checking has been exploited by (61), but only at the process model level in order to achieve performance gains for the compliance checks. The article at hand proposes to decompose the GCR to distribute the compliance checks to the partners for maintaining the confidentiality of their private tasks. This section focuses on the decomposition algorithms and explains the steps to derive assertions by applying the theorems introduced in Section 3.

Figure 10 provides a high-level description of the steps required by partners involved in a GCR to identify a valid decomposition. Algorithm 1 provides a more detailed view on how this can be achieved in practice with a particular focus on compliance rules that include one antecedence pattern.

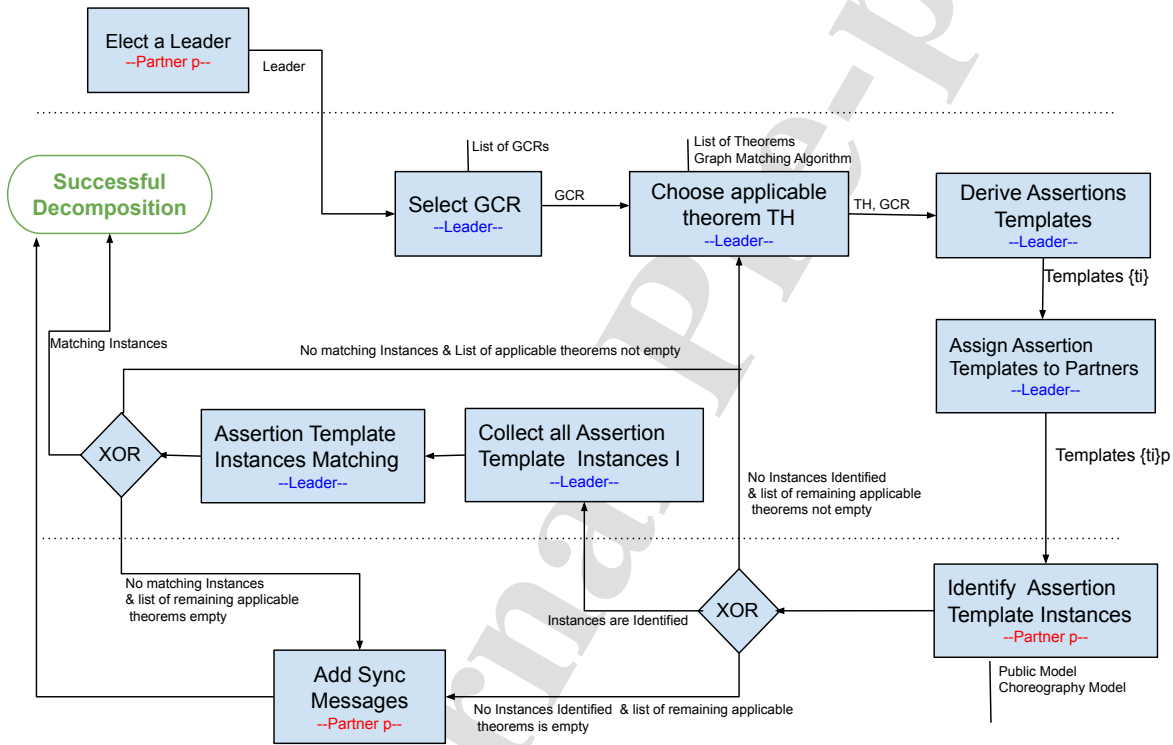


Figure 10: Decomposition Process

Given a GCR, the decomposition process starts by selecting a leader among the partners involved in the GCR. The leader will be responsible for identifying the pattern corresponding to the GCR (e.g., rightward chaining pattern or between pattern). This is trivial and can be accomplished with a simple exact graph matching algorithm (e.g., comparing node types and connectors). Once the pattern is identified, several decomposition theorems may be applicable. For example, in the case of the between pattern, Theorems 5, 6, and possibly 7 (if the processes are loop free) may be applied. The leader will then select and apply a theorem, derive the assertion templates accordingly, and send each of them to the corresponding partner—this step can be easily automated. An assertion template, in turn, is a derived assertion output from the theorem, where the actual interactions (i.e., message exchanges) have not yet been defined. For example, for $\langle \text{---} \text{---} \text{---} \rangle$ (where M shall be an activity interacting with $\langle \text{---} \text{---} \text{---} \rangle$) is a derived assertion template of the global compliance rule $\langle \text{---} \text{---} \text{---} \rangle \langle \text{---} \text{---} \text{---} \rangle$, as has not been defined yet.

Verifying Compliance in Process Choreographies

Next, each partner will try to identify an assertion template instance that conforms with the derived assertion template—an instance of an assertion template corresponds to a template with actual activities. τ is a valid option in this example as it interacts with τ fulfilling the template constraint (i.e., the template instance corresponds to τ). This can be automated by having each partner responsible for an assertion template iterating over the message exchanges in its private process model. Depending on the assertion template structure, one or several message exchanges may be selected for constructing an assertion instance candidate. The latter must conform with the assertion template. Moreover, it needs to be ensured that the process in question is compliant with it. Existing design-time compliance checking techniques can be employed in this regard (3; 4; 35; 6).

Several assertion instance candidates can be identified for the same partner, which may increase the probability of finding a collective solution among all partners. Afterwards, partners either collectively enter a negotiation phase and exchange their assertion instance candidates, or rely on the leader to collect all proposals and run a matching algorithm that selects assertion candidates, which replicate the templates derived by the decomposition. Indeed, two assertion templates may require that they use the same message exchange. Therefore, the matching algorithm will select the assertion instance candidates of different partners that have the same message exchange in common (using function match or match' to ensure that the mapping is correct). While having the leader collecting the assertion proposals and doing the matching can be more efficient and reduce the communication overhead between partners, conducting the negotiation in a distributed manner reduces trust assumptions.

Note that it is possible to run the entire process in a distributed manner, without need for a leader. In this case, all partners will have to run the matching algorithms for identifying the GCR pattern. Moreover, they have to agree on the decomposition theorem to be applied (e.g., using a majority vote) and collectively execute the matching of assertion instances. If a matching solution is found, each partner will use the selected assertion instance for future run-time checking. Unless a solution is found, the next applicable theorem will be explored in the same way. If no solution could be found after trying all applicable decomposition theorems, synchronization messages become necessary for enabling distributed run-time checking of the GCRs. At run-time, no additional communication with other partners is needed for checking the GCR, unless a violation occurs. Similar to assertion and local compliance rules monitoring, each partner is responsible for complying with the derived assertions. This can be enforced using post-auditing processes by the respective legal entities, e.g., external audits conducted by data protection officers on hospitals that participate in a collaborative study on COVID-19 vaccines' efficiency. Indeed, in the healthcare sector, new methods exist, where it is possible to conduct a research study using federated machine learning⁶. In this setting, the ML application is conducted locally within each healthcare data provider infrastructure (e.g., hospitals or bio banks), and the resulting output models are aggregated instead. This prevents data of different participants from being merged in a central repository, and thus being subject to different and complex regulatory issues. These locally executed processes, nevertheless, also need to comply with data protection rules, where, for example, the data used for the aggregated model must have patient consents beforehand. Therefore, external audits become necessary, at each site, to check whether each federated model used for the aggregation is indeed compliant with the GDPR requirements for data consents. Note that this also prevents collecting all participants' consents in one central repository for the purpose of compliance checks on the aggregated model.

Algorithm 1 realizes GCR decomposition as set out in Definition 2. It assumes that each node of the GCR is assigned to one partner being responsible for it. Further on, we assume the input GCR to be consistent and satisfiable (for dealing with unsatisfiable and inconsistent rules we refer interested readers to (11)). In the following, we first explain Algorithm 1 step by step and then illustrate the entire algorithm along Example 1 (see below).

Starting from the \square node (cf. Def. 2), Algorithm 1 walks outwards through all other nodes of the GCR. Thereby, the visited parts are copied and become assertions. Wherever the algorithm walks over a connector between two nodes α and β , which are assigned to different partners $\pi(\alpha)$ and $\pi(\beta)$, the GCR is split at this position as this dependency cannot be evaluated by a single partner. Next, the algorithm tries to replicate the connector where the GCR was split through (transitive) message exchanges between both affected partners by applying the transitive relationships from Section 2. Therefore, sets \mathcal{M}_α , \mathcal{M}_β , and $\mathcal{M}_{\alpha\beta}$ are calculated. Depending on the pattern of α (cf. Def. 2), \mathcal{M}_α and \mathcal{M}_β contain the messages succeeding or preceding α and β , respectively. Note that these calculations have to be accomplished in a decentralized manner by $\pi(\alpha)$ and $\pi(\beta)$ themselves as α and β may be private tasks. Next, $\mathcal{M}_{\alpha\beta}$ combines those messages of \mathcal{M}_α and \mathcal{M}_β that can be combined.

⁶<https://featurecloud.eu/>

Verifying Compliance in Process Choreographies

Algorithm 1: GCR decomposition DECOMPOSE(gcr)

- Global compliance rule $\mathcal{G} = (\mathcal{P}, \mathcal{M}, \mathcal{A}, \mathcal{R})$
- Choreography model \mathcal{C} , and \mathcal{P} as the set of all partners' message nodes.
- We assume that \mathcal{G} also returns the partner private model of a node p .

```

select the only  $p$  with  $\mathcal{G}(p) = \square$ 
initialize queue  $Q = \{ \}$ 
create (incomplete) Assertion  $\mathcal{A}$  for the partner associated with  $(p)$ 
foreach (  $p \in \mathcal{P}$  ) do
  foreach (  $m \in \mathcal{M}(p)$  ) do
    { }
    if (  $\mathcal{G}(m) = \square$  ) then
      // and involve the same partner
      initialize  $\mathcal{A}$  as reference on
      if (  $\mathcal{G}(m) = \otimes$  ) then extend with  $\otimes$ 
      if (  $\mathcal{G}(m) = \circ$  ) then extend with  $\circ$ 
      if (  $\mathcal{G}(m) = \otimes$  ) then extend with  $\otimes$ 
    else
      // and involve different partners
      if (  $\mathcal{G}(m) = \circ$  ) then
        {  $\mathcal{G}(m)$ ,  $\mathcal{G}(m)$ ,  $\square$ ,  $\circ$  }
        {  $\mathcal{G}(m)$ ,  $\mathcal{G}(m)$ ,  $\square$ ,  $\circ$  }
        { (  $\mathcal{G}(m)$ ,  $\mathcal{G}(m)$  )  $\square$ ,  $\circ$  }
      if (  $\mathcal{G}(m) = \otimes$  ) then
        {  $\mathcal{G}(m)$ ,  $\mathcal{G}(m)$ ,  $\circ$ ,  $\square$  }
        {  $\mathcal{G}(m)$ ,  $\mathcal{G}(m)$ ,  $\square$ ,  $\otimes$  }
        { (  $\mathcal{G}(m)$ ,  $\mathcal{G}(m)$  )  $\circ$ ,  $\square$  }
      if (  $\mathcal{G}(m) = \otimes$  ) then
        {  $\mathcal{G}(m)$ ,  $\mathcal{G}(m)$ ,  $\circ$ ,  $\square$  }
        {  $\mathcal{G}(m)$ ,  $\mathcal{G}(m)$ ,  $\square$ ,  $\otimes$  }
        { (  $\mathcal{G}(m)$ ,  $\mathcal{G}(m)$  )  $\circ$ ,  $\square$  }
      if (  $\mathcal{G}(m) = \{ \}$  ) then
        //No implicit dependency between  $p$  and
        add message between  $p$  and
        update models  $\mathcal{M}_1, \dots, \mathcal{M}_n$  and
        recalculate  $\mathcal{G}$ , and
      else
        //implicit dependency between  $p$  and exists
        select (  $\mathcal{G}(m)$ ,  $\mathcal{G}(m)$  )2
        if (  $\mathcal{G}(m) = \circ$  ) then
          extend with  $\circ$ 
          create Assertion  $\mathcal{A}$   $\square$ ,  $\circ$  for (  $p$  )
        if (  $\mathcal{G}(m) = \otimes$  ) then
          extend with  $\otimes$ 
          create Assertion  $\mathcal{A}$   $\square$ ,  $\otimes$  for (  $p$  )
        if (  $\mathcal{G}(m) = \otimes$  ) then
          create Assertion  $\mathcal{A}$   $\square$ ,  $\otimes$  for (  $p$  )
    foreach ( (  $p'$  with (  $\mathcal{G}(p')$  ) ) ) do
      //Same as for each (  $p'$  ) above
      //but with flipped directions
    foreach ( partner ) do
      foreach ( (  $p'$ ,  $\mathcal{G}(p')$  ) of partner ) do
        if (  $\mathcal{G}(p)$ ,  $\mathcal{G}(p')$  have the same  $\square$  pattern ) then
          merge  $\mathcal{A}$  and  $\mathcal{A}'$  based on the  $\square$  pattern
    foreach ( Assertion ) do
      if ( has empty  $\circ$  and  $\otimes$  patterns ) then remove

```

Verifying Compliance in Process Choreographies

If n is a \bigcirc node (i.e., n must follow m), n contains message tuples (m_1, m_2) that ensure that n is always followed by m_1 , m_1 by m_2 (unless $m_1 = m_2$), and m_2 by n . Any pair (m_1, m_2) can then be used to complement the created assertions, i.e., m_1 becomes a placeholder for n within the assertion of (m_1, m_2) , whereas m_2 replaces n for (m_1, m_2) . Regarding \otimes nodes (i.e., n must not follow m), all pairs of messages (m_1, m_2) ensure that n is preceded by m_1 in any case, and m_1 is preceded by m_2 (unless $m_1 = m_2$), whereas n never follows m_2 . Finally, for \square nodes an occurrence of n after m allows ignoring the rule. Hence, \square nodes result in pairs (m_1, m_2) such that m_1 may only occur after m and m_2 may only occur after n (unless $m_1 = m_2$). However, there should be at least one case in which m_2 is followed by n (i.e., n is not always preceded by m_2).

Finally, all assertions of the same partner, which depend on the same \square message, are merged to reduce the number of assertions. Remaining assertions without consequences are removed as they result from the processing of \otimes nodes, but have not been merged in the previous step. Remember that ignoring \otimes nodes is allowed as this makes rules even more strict.

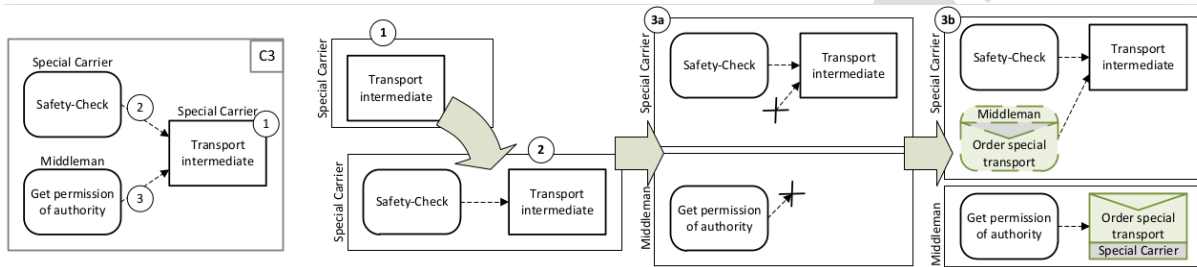


Figure 11: Decomposition process of GCR 3

Example 1. Let us apply Algorithm 1 to GCR 3 (cf. Fig. 11) from the running example that we introduced in Section 1 (cf. Figs. 1 - 4). Let the responsibilities be (Safety Check) = Special Carrier, (Get permission of authority) = Middleman, and (Transport intermediate) = Special Carrier. After assigning responsibilities, Algorithm 1 starts with \square node Transport intermediate and creates a new assertion for the Special Carrier who is responsible for this activity. Then the Safety check is discovered and added to the assertion, since it belongs to the same partner. In turn, another partner (i.e., Middleman) is responsible for activity Get permission of authority. Hence, the algorithm cuts the respective connector and creates a new assertion for the respective partner. Next, Special Carrier and Middleman determine $\mathcal{M}_1 = \{\text{Waybill, T. Details, Req. Details, Order ST}\}$ and $\mathcal{M}_2 = \{\text{Order ST}\}$ to calculate those message pairs $\mathcal{M} = \{(\text{Waybill, Order ST}), \dots, (\text{Order ST, Order ST})\}$ that can be used to transitively replicate the connector where the GCR was split. Finally, the algorithm places the selected messages into both assertions in such a way that the correctness of the original rule is preserved through the (leftwards) transitivity of eCRGs. Note that the Special Carrier could use message Waybill instead as (Waybill, Order ST) holds.

For the same GCR, it is possible to infer several decomposition alternatives, depending on which interactions are used to find a transitive control flow relationship between the nodes of the GCR. It is also possible that no direct link can be identified between two partners' GCR nodes (i.e., there is no interaction between these two partners). As such, interactions with *intermediary partners* can be used to find an indirect link (i.e., transitive interactions). As aforementioned, if no transitivity is identified between the GCR nodes of two partners (even not through intermediary partners), it becomes necessary to exchange additional execution data between the partners involved in the GCR, by, for example, adding sync messages. Sync messages are a specific type of messages communicated between partners to inform about the state of a given task (e.g., terminated, started, not executed). Although sync messages are not preferred as they expose private data about the exact execution time of a private task, they become necessary when the GCR cannot be decomposed into assertions, i.e., no transitive relations can be identified.

In the following, we discuss the complexity of the GCR decomposition in Algorithm 1. Results on checking regulatory compliance in general have been provided in (60). The first and second loops iterate over the nodes of the compliance rule. If we consider that two nodes can only have one flow connector, the number of required operations will be $\frac{-1}{2}$, otherwise (-1) . In both cases complexity corresponds to (-2) . The first statement is (1), whereas the

Verifying Compliance in Process Choreographies

statement calculates τ , σ and ρ each with a worst case complexity of $\mathcal{O}(n^2)$. The second inner loop has the same complexity as the first inner loop. The third nested inner loop iterates over partners and compare assertions within the same partner with a number of operations equal to $n \times \frac{n-1}{2}$, which gives a complexity of $\mathcal{O}(n^3)$. Finally, the last inner loop has a complexity of $\mathcal{O}(n)$. Obviously, the overall worst case complexity of the algorithm is polynomial $\mathcal{O}(n^4)$; i.e., outer loop combined with the third nested inner loop.

5. Verifying GCR Decomposition

The GCR decomposition algorithms from Section 4 are based on the theorems we have presented and proven in Section 3. Although these theorems support most control flow (i.e., behavior) compliance patterns known from literature (57; 19), they cannot cover every possible scenario. There may be two reasons for this: (i) either the structure of a GCR is not covered by Theorems 1–8 or (ii) none of the proposed decompositions is applicable. In both cases, it might not only become necessary to find a novel decomposition of a GCR, but also to verify the latter, i.e., to prove the correctness of the decomposition. One approach to accomplish this would be to apply the eCRG semantics and to formally prove correctness (cf. Section 3). However, this is far from being trivial. Therefore, we introduce Algorithm 2 that enables the automated verification of GCR decompositions based on eCRG model checking.

Algorithm 2: Verification of Decompositions (Assertions $\sigma_1, \dots, \sigma_n$, GCR ρ)

```

1 Let  $\tau$  be a function that translates an eCRG into a corresponding finite-state automaton
2    $\tau(\sigma_1)$  ( )
3    $\tau(\sigma_n)$  ( )
4 if  $\tau(\rho) = \tau(\sigma_1) \cap \dots \cap \tau(\sigma_n)$  then
5   //Decomposition is correct
6   return true
7 else
8   //Decomposition is incorrect
9   return any arbitrary trace through  $\tau(\rho)$  as counter example.
```

The main idea of Algorithm 2 is to interpret a GCR decomposition as declarative process model and to verify whether it solely allows for execution traces that comply with the original GCR. Thus, techniques that are known from the verification of declarative process models (51; 53) can be applied: First, all assertions $\sigma_1, \dots, \sigma_n$ of the decomposition are translated into finite state automatons $\tau(\sigma_1), \dots, \tau(\sigma_n)$. Their intersection $\tau(\sigma_1) \cap \dots \cap \tau(\sigma_n)$ corresponds to an automaton that only accepts traces that comply with every assertion. In turn, $\tau(\rho)$ denotes the automaton that accepts solely traces violating the original GCR. If the intersection of these two automatons is empty, the decomposition is correct as it only allows for traces that do not violate the original GCR and, thus, comply with it.

$$\tau(\sigma_1) \cap \dots \cap \tau(\sigma_n) = \tau(\rho) \cap \tau(\sigma_1) \cap \dots \cap \tau(\sigma_n)$$

For any choreography y , whose partners ensure $\sigma_1, \dots, \sigma_n$, we can now directly conclude:

$$\tau(y) \cap \tau(\sigma_1) \cap \dots \cap \tau(\sigma_n) = \tau(y) \cap \tau(\rho) \cap \tau(\sigma_1) \cap \dots \cap \tau(\sigma_n), \text{ i.e., } y \text{ complies with } \rho$$

6. Implementation

The presented approach is implemented as part of the C³Pro framework⁷, which deals with change and compliance in process choreographies (21). The framework provides sophisticated functions for defining, propagating and negotiating changes in the context of process choreographies. Furthermore, it comprises a modeling component as one of its core components for editing and changing public and private process models as well as process choreography models. This component further enables the visualisation of change propagations. In the context of the present work, the three-layer architecture of the framework (21) (i.e., process modeling, change, and execution) was extended with additional

⁷<http://www.wst.uniwi.ac.at/communities/c3pro/>

Verifying Compliance in Process Choreographies

components for dealing with **process compliance**. In detail, these new components include (i) an eCRG modelling tool, (ii) an automated generator of compliant choreographies, (iii) a model checker, and (iv) a GCR decomposition tool.

Figure 12 depicts the main components of the C³Pro framework. The **compliance (CME) and process modeling (PME) environments** allow defining and editing *compliant* process choreography models (39; 35) and decomposing global compliance rules. *Compliance* was introduced as “a semantic correctness criterion to be considered when designing interaction models. It ensures that interaction models do not conflict with the set of imposed global compliance rules” (39). At design time, it is ensured that the created choreography models are compliant with the defined compliance rules. Using both PME and CME, it becomes possible to parameterize and automatically generate compliant choreographies, which then can be used for testing and simulation purposes. A user, therefore, can specify the number and patterns of the GCRs as well as the parameters of the process and choreography models (e.g., number of private and public tasks, number of partners and interactions, or number and types of the control flow patterns that shall be covered by the processes) (6). Although the generated models represent synthetic processes without real-world semantics (i.e., these models do not reflect actual use cases such as a manufacturing collaborative process), they may serve as a support for simulation and research work evaluation, e.g., model executions can result in distributed logs of synthetic data, which are useful for evaluating the efficiency of specific mining techniques. In the context of this work, this tool can be used to test the feasibility and applicability of the decomposition process on more complex choreographies and corresponding GCRs. Currently, the automated generator tool only supports four basic compliance patterns. However, other GCR patterns can be directly inferred from the models and be used for testing. The tool is integrated in the C³Pro framework and can be tested. A data set of automatically generated models and the corresponding GCRs are made available in the same repository.⁷ Finally, the change editor allows defining and editing changes of process models and compliance rules, respectively.

The **Compliance Management Service** represents the *main extension* to C3Pro related to this work, and handles the defined compliance rules and implements the theorems as well as the GCR decomposition algorithm (cf. Section 4). As process execution engine, the Cloud Process Execution Engine⁸ is utilized. Most functions of the C³Pro framework are provided as a RESTful service, which enables unified access from any client being able to communicate via HTTP. Finally, the Compliance Management Service serves as a pluggable middleware that may be used to integrate other process execution engines.

For testing the framework, we edited BPMN 2.0 choreography and collaboration models using Signavio⁹ and exported them to the C³Pro framework as XML files. Examples are extended with GCRs, which are then decomposed into derived assertions using Algorithm 1. To this endeavor, mainly the CME, PME and the compliance management service were used.

In addition, the C³Pro framework was *extended with* a novel eCRG model checker that was published on github¹⁰. Its command line interface enables the specification and verification of both global and local compliance rules (GCR and LCR) as well as process models and choreographies. In particular, the eCRG model checker supports the verification of

- GCR decompositions, i.e., it allows verifying whether GCRs can be concluded from a given decomposition,
- local compliance, i.e., it allows verifying whether a single process model complies with a given compliance rule (CR), and
- global compliance, i.e., it allows verifying whether a process choreography complies with a given GCR.

Moreover, the eCRG checker enables the automated decomposition of tree-structured GCRs with a single antecedence.

In order to verify GCR decompositions and the various kinds of compliance respectively, the eCRG model checker translates global and local compliance rules as well as process models and choreographies into automaton, which are then combined and intersected. Depending on whether the resulting automaton is empty or not, the verification has

⁸<http://cpee.org>

⁹<http://academic.signavio.com/>

¹⁰<https://github.com/daviddknupl/esch/SCV>

Verifying Compliance in Process Choreographies

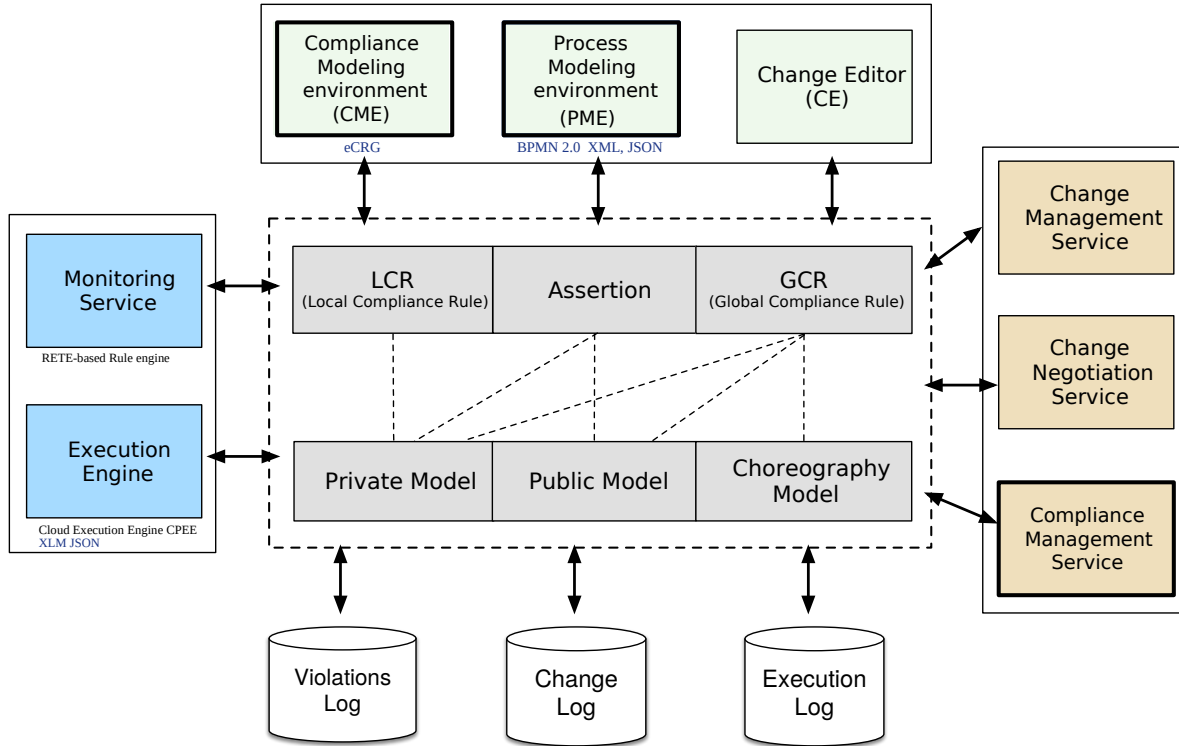


Figure 12: C3Pro Prototype Architecture

been successful or a counter example trace has been produced. The eCRG model checker has been written in Java 8 on top of the dk.brics.automaton framework¹¹.

7. Applicability

This section demonstrates and discusses how the presented decomposition algorithm can be applied in real-world settings. Here, a GCR may be imposed on process choreographies by *external* sources, e.g., considering regulatory documents such as the GDPR or standards such as ISO norms (64). A GCR may also reflect *internal* compliance rules expressing, for example, implicit dependencies between the partners that are crucial for (re-)scheduling the process activities for one partner or across multiple partners in the choreography. The visibility of activities and compliance rules in real-world settings depends on the contractual situation and the roles of the partners in the choreography. In supply chains in automotive domain, for example, an Original Equipment Manufacturer (OEM) might demand insights into certain specifics of the private processes of its suppliers and the connected (internal) compliance rules.

7.1. Manufacturing Use Case

In the supply chain example presented in the Section 1, GCR 3 (cf. Fig. 2) reflects an externally imposed GCR on safety in manufacturing and logistics processes. GCR 1, in turn, might reflect an internal quality assurance rule that is solely verified by the *Manufacturer*, but is also made visible to the other partners in order to, e.g., create trust. The real-world use case from manufacturing depicted in Fig. 13 demonstrates how the decomposition algorithms can be employed to lift implicit (data) connections to explicitly modeled assertions. The use case covers a process choreography between Partner 1 (i.e., a car manufacturer), Partner 2 (doing injection molding), and Partner 3 (i.e., the electro plater that coats parts in a correct color scheme). The choreography is designed and implemented using the CPEE (Cloud Process Execution Engine)¹². Figure 13 shows the public activities of all partners, e.g., activity pl ace

¹¹<https://www.brics.dk/automaton/>

¹²<https://cpee.org/>

Verifying Compliance in Process Choreographies

order for Partner 1 and private activities, e.g., activity wait for order completion for Partner 1. The public activities send or receive messages, e.g., activity place order sends a message received by activity receive order. Note that the scenario abstracts from the details of the public and private activities, which are modelled as sub-processes activities, except for the electroplating task where the corresponding sub-process is depicted. The sub-process describes the measuring of the quality of the bath and the glossiness; both measures are then forwarded to the partner. Depending on the measures bath maintenance is conducted (alternative default branch). All sub-processes are of different complexity, i.e., they might contain decisions and loops as well. The complex activity wait for order completion, in particular, comprises a set of sub-activities and is signifying the scheduling between the activities of the different partners.

During the design of the choreography the partners specified implicit connections, i.e., dependencies between (private) activities that are not covered by message exchanges and express mostly data dependencies. For example, activity wait for order completion (private, Partner 1) implicitly depends on the data produced by activities prepare for manufacturing, manufacturing of parts, quality control (all private, Partner 2).

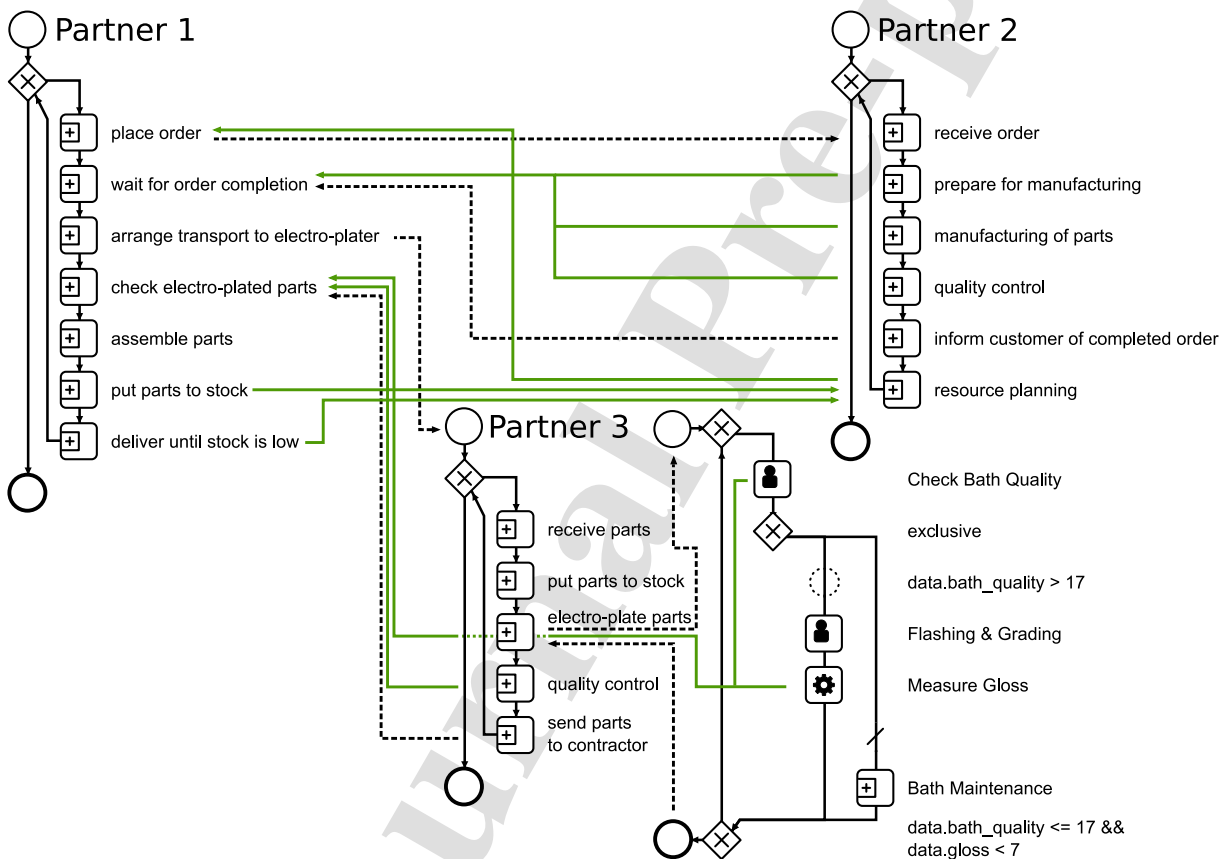


Figure 13: Manufacturing Use Case

These implicit connections refer to two main “functions” of the manufacturing setting, i.e., (i) *resource planning* and (ii) *quality control*.

(i) Proper resource planning is part of the contract between the partners. There are implicit rules regarding how fast Partner 2 has to react to an order. This depends on assumptions how fast the stock drops for Partner 1. This manifests as follows: If activity place order (Partner 1) occurs, then activity resource planning Partner 2 has been done before (i.e., resource planning data received) and activity resource planning (Partner 2) was

Verifying Compliance in Process Choreographies

based on activities put parts to stock and deliver until stock is low (Partner 1). Understanding this as a compliance task, we can say that when the above activity information matches, the compliant ordering can be ensured.

The question is how such implicit connections can be checked without revealing information on the private activities. Here, the presented decomposition algorithm can help. The idea is to express the implicit connections by GCR and verify them based on assertions. Take, for example, the implicit connection between private activities place order (Partner 1) and resource planning (Partner 2). We can formulate this implicit connection as GCR

Using Algorithm 1, \mathcal{G}_1 can be decomposed into the following assertions:

- \mathcal{A}_1 : $\text{place_order} \rightarrow \text{resource_planning}$
- \mathcal{A}_2 : $\text{resource_planning} \rightarrow \text{place_order}$

Note that doing so, the implicit connection is lifted up to an explicit one by sending message m_1 .

(ii) Consider now the more complex GCR covering the overall quality control that involves all three partners. Partner 1 has to do the final inspection, Partner 2 has to ensure the quality of the injection molded parts (no cracks, no holes), and Partner 3 has to ensure the quality.

It is assumed that data for checking quality individually has been delivered whenever an order is finished. In this case, activity wait for order completion by Partner 1 yields all information about the quality of the injection molded parts and the electroplating process, whereas during activity check electro-plated parts by Partner 1 all additional information about the molding process is available. Overall, if activity check electro-plated parts (Partner 1) occurs, then it has to be checked whether or not activities electro-plate parts and quality control by Partner 3 were executed. Before that, for activity wait for order completion, activities prepare for manufacturing (e.g. machine calibration data, material information), manufacturing of parts, and quality control by Partner 2 were executed. If all of the above information has been received, the manufacturing process was compliant, i.e., all required steps seem to have taken place.

7.2. Healthcare Use Case: FeatureCloud¹³

In the healthcare sector, data is usually distributed over multiple hospitals or biobanks. As such, conducting a clinical trial requires all data to be collected and accessed centrally, e.g., in a central cloud. The advent of federated machine learning (ML) has overcome this issue by providing means for third parties (e.g., research institutions) to conduct ML studies while ensuring that the data remains legally and technically within the data provider infrastructures, and where locally learned models are aggregated instead of the actual data (46). While, in theory, this significantly improves patient data privacy, in practice it provides no guarantees that the local processes and the locally learned models comply with the overall global policies. To deal with this, the local process of each party (e.g., hospital) can be equipped with assertions that help to facilitate audits and consent tracking. In such a setting with multiple stakeholders, multiple global compliance rules are put in place. The following GCRs are real examples taken from the FeatureCloud project in the healthcare domain (46). The example includes multiple processes related to participants in a study (e.g., hospitals), the study coordinator (e.g., pharmaceutical company or university), patients (e.g., process for managing consents), and legal entities (e.g., for approving studies):

- Before a participant (hospital) executes a clinical trial (participant process) upon invitation by the coordinator (e.g., research entity), the trial has to be approved by a legal health authority, e.g., ethics committee (approval process of the authority).
- After a patient revokes a consent (i.e., patient process), the hospital must delete the information that the latter had given a consent to use his or her data (i.e., hospital process).
- After inviting a participant for a study (i.e., by the study coordinator), a hospital has to collect patient consents specific to that type of study (e.g., breast cancer) unless board consent had been given before.

¹³<https://featurecloud.eu/>

Verifying Compliance in Process Choreographies

- Before model aggregation by the study coordinator takes place (e.g., research entity process), each participant (hospital process) has to commit the inputs/ outputs of the locally trained model to an immutable storage.

Note that most of the tasks involved in these GCRs are private, and federated ML prevents the data from being accessible or from being moved to other parties, which makes it impossible for the study coordinator, for example, to check whether or not the hospital has already collected consents after being invited. Moreover, it is not possible for a patient to check whether the hospital had really deleted the data after being revoked. Additionally, it is possible that the study coordinator requests approval from a legal authority, but continues with executing the federated ML process on patient data before receiving the approval. Therefore, decomposing such GCRs into local assertions that have to be checked and audited locally within each process is required as no global or shared execution logs exist that enable the monitoring of GCR violations. The motivation behind GCR decomposition is therefore not solely restricted to improving performance by minimising message exchanges, but also to deal with the impossibility of creating shared logs due to privacy concerns. Further assume that it is possible to create a shared log that collects information about the execution times of private tasks of different partners composing a GCR (to centrally monitor GCRs). Then, the more GCRs a partner is involved in, the more information about its overall process can be learned by other parties. Finally, in the case where privacy is not a concern, GCR decomposition can drastically enhance performance by reducing the number of additional messages that need to be exchanged between partners when relying on existing message exchanges. This becomes highly interesting when the number of process instances is very high (e.g., car manufacturing, supply chain management) as the GCR then needs to be checked for each instance. For example, in the automotive domain (cf. Section 7.1), new reports¹⁴ estimate the production of automotive brake pads to reach more than 1.3 billion OEM per year by 2026. If we consider that each OEM manufacturing corresponds to a process instance that has to comply with at least one GCR, the number of required message exchanges needed to check the GCR (without decomposition) will at least match the same number.

7.3. Discussion

In the introduction, the following research question was stated:

RQ: How to verify GCRs in a decentralized setting of a process choreography where no central coordinator with complete knowledge on the private and public tasks of all partners exists?

In the following, we discuss how far the article at hand has addressed RQ and which open questions still remain. For this purpose, we sketch the end-to-end application of the decomposition algorithm along the following steps:

1. Check whether the GCR can be verified at choreography level, i.e., solely referring to interactions.
2. Check whether the GCR can be verified at the public process of one partner.
3. Check whether the GCR can be verified on the public processes of at least two partners: partners have to check for the absence/presence of GCR-related activities and activity orders based on choreography and interactions if possible. Otherwise, verification has to be postponed to the runtime by additionally synchronizing activities OR compliance has to be verified in an ex post way based on logs if available.
4. Otherwise: The GCR refers to the private processes of one or several partners and a decomposition has to be applied.

As private activities and their dependencies are not visible to all collaborators, parts of the decomposition algorithm are executed locally by all partners involved in this GCR in order to identify possible transitive relations between their corresponding private activities and possible public activities, or interactions that replicate the connector where the GCR was split. This results in multiple derived assertion alternatives, which are then aggregated to alternatives from other partners in order to find a combination that recreates the original rule as described in Section 4 (cf. Example 1). Once the GCR is decomposed and the corresponding assertions are derived, each partner locally checks its derived assertions at runtime.

Overall, RQ has been addressed in breaking down the problem of GCR verification on distributed processes into the steps outlined above. Moreover, a sophisticated decomposition algorithm for GCR that refers to private activities of one or multiple partners has been provided. This enables distributed compliance verification for the supply chain and manufacturing use cases discussed in Section 7.1.

¹⁴<https://www.gminsights.com/industry-analysis/automotive-brake-pads-market>

Verifying Compliance in Process Choreographies

Limitations and open questions: This paper focuses on structural compliance, i.e., a GCR solely refers to control flow patterns. Compliance patterns that deal with, for example, data and resources (cf. (62)) are future work and will add substantially to the complexity of the approach. Moreover, we have applied our approach to two use cases from the manufacturing and healthcare domains. However, additional studies in other domains, such as supply chains and logistics, become necessary to evaluate the generalizability and broad applicability of the approach. In addition, the presence of *XOR* branches in the processes (where sending of messages on these branches is optional) does not affect the correctness of the decomposition as long as the processes are *compliant* with the original GCR (39). Indeed, we assume the soundness of the different process models (i.e., consistency and compatibility) and their compliance to the original GCR. This means that original GCRs are correctly specified, and the decomposition enables their checking in a distributed way. In this case, transitivity ensures correct decomposition of GCR even at the presence of *XOR* branches. If no transitive relations are identified, sync messages are required. Further on, in the end-to-end approach outlined above, Step 3 still offers the challenge on how to check GCR on public processes of multiple partners.

As a limitation of the evaluation, the proposed approach has not yet been integrated with off-the-shelf information systems (e.g., a manufacturing execution system, supply management systems) and, thus, has not been deployed to real-world environments and been evaluated over a longer period of time. Due to this missing practical use in a real setting, however, it is difficult to evaluate or estimate the amount of efforts (e.g. in terms of time or money) that is spent on manually checking compliance in process choreographies and to compare it with the effort required with the proposed approach. We have insights into other real-world use cases (e.g., change management in automotive supply chains, cross-organizational patient treatment processes) where these efforts are huge according to the parties involved. However, to be able to quantitatively assess these manual efforts and to compare them with the ones of an automated approach, an empirical study is required that systematically assesses use cases from various application domains. Note that this constitutes a research endeavor on its own. Obviously, any approach automating compliance verification in a decentralized setting offers benefits compared to the manual efforts and audits that would be needed instead. From a qualitative perspective the benefits for enterprises, which are enabled by the proposed approach, are as follows:

- Privacy. The proposed approach prevents a partner from gaining insights into the private tasks of other partners, e.g., there is no need for exchanging messages to tell other partners when own private tasks are executed.
- Decentralization. There is no need for a shared common log where partners collect all messages / events related to the execution of tasks involved in the GCR. Instead compliance checking is accomplished in a decentralized fashion.
- Minimizing the message exchange overhead. No additional message exchanges are required for checking a GCR.

8. Related Work

The work presented in this article can be positioned at the interface between process choreographies and business process compliance. Section 8.1 summarizes basic works from these two research fields, whereas Section 8.2 discusses approaches that address issues at the interface between them.

8.1. Basic Research Fields

Section 8.1.1 gives backgrounds on process choreography research, whereas Section 8.1.2 summarizes basic works dealing with business process compliance.

8.1.1. Process Choreography

Process choreography research has mainly dealt with the modeling of process choreographies and the verification of correctness properties. For this purpose, specific choreography modeling languages like *Let's dance*, *Interaction Petri nets*, and *BPMN choreography diagrams* are proposed, which support the modeling of collaborative process behavior. A particular focus of existing works has been put on correctness properties of choreography models (e.g., realizability), which have been intensively studied in literature (17; 16)—for an overview we refer interested readers to (63).

8.1.2. Business Process Compliance

Business process compliance, in turn, has been investigated for more than a decade and several surveys exist (e.g., (4; 24)). Contemporary approaches have focused on compliance rule languages, including visual notations (3; 40; 38), logic-based formalisms (45; 43), and Event Calculus (50). Moreover, several approaches enable compliance checking

Verifying Compliance in Process Choreographies

at both build- and run-time (e.g., (4; 44; 40)) or cover different process perspectives of compliance rule checking such as behavior, data, time, and resources (59; 40). Finally, characteristic patterns for business process compliance are proposed by (19).

A formal approach that verifies local process behavior (i.e., WS-BPEL process models) against legal constraints, specified in terms of the *Compliance Request Language*, is proposed by (19). This approach and similar works focus on local compliance rules, which can be checked for a given (partner) process model. By contrast, we consider verifying global compliance rules (GCR) in a process choreography based on their correct and lossless decomposition into assertions that can be checked locally by each concerned partner.

Another related field deals with checking compliance of a (local) process model against its *refinement* or *implementation*. An approach that enables checking compliance of a (local) process model against its refinement is presented in (54). More specifically, this approach deals with the automated verification of lower-level against higher-level UML activity charts. Behavioural containment is established to ensure that a lower-level chart constitutes a valid refinement of the higher-level one. Similarly, (18) presents an approach for enforcing compliance of hierarchical business processes with visually specified security constraints. An approach that enables checking compliance of a (local) process model against its implementation is presented in (8), which derives the specification of a web application from a (local) process model followed by the verification of web execution logs against derived LTL formulae. Although the problem addressed by these approaches is different from the one considered in our paper, the techniques could be of interest for global compliance checking in choreographies as well.

8.2. Interface between Process Choreography and Process Compliance

There exist several approaches that address issues at the interface between process choreography and process compliance. Section 8.2.1 discusses centralized and distributed approaches for checking compliance in multi-party processes (i.e., process choreographies). In turn, approaches that map global contracts (i.e., sets of global compliance rules) to compliant process choreographies are presented in Section 8.2.2. The conformance between process choreography and local partner processes are considered in Section 8.2.3. Finally, issues related to business process compliance in the context of dynamically evolving business partner networks are discussed in Section 8.2.4.

8.2.1. Compliance Checking in a Process Choreography – Centralized vs. Distributed Approaches

Compliance checking mechanisms assuming a trusted party are proposed by (27). In (26) the same authors present a service-oriented approach that relies on a central *integration platform* in order to enable cross-organizational service interactions, while at the same time meeting global compliance rules. As opposed to our work, this approach relies on a central component (i.e., the integration platform) to ensure that global compliance rules can be checked.

(33) advocates compliance checking in a distributed process (i.e., process choreography) as crucial, but it cannot be assessed in how far the approach deals with the restricted visibility and availability of process information as we do. In prior work, we have introduced the criterion of *compliance* (39) that captures the ability of a choreography to comply with a given set of compliance rules at all and how to check it (35). The approach presented in (37) enables checking the effects of changes on the compliance in process choreographies based on dependency graphs between global and local compliance rules as well as assertions. Finally, (22) provides an overview on the challenges, related approaches, and possible solutions at the interplay of compliance, change, and choreographies.

Distributed approaches that rely on IoT technology are proposed by (52; 49). The approach suggested by (52) considers the flow of physical objects between the parties of a multi-party process. In particular, this approach exploits the sensing capabilities of smart devices to improve process compliance checking. For this purpose, commitments define mutual contractual relationships between parties in a declarative way and drive the configuration of smart devices, which check their satisfaction and, in case of misalignment, inform the concerned parties timely.

This multi-party process compliance monitoring approach is conceptually enhanced by (49) through IoT-enabled artifacts. This approach proposes a decentralized solution switching from control- to artifact-based monitoring, where physical objects can monitor their own conditions as well as the process activities in which they participate, i.e., compliance monitoring is distributed among the physical artifacts interacting with the global process. To instruct these smart objects, BPMN process models are translated into a set of artifact-centric process models, rendered in Extended-Guard-Stage-Milestone (E-GSM) notation. In particular, this work shows that artifact-based modeling approaches have a high potential in respect to multi-party process management involving physical objects, which has not been the focus of our work.

Finally, (47) discusses legal, organizational, human-centered, technical and economic challenges to be tackled in the

Verifying Compliance in Process Choreographies

field of business process compliance when enacting the (cross-organizational) business processes on the Ethereum blockchain. For example, at the implementation level, the immutability of illegal content or the error-proneness and zero-defect tolerance of smart contracts raise challenging issues in this context. Although this work does not deal with a concrete compliance verification approach, it indicates directions for future research on process compliance when using blockchain infrastructures for enacting multi-party business processes.

8.2.2. Mapping Global Contracts to Process Choreographies

Contract languages allow specifying obligations, permissions and prohibitions in business contracts, which govern the interactions between business partners. (28) provides means to model corresponding contract constraints. An early approach that extends choreographies with such contract constraints is provided by (5). This approach transforms the contract constraints into expressions of a choreography language, i.e., contract terms are translated into choreography expressions that govern the global process (i.e., choreography) to ensure compliance. In particular, it is shown how cross-organizational business processes can be monitored and enforced according to business contract specifications through the transformation of the contract definition to constraints on (global) process behavior. However, this approach is less powerful than ours as it tightly couples compliance constraints with choreography models, which aggravates the evolution of both choreography model and contract constraints. Besides, this approach does not consider local compliance checking (i.e., locally checking assertions derived from the decomposition of global compliance rules), which limits its applicability at the presence of more complex compliance requirements.

(31) advocates Dynamic Condition Response (DCR) Graphs for decomposing global contracts into local processes. More precisely, (31) shows how a timed DCR Graph can be used to describe the global contract for a timed multi-party process (i.e., choreography), which can then be distributed as a network of communicating timed DCR Graphs (i.e., local processes) describing the local contract for each party. As opposed to our work, this approach relies on a declarative process modeling approach with a focus on discrete time deadlines.

8.2.3. Conformance between Process Choreography and Local Partner Processes

Several proposals have been made to ensure *conformance* between choreography (i.e., the global process) and the local processes of the involved business partners. In (1), conformance checking of the event logs of local processes against a given choreography model is addressed. As such an event log is not available at design time, (43) relies on a graph transformation tool—GROOVE (GRaphs for Object-Oriented VERification)—to automatically verify that a local process of a partner involved in a choreography complies with the globally specified behavior of that choreography. LTL semantics of the choreography model is employed and token semantics of the local process model, which is expressed in terms of a BPMN collaboration diagram, is used to verify conformance.

(9) relates the theory of contracts with the notion of *choreography conformance*, i.e., it is checked whether an aggregation of services correctly behaves according to a high level specification of their possible conversations. For this purpose, a method based on the composition of choreography projection and contract refinement is presented that allows verifying that a service with a given contract can correctly play a specific role within a choreography. Finally, (14) presents an approach for ensuring conformance between a set of BPMN collaboration diagrams (i.e., local process models) and a BPMN choreography diagram (i.e., choreography model).

As opposed to these approaches, our decomposition-based method verifies the compliance of a choreography model with global compliance rules and regulations that cover multiple process perspectives. However, conformance between the choreography and the participating partner processes can be considered as a prerequisite of our approach.

8.2.4. Ensuring Compliance in Dynamically Evolving Partner Networks

(65) assumes that the partners in a business networks try to provide wrong information and, hence, introduce the notion of accountability. Compliance requirements also need to be met in dynamic business networks (12; 13). In such a network, the partners may join and leave the collaboration dynamically and tasks over which compliance rules may be specified then have to be distributed or delegated to new partners or be backsourced by network participants in order to avoid compliance issues. In (12; 13), a conceptual model for aligning the compliance requirements in a business network with the monitoring requirements they induce on the network participants, particularly when the network changes or evolves, is presented. Various techniques (e.g., task delegation and in-house back-sourcing) for ensuring the consistency between the monitoring and compliance requirements as well as metrics for evaluating the status of a collaboration in respect to compliance monitorability are discussed. Obviously, this approach lacks a process perspective, but is complementary to our work with a focus on business network changes and their effects on compliance

requirements.

9. Conclusions

The interplay between interoperability and business process compliance poses a tremendous challenge on companies. In this problem space, the work at hand addresses the question on how to verify global compliance rules (GCR), i.e., rules that span multiple partners in a multi-party process (i.e., process choreography), in a decentralized manner where certain tasks of one partner process might not be visible to the other partners due to confidentiality reasons. The presented approach focuses on the decomposition of a GCR such that the decomposed parts, i.e., assertions, can be checked by the partners locally. Consequently, compliance verification is shifted from a global to a decentralized manner.

The presented decomposition approach addresses research question *RQ*, which we introduced in Section 1. In particular, the presented decomposition algorithm exploits transitivity properties of the GCR for finding the correct decompositions. The correctness is formally proven. Moreover, the complexity of the decomposition algorithm is formally analyzed and also illustrated based on specific scenarios. The feasibility of the decomposition algorithm is shown based on a prototypical implementation, including a model checker for ensuring decomposition correctness. The applicability of the approach is demonstrated through a use case from the manufacturing domain.

Future work targets at two directions, (1) GCR language and (2) applications. First, we want to extend the decomposition based on the eCRG formalism, as used in this work, to arbitrary GCR and adapt the decomposition to other compliance rule languages such as Declare (56), PENELOPE (25), or BPMN-Q (3). Second, we plan to further investigate the presented use cases in manufacturing and healthcare and to identify additional use cases, e.g., in logistics.

Overall, the presented approach provides a fundamental brick in enabling process collaborations between different partners by infringing neither their privacy nor any regulations.

Acknowledgment

This work is jointly supported by (i) the competence center SBA Research (SBA-K1) funded within the framework of COMET Competence Centers for Excellent Technologies by BMVIT, BMDW, and the federal state of Vienna, managed by the FFG, (ii) The FFG ICT of the Future project 874019 dIdentity & dApps, (iii) the "Austrian Competence Center for Digital Production" (CDP) under contract number 85418, and (iv) the European Union's Horizon 2020 research and innovation programme under grant agreement No 826078 (FeatureCloud). We thank Juergen Mangler for his input on the manufacturing use case.

References

- [1] van der Aalst, W.M.P. and Weske, M.: Reflections on a decade of interorganizational workflow research. In: Bubenko, J., Krogstie, J., Pastor, O., Pernici, B., Rolland, C., Sølvberg, A. (eds.) *Seminal Contributions to Information Systems Engineering: 25 Years of CAiSE*. p. 307â313. Springer (2013)
- [2] van der Aalst, W., Lohmann, N., Massuthe, P., Stahl, C., Wolf, K.: From public views to private views – Correctness-by-Design for services. In: *WS-FM'08*. pp. 139–153 (2008)
- [3] Awad, A., Weidlich, M., Weske, M.: Visually specifying compliance rules and explaining their violations for business processes. *Journal of Visual Languages & Computing* 22(1), 30–55 (2011)
- [4] Becker, J., Delfmann, P., Eggert, M., Schwittay, S.: Generalizability and applicability of model-based business process compliance-checking approaches - a state-of-the-art analysis and research roadmap. *Bus Res* 5(2), 221–247 (2012)
- [5] Berry, A., Milosevic, Z.: Extending choreography with business contract constraints. *Int'l Journal on Cooperative Information Systems* 14(2-3), 131–179 (2005)
- [6] Bischof, F., Fdhila, W., Rinderle-Ma, S.: Generation and transformation of compliant process collaboration models to BPMN. In: *Advanced Information Systems Engineering*. pp. 462–478 (2019)
- [7] Borkowski, M., Fdhila, W., Nardelli, M., Rinderle-Ma, S., Schulte, S.: Event-based failure prediction in distributed business processes. *Inf. Syst.* 81, 220–235 (2019)
- [8] Brambilla, M., Deutsch, A., Sui, L., Vianu, V.: The role of visual tools in a web application design and verification framework: A visual notation for IIT formulae. In: Lowe, D., Gaedke, M. (eds.) *Web Engineering*. pp. 557–568. Springer (2005)
- [9] Bravetti, M., Zavattaro, G.: Foundational aspects of contract compliance and choreography conformance. In: *The Rise and Rise of the Declarative Datacentre*. Microsoft Research Cambridge, UK (2018)
- [10] Ciccio, C.D., van der Aa, H., Cabanillas, C., Mendling, J., Prescher, J.: Detecting flight trajectory anomalies and predicting diversions in freight transportation. *Decis. Support Syst.* 88, 1–17 (2016)

Verifying Compliance in Process Choreographies

- [11] Ciccio, C.D., Maggi, F.M., Montali, M., Mendling, J.: Resolving inconsistencies and redundancies in declarative process models. *Inf. Syst.* 64, 425–446 (2017)
- [12] Comuzzi, M.: Aligning monitoring and compliance requirements in evolving business networks. In: *On the Move to Meaningful Internet Systems: OTM 2014 Conferences*. pp. 166–183. Springer (2014)
- [13] Comuzzi, M.: Alignment of process compliance and monitoring requirements in dynamic business collaborations. *Enterprise Information Systems* 11(6), 884–908 (2017)
- [14] Corradini, F., Morichetta, A., Polini, A., Re, B., Tiezzi, F.: Collaboration vs. choreography conformance in bpmn 2.0: From theory to practice. In: *IEEE 22nd International Enterprise Distributed Object Computing Conference (EDOC 2018)*. pp. 95–104 (2018)
- [15] Decker, G., Weske, M.: Behavioral consistency for B2B process integration. In: *CAiSE'07*. pp. 81–95 (2007)
- [16] Decker, G., Weske, M.: Interaction-centric modeling of process choreographies. *Information Systems* 36(2), 292–312 (2011)
- [17] Decker, G., Weske, M.: Local enforceability in interaction petri nets. In: *5th Int Conference on Business Process Management (BPM 2007)*. pp. 305–319. Springer (2007)
- [18] Duan, L., Zhang, Y., Sun, C., et al.: Enforcing compliance of hierarchical business process with visual security constraints. *Int J Syst Assur Eng Manag* 95(1), 703–715 (2018)
- [19] Elgammal, A., Tureken, O., van den Heuvel, W.J., Papazoglou, M.: Formalizing and applying compliance patterns for business process compliance. *Software and Systems Modeling* 15(1), 119–146 (2016)
- [20] Fdhila, W., Gall, M., Rinderle-Ma, S., Mangler, J., Indiono, C.: Classification and formalization of instance-spanning constraints in process-driven applications. In: *BPM*. pp. 348–364 (2016)
- [21] Fdhila, W., Indiono, C., Rinderle-Ma, S., Reichert, M.: Dealing with change in process choreographies: Design and implementation of propagation algorithms. *Inf Sys* 49, 1 – 24 (2015)
- [22] Fdhila, W., Rinderle-Ma, S., Knuplesch, D., Reichert, M.: Change and compliance in collaborative processes. In: *SCC*. pp. 162–169 (2015)
- [23] Fdhila, W., Rinderle-Ma, S., Knuplesch, D., Reichert, M.: Decomposition-based verification of global compliance in process choreographies. In: *24th IEEE International Enterprise Distributed Object Computing Conference, EDOC 2020, Eindhoven, The Netherlands, October 5-8, 2020*. pp. 77–86 (2020)
- [24] Fellmann, M., Zasada, A.: State-of-the-art of business process compliance approaches. In: *ECIS* (2014)
- [25] Goedertier, S., Vanthienen, J.: Designing compliant business processes with obligations and permissions. In: *BPM'06 Workshops*. pp. 5–14 (2006)
- [26] González, L., Ruggia, R.: Towards a compliance-aware inter-organizational service integration platform. In: *On the Move to Meaningful Internet Systems: OTM 2014 Workshops*. pp. 8–17. Springer (2014)
- [27] González, L., Ruggia, R.: A comprehensive approach to compliance management in inter-organizational service integration platforms. In: *ICSOFT*. pp. 722–730 (2018)
- [28] Governatori, G., Idelberger, F., Milosevic, Z., Riveret, R., Sartor, G., Xu, X.: On legal contracts, imperative and declarative smart contracts, and blockchain systems. *Artif. Intell. Law* 26(4), 377–409 (2018)
- [29] Hashmi, M., Governatori, G., Lam, H., Wynn, M.T.: Are we done with business process compliance: state of the art and challenges ahead. *Knowl. Inf. Syst.* 57(1), 79–133 (2018)
- [30] Hashmi, M., Governatori, G., Wynn, M.T.: Normative requirements for regulatory compliance: An abstract formal framework. *Information Systems Frontiers* 18(3), 429–455 (2016)
- [31] Hildebrandt, T.T., Mukkamala, R.R., Slaats, T., Zanitti, F.: Contracts for cross-organizational workflows as timed dynamic condition response graphs. *J. Log. Algebraic Methods Program.* 82(5-7), 164–185 (2013)
- [32] Kasse, J.P., Xu, L., de Vrieze, P., Bai, Y.: The need for compliance verification in collaborative business processes. In: *PRO-VE*. pp. 217–229 (2018)
- [33] Kasse, J.P., Xu, L., de Vrieze, P., Bai, Y.: Verifying for compliance to data constraints in collaborative business processes. In: *PRO-VE*. pp. 259–270 (2019)
- [34] Knuplesch, D., Reichert, M., Mangler, J., Rinderle-Ma, S., Fdhila, W.: Towards compliance of cross-organizational processes and their changes. In: *BPM Workshops*. pp. 649–661 (2013)
- [35] Knuplesch, D., Reichert, M., Pryss, R., Fdhila, W., Rinderle-Ma, S.: Ensuring compliance of distributed and collaborative workflows. In: *CollaborateCom'13*. pp. 133–142 (2013)
- [36] Knuplesch, D.: *Enabling Multi-Perspective Business Process Compliance*, University of Ulm, Germany (2019)
- [37] Knuplesch, D., Fdhila, W., Reichert, M., Rinderle-Ma, S.: Detecting the effects of changes on the compliance of cross-organizational business processes. In: *ER'15*. pp. 94–107 (2015)
- [38] Knuplesch, D., Reichert, M.: A visual language for modeling multiple perspectives of business process compliance rules. *Software and Systems Modeling* 16(3), 715–736 (2017)
- [39] Knuplesch, D., Reichert, M., Fdhila, W., Rinderle-Ma, S.: On enabling compliance of cross-organizational business processes. In: *BPM*, pp. 146–154 (2013)
- [40] Knuplesch, D., Reichert, M., Kumar, A.: A framework for visually monitoring business process compliance. *Inf. Syst.* 64, 381–409 (2017)
- [41] Knuplesch, D., Reichert, M., Ly, L.T., Kumar, A., Rinderle-Ma, S.: On the formal semantics of the extended compliance rule graph. *Tech. Rep. 2013-05*, Ulm University (2013)
- [42] Knuplesch, D., Reichert, M., Ly, L.T., Kumar, A., Rinderle-Ma, S.: Visual modeling of business process compliance rules with the support of multiple perspectives. In: *ER'13*. pp. 106–120 (2013)
- [43] Kwantes, P.M., Van Gorp, P., Kleijn, J., Rensink, A.: Towards compliance verification between global and local process models. In: *Graph Transformation*. pp. 221–236. Springer (2015)
- [44] Ly, L.T., Maggi, F.M., Montali, M., Rinderle-Ma, S., van der Aalst, W.M.P.: Compliance monitoring in business processes: Functionalities, application, and tool-support. *Inf. Syst.* 54, 209–234 (2015)

Verifying Compliance in Process Choreographies

- [45] Maggi, F., Montali, M., Westergaard, M., van der Aalst, W.M.P.: Monitoring business constraints with linear temporal logic: an approach based on colored automata. In: BPM. pp. 132–147 (2011)
- [46] Matschinske, J.O., Späth, J., Nasirigerdeh, R., Torkzadehmahani, R., Hartebrodt, A., Orb'an, B., Fej'er, S., Zolotareva, O.I., Bakhtiari, M., Bihari, B., Bloice, M.D., Donner, N.C., Fdhila, W., Frisch, T., Hauschild, A.C., Heider, D., Holzinger, A., Hötendorfer, W., Hospes, J., Kacprowski, T., Kastelitz, M., List, M., Mayer, R., Moga, M., Müller, H., Pustozero, A., Röttger, R., Saranti, A., Schmidt, H.H.H.W., Tschohl, C., Wenke, N.K., Baumbach, J.: The featurecloud ai store for federated learning in biomedicine and beyond. ArXiv abs/2105.05734 (2021)
- [47] Meironke, A., Sey arth, T., Damarowsky, J.: Business process compliance and blockchain: How does the ethereum blockchain address challenges of business process compliance? In: Human Practice. Digital Ecologies. Our Future. 14. Internationale Tagung Wirtschaftsinformatik (WI 2019). pp. 1880–1891 (2019)
- [48] Mendling, J., et al.: Blockchains for business process management - challenges and opportunities. ACM Trans. Management Inf. Syst. 9(1), 4:1–4:16 (2018)
- [49] Meroni, G., Baresi, L., Montali, M., Plebani, P.: Multi-party business process compliance monitoring through iot-enabled artifacts. Inf. Syst. 73, 61–78 (2018)
- [50] Montali, M., Maggi, F.M., Chesani, F., Mello, P., van der Aalst, W.M.P.: Monitoring business constraints with the event calculus. Trans on Intelligent Sys and Tech 5(1), 17.1–17.30 (2014)
- [51] Montali, M., Pestic, M., Aalst, W.M.P.v.d., Chesani, F., Mello, P., Storari, S.: Declarative specification and verification of service choreographies. ACM Trans on the Web 4(1), 1–62 (2010)
- [52] Montali, M., Plebani, P.: Iot-based compliance checking of multi-party business processes modeled with commitments. In: 6th European Conference on Service-Oriented and Cloud Computing(ESOCC). pp. 179–195. Springer (2017)
- [53] Montali, M., Torroni, P., Chesani, F., Mello, P., Alberti, M., Lamma, E.: Abductive logic programming as an effective technology for the static verification of declarative business processes. Fundam. Informaticae 102(3-4), 325–361 (2010)
- [54] Muram, F., Tran, H., Zdun, U.: Automated mapping of uml activity diagrams to formal specifications for supporting containment checking. In: 11th Int Workshop on Formal Engineering Approaches to Software Components and Architectures. p. 93â 107 (2014)
- [55] Oyekola, O., Xu, L.: Verification and compliance in collaborative processes. In: Boosting Collaborative Networks 4.0 - 21st IFIP WG 5.5 Working Conference on Virtual Enterprises, PRO-VE 2020, Valencia, Spain, November 23-25, 2020, Proceedings. pp. 213–223 (2020)
- [56] Pestic, M., Schonenberg, H., van der Aalst, W.M.P.: DECLARE: full support for loosely-structured processes. In: EDOC. pp. 287–300 (2007)
- [57] Ramezani, E., Fahland, D., van der Aalst, W.M.P.: Where did i misbehave? Diagnostic information in compliance checking. In: BPM'12. pp. 262–278 (2012)
- [58] Schunselaar, D.M.M., Maggi, F.M., Sidorova, N.: Patterns for a log-based strengthening of declarative compliance models. In: Integrated Formal Methods. pp. 327–342 (2012)
- [59] Taghiabadi, E.R., Fahland, D., van Dongen, B.F., van der Aalst, W.M.P.: Diagnostic information for compliance checking of temporal compliance requirements. In: Advanced Information Systems Engineering. pp. 304–320 (2013)
- [60] Tosatto, S.C., Governatori, G., van Beest, N.: Checking regulatory compliance: Will we live to see it? In: Business Process Management. pp. 119–138 (2019)
- [61] Tosatto, S.C., Governatori, G., van Beest, N., Olivieri, F.: Efficient full compliance checking of concurrent components for business process models. FLAP 6(5), 963–998 (2019)
- [62] Voglhofer, T., Rinderle-Ma, S.: Collection and elicitation of business process compliance patterns with focus on data aspects. Bus. Inf. Syst. Eng. 62(4), 361–377 (2020), <https://doi.org/10.1007/s12599-019-00594-3>
- [63] Weske, M.: Business Process Management - Concepts, Languages, Architectures, Third Edition. Springer (2019)
- [64] Winter, K., Stertz, F., Rinderle-Ma, S.: Discovering instance and process spanning constraints from process execution logs. Inf. Syst. 89, 101484 (2020), <https://doi.org/10.1016/j.is.2019.101484>
- [65] Yao, J., Chen, S., Levy, D.: Accountability-based compliance control of collaborative business processes in cloud systems. In: Security, Privacy and Trust in Cloud Systems, pp. 345–374 (2014)
- [66] Zaman, R., Hassani, M.: On enabling GDPR compliance in business processes through data-driven solutions. SN Comput. Sci. 1(4), 210 (2020)

Highlights

- Decomposition enables decentralized checking of global compliance rules.
- Global compliance decomposition avoids trust in a central coordinator with knowledge on private tasks.
- Global compliance decomposition improves privacy of partner private processes
- Global Compliance decomposition reduces compliance-related partner communication.
- Proof-of-concept and practical applications

Conflict of interest

No conflict of interests.

Journal Pre-proof