

# Untrusted Business Process Monitoring and Execution Using Blockchain

Ingo Weber<sup>1,2(✉)</sup>, Xiwei Xu<sup>1,2</sup>, Régis Riveret<sup>3</sup>, Guido Governatori<sup>3</sup>,  
Alexander Ponomarev<sup>1</sup>, and Jan Mendling<sup>4</sup>

<sup>1</sup> Data61, CSIRO, Eveleigh, NSW, Australia

{Ingo.Weber, Xiwei.Xu, Alexander.Ponomarev}@data61.csiro.au

<sup>2</sup> School of Computer Science and Engineering, UNSW, Sydney, Australia

<sup>3</sup> Data61, CSIRO, Spring Hill, QLD, Australia

{Regis.Riveret, Guido.Governatori}@data61.csiro.au

<sup>4</sup> Wirtschaftsuniversität Wien, Vienna, Austria

jan.mendling@wu.ac.at

**Abstract.** The integration of business processes across organizations is typically beneficial for all involved parties. However, the lack of trust is often a roadblock. *Blockchain* is an emerging technology for decentralized and transactional data sharing across a network of untrusted participants. It can be used to find agreement about the shared state of collaborating parties without trusting a central authority or any particular participant. Some blockchain networks also provide a computational infrastructure to run autonomous programs called *smart contracts*. In this paper, we address the fundamental problem of trust in collaborative process execution using blockchain. We develop a technique to integrate blockchain into the choreography of processes in such a way that no central authority is needed, but trust maintained. Our solution comprises the combination of an intricate set of components, which allow monitoring or coordination of business processes. We implemented our solution and demonstrate its feasibility by applying it to three use case processes. Our evaluation includes the creation of more than 500 smart contracts and the execution over 8,000 blockchain transactions.

**Keywords:** Business process · Blockchain · Choreography · Orchestration

## 1 Introduction

The integration of business processes, e.g., along the supply chain, has been found to contribute both to better operational and business performance [4, 10]. A lack of trust, however, may hamper the innovativeness of further developing the collaborative process and its performance altogether [13]. Once service-level agreements are in place, it becomes a highly delicate question which partner should serve as a hub for controlling the collaborative process of several parties, or where a mediator process is hosted. While control asymmetries can be avoided

by a decentralized choreography instead of central orchestration, it does not solve the general problem of trust in controlling the collaborative business process.

The described lack-of-trust problem can be addressed with novel blockchain technology. Instead of agreeing on one trusted party, participants share transactional data across a large network of *untrusted* nodes (i.e., machines). This is achieved using a timestamped list of blocks which record, share, and aggregate data about transactions that have ever occurred within the blockchain network. Cryptographic proofs make this data storage immutable. As long as a majority share of the blockchain is not compromised, transactions can only be inserted; updating or deleting existing transactions is prohibitively expensive, making the blockchain tamper-proof. Blockchain also provides a global computational infrastructure, which can run programs: so-called *smart contracts* [12] execute across the blockchain network and automatically enforce the conditions defined in the transactions to enable, for example, conditional payment.

In this paper, we adopt blockchain technology to address the lack-of-trust problem in collaborative business processes. More specifically, we develop an approach to map a business process onto a peer-to-peer execution infrastructure that stores transactions in a blockchain, offering the following benefits. First, we provide a monitoring facility that integrates an automatic and immutable transaction history. Second, smart contracts can be used as a direct implementation of the mediator process control logic. Third, we obtain an audit trail for the complete collaborative business processes, for which payments, escrow, and conflict resolution can be enforced automatically. Our contribution is the first approach and implementation that leverages blockchain for collaborative process execution and monitoring. We evaluate our approach for feasibility by prototyping three use case processes on top of it. To this end, we ran of more than 500 process instances by creating as many smart contracts, and executed over 8,000 blockchain transactions that interact with the smart contracts.

The paper proceeds with a discussion of the research problem, related work, and blockchain technology in Sect. 2. Section 3 presents the details of our approach. Section 4 evaluates our approach using several real-world business scenarios, and Sect. 5 concludes. Technical details and evaluation use cases are described in a technical report (TR) [23]. Finally, a screencast video is available.<sup>1</sup>

## 2 Background

This section discusses the research problem we address, related work, and the background of blockchain technology as a solution.

### 2.1 Challenges of Collaborative Business Process Execution

We illustrate challenges of executing collaborative business processes by the help of a supply chain scenario reported in [3] that we simplify in Fig. 1. The process

<sup>1</sup> <https://youtu.be/1SNn9c5HHQs>.

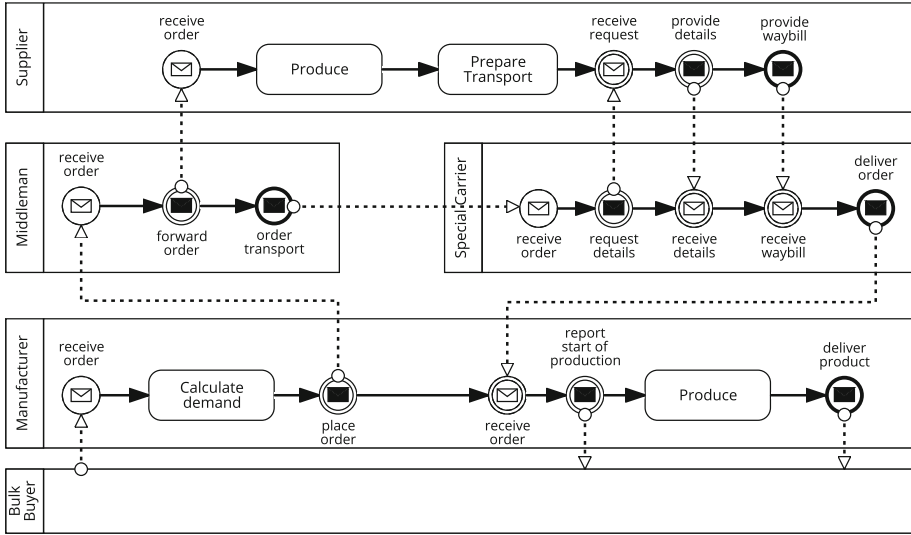


Fig. 1. Supply chain scenario from [3] (simplified)

starts with the Bulk Buyer placing an order with the Manufacturer. The latter calculates the demand and places an order for materials via a Middleman. This Middleman forwards the order to a Supplier and arranges transportation by a Special Carrier. Once the materials are produced, the Carrier picks them up at the Supplier site and delivers them to the Manufacturer. The Manufacturer produces the goods and delivers to the Bulk Buyer. The process is a *choreography* since there is no party that sees all messages. If all messages were sent and received by the Manufacturer, it would be an *orchestration* with the Manufacturer serving as a mediator [7].

**Conflict Example.** This simple scenario already involves five participants who would likely blame each other in case of delays and errors. Consider the case that the Manufacturer receives the materials three days later than agreed, with eight pallets being delivered instead of ten. The Supplier might argue that this is exactly in line with what was ordered by the Middleman while the Middleman would claim the fault to be on the side of the Supplier. The situation is delicate for the Carrier since the Manufacturer refuses to accept the delivery. The Carrier is now eligible for a compensation by the Supplier or the Middleman depending on who is responsible for the fault.

## 2.2 Prior Research on Collaborative Business Processes

Prior research on collaborative business processes has intensively investigated different notions of compatibility between the local processes of different partners and between local processes and a global process. Such compatibility can be

achieved by design, for instance using a P2P approach [20], transformations from a global choreography [7, 22], or interaction modeling [2].

Business processes involve different trust issues (see e.g. [21] for a summary) which can be addressed in different ways. For example, [1] relaxed the assumption that the broker hosting the process engine has to be trusted: using selective encryption, data access for both the broker and the service partners can be restricted. [8] designed a trust service for cross-company collaboration based on a hybrid architecture mixing a trusted centralized control with untrusted peer-to-peer components. [6] put forward an agent-based architecture that can remove the scalability bottleneck of a centralized orchestration engine, and provides more efficiencies by executing portions of processes close to the data they operate on. In virtual organizations, [15] proposed to select partners on the basis of disclosure policies and credentials (i.e. identity attributes issued by a “Credential Authority”).

Various important concepts such as conformance [19], reliability [16] and quality of services [24] have been investigated for centrally controlled business process execution. However, these works do not solve the trust issue: a collaborating party might have corrupted their historic files to their advantage. Technologies such as shared data stores provide solutions via consensus protocols to synchronize replicas [5] in a fully trusted environment. In this paper, we build our approach on blockchain technology for reasons explained next.

### 2.3 Blockchain Technology

Blockchain is the technology that supports Bitcoin [9]. The Bitcoin blockchain is a public ledger, which stores all transactions of the Bitcoin network. This concept has been generalized to distributed ledger systems that verify and store any transactions without coins or tokens [17]. A key feature of a blockchain-based system is that it does not rely on any central trusted authority, like traditional banking or payment systems. Instead, trust is achieved as an emergent property from the interactions between nodes within the network.

The blockchain data structure is an ordered list of blocks. *Blocks* are containers aggregating transactions. Every block is identifiable and linked to the previous block in the chain. *Transactions* are identifiable data packages that store parameters (such as monetary value in case of Bitcoin) and results of function calls in smart contracts. The integrity is ensured by cryptographic techniques. Once created, a transaction is signed with the signature of the transaction’s initiator, which indicates e.g. the authorization to spend the money, create a smart contract, or pass the data parameters associated with the transactions.

If the signed transaction is properly formed, valid and complete, it is sent to a few other nodes on the blockchain network, which will further validate it and send it to their peers until it reaches every node in the network. This *flooding approach* guarantees that a valid transaction will reach all the connected nodes in the network within a few seconds. The senders do not need to trust the nodes they use to broadcast the transactions, as long as they use more than one to ensure that it propagates. The recipient nodes do not need to trust the sender

either because the transaction is signed. When a transaction reaches a *mining node*, it is verified and included in a block. Blockchain networks rely on miners to aggregate transactions into blocks and append them to the blockchain. Once the transaction is confirmed by a sufficient number of blocks, it becomes a permanent part of the ledger and is accepted as valid by all nodes.

A *smart contract* is a user-defined program executed on the blockchain network [12]. It can be used to reach agreement and solve common problems. Smart contracts can be enforced as part of transactions, and are executed across the blockchain network by all connected nodes. The blockchain platform Ethereum views smart contract as a first-class element, and offers a built-in Turing-complete scripting language for writing smart contracts, called *Solidity*. Its execution environment, the *Ethereum Virtual Machine (EVM)*, comprises all full nodes on the network and executes bytecode compiled from Solidity scripts. Trust in the correct execution of smart contracts extends directly from regular transactions, since (i) they are deployed as data in a transaction, and hence immutable; (ii) all their inputs are through transactions; and (iii) their execution is deterministic. Deployed contracts should be tested. Whether the bytecode can be trusted is a separate matter, which we discuss for our approach in Sect. 4.5.

### 3 Blockchain-Based Collaborative Process Execution

In the following, we propose a blockchain-based system to address the lack-of-trust problem in collaborative business processes. A number of technical challenges arise during the adoption of blockchain for this purpose. For example, since transactions, computation, and data storage in blockchain platforms are not cost-free, not all aspects of collaborative processes should be dealt with inside smart contracts. However, smart contracts cannot call external APIs outside the blockchain environment or directly create blockchain transactions. This section presents our approach and how it addresses the challenges encountered.

#### 3.1 Overview of the Approach

An overview of our approach is shown in Fig. 2. We use blockchain to facilitate the collaborative processes in either of two ways:

(i) As a *choreography monitor*, it stores the process execution status across all involved participants by observing the message exchanges. In this setting, blockchain serves as an immutable data storage to share the process execution status and create an audit trail. Smart contracts check if interactions are conforming to the choreography model. In addition, a choreography monitor can be used to manage automated payment points and escrow.

(ii) As an *active mediator* among the participants, it coordinates the collaborative process execution. This includes all the above as well as using smart contracts to drive the process and implement data transformation or calculations.

These options are supported by the following main components:

- At design time, a **translator** derives from a process specification described in, e.g., Business Process Model and Notation (BPMN), a smart contract in a script language (such as Solidity). The generated smart contract is a factory for mediators or choreography monitors.
- For Option (i), a **Choreography monitor** or **C-Monitor** uses smart contracts *to monitor* the collaborative business processes. The C-Monitor is split into a factory and case-specific instance C-Monitors. The factory instantiates the case-specific monitors as needed, and contains the blueprint for instance C-Monitors. The C-Monitor instance tracks the interactions of a choreography instance and combines them into a consolidated view of the current state of the execution. Optionally, it can trigger automatic conditional payment from escrow, when certain points in the choreography are reached.
- For Option (ii), an active **mediator** uses a smart contract to *implement* the collaborative business processes. As with the C-Monitor, it is split between a factory and a set of instances and offers a consolidated view of the process state. In contrast to the C-Monitor, the mediator always plays an active role, receiving and sending messages according to the business logic defined in the process model. It also may transform data or execute other computations.
- **Interfaces** or **triggers** connect the process executing on blockchain and the external world. Because smart contracts cannot directly interact with the world outside the blockchain, a trigger plays the role of an organization’s agent. It holds confidential information and runs on a full blockchain node, keeping track of the execution context and status of running business processes. The trigger calls external APIs if needed, receives API calls from external components, and updates the process state in the blockchain based on external observations. It further keeps track of data payload in API calls and keeps the data in an external database when appropriate.

By the help of these components, we achieve that (i) participants can execute collaborative processes over a network of untrusted nodes, (ii) only conforming messages advance the state of the process, (iii) payments and escrow can be coded into the process, and (iv) an immutable ledger keeps a log of all transactions, successful or not. Next, we explain the above components in more detail. Additional details are available in a technical report [23].

### 3.2 Design Time: Translator

The translator is used at design time: it takes an existing business process specification as input and generates smart contracts. These implement the C-Monitor or mediator and can be deployed and executed on the blockchain.

In a collaborative process, this functionality must be split and distributed between the smart contract and the triggers. The translator creates the artifacts in such a way that the triggers and the smart contract can collaborate directly with each other over the blockchain network.

When the translator is called, it may not be known which participants will play which roles. Therefore, the translator outputs only a **factory contract**,

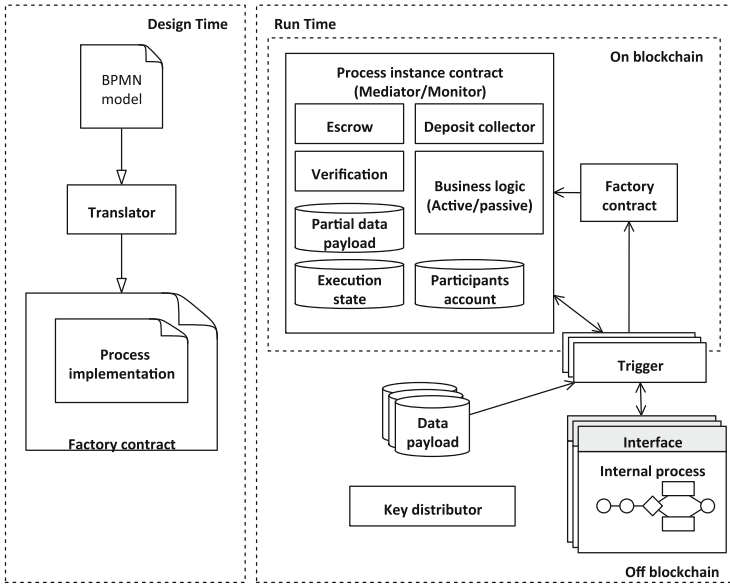


Fig. 2. Overview of our approach

which in turn contains all information needed for instantiating the process. The factory contract includes the methods for instantiation and two types of artifacts: (i) an interface specification per role (e.g., buyer, manufacturer, and shipper) in a collaborative process, to be distributed to the respective triggers, and (ii) a process instance contract, which is deployed to the blockchain when the process is instantiated. The process instance contract contains the implementation of the business logic and takes the form of a C-Monitor or mediator, depending on the content of the original process specification.

The overall translation algorithm has two phases. First, the translator parses the input process model and iterates through all its elements, where it generates two lists per element in the process model: one list of previous elements and one of next elements. Then, the translator translates each element with its respective links, generating Solidity code based on the translation rules for different types of elements as detailed in the TR [23]. Note that, in the current implementation, only some combinations of consecutive gateways can be connected to each other without tasks in between. The previous element list is used by the translator to determine which other elements need to be deactivated when the current element is executed; the next element list specifies which elements need to be activated after the current element is executed.

The selection methods for the two lists are shown in Algorithm 1. *NextElements* of an element includes all the tasks that directly follow the element, or the outgoing edge if the target of that edge is an AND-Join. If a next element is a Split or XOR-Join gateway, the tasks / edges that connect to it are added into *NextElements* through a recursive call. *PreviousElements* of

an element includes the element itself. If an XOR-Split gateway  $Split_i$  precedes the current element, the tasks that follow it are added to  $PreviousElements$ . In the case of an AND-Join gateway, all incoming edges are added to  $PreviousElements$ .

---

**Algorithm 1.** Calculating PreviousElements and NextElements.

---

```

1: function SELECTNEXTELEMENTS(Element, NextElements[])
2:   for all  $Edge_j \in outgoingEdges[Element]$  do
3:     if  $Edge_j.targetElement$  is Task then
4:        $NextElements \leftarrow Edge_j.targetElement$ 
5:     else if  $Edge_j.targetElement$  is AND-Join gateway then
6:        $NextElements \leftarrow Edge_j$ 
7:     else if  $Edge_j.targetElement$  is Split or XOR-Join gateway then
8:        $SelectNextElements(Edge_j.targetElement, NextElements[])$ 
9:     end if
10:  end for
11: end function
12:
13: function SELECTPREVIOUSELEMENTS(Element, PrevElements[])
14:   $PrevElements \leftarrow Element$ 
15:  if Element is Task then
16:    for all  $Edge_i \in incomingEdges[Element]$  do
17:      if  $Edge_i.sourceElement$  is XOR-Split gateway then
18:         $SelectNextElements(Edge_i.sourceElement, PrevElements[])$ 
19:      end if
20:    end for
21:  else if Element is AND-Join gateway then
22:    for all  $Edge_i \in incomingEdges[Element]$  do
23:       $PrevElements \leftarrow Edge_i$ 
24:    end for
25:  end if
26: end function

```

---

The **generator** is based on the workflow patterns [18]. Some patterns can be directly translated, some have to be supported off-chain, and other are unnecessary in our case. Our focus is not on supporting all elements of BPMN, but we start from the 5 basic control flow patterns [18], which are among the most frequently used elements in process models [25]. For brevity, we give an overview of the translation rules in Table 1. These make use of the two lists derived above, for activation / deactivation. After generating the smart contracts, the translator also calculates the cost range for executing the resulting smart contract. This serves as an indication of how much crypto-coins have to be spent in order to execute process instances over the blockchain.



**Table 1.** Translation rule summary. During traversal of the process model, when the translator encounters a pattern (left column), it inserts code according to the right column into the smart contract code. Scope concerns which variants the pattern applies to (M: mediator; CME: C-Monitor with escrow)

| BPMN element              | Scope  | Solidity code summary  |
|---------------------------|--------|--|
| All patterns              | All    | On execution, deactivates itself and activates the subsequent element.   |
| Parallel-Split            | All    | Executes on activation, activates <i>all</i> subsequent elements.  |
| Parallel-Join             | All    | Executes on activation of <i>all</i> incoming edges.   |
| XOR-Split                 | All    | Executes on activation, conditionally activates all subsequent elements. If one of them is executed, it deactivates all others.  |
| XOR-Join                  | All    | Executes on activation of <i>one</i> incoming edge.  |
| Choreography Task         | All    | Executes when the respective message is received (as blockchain transaction), and if the task is activated (message conforms with process). If conforming, the message is forwarded (as smart contract log entry); else, an alert is broadcasted.  |
| Task: Payment             | M, CME | Execution and conformance check as above. If conforming, payment into or from escrow is processed. Incoming payment is through a transaction, which has the desired effect already. Outgoing payment is sent to the account of the specified role. |
| Task: Data Transformation | M      | Execution and conformance check as above. Mediator-internal logic on data transformation, to be handled on-chain by the mediator or off-chain by a designated trigger.   |

### 3.3 Runtime Environment: Executing Processes as Smart Contracts

The translator generates all artifacts needed for runtime execution. We start by describing C-Monitors, which allow passive monitoring of choreographies and optionally escrow. Active mediators can be seen as an extension of C-Monitors, and the additions are explained subsequently. The third important concept for runtime, the triggers, and the interaction between triggers and smart contracts are covered afterwards. Finally, we describe how technical challenges like key distribution are handled.

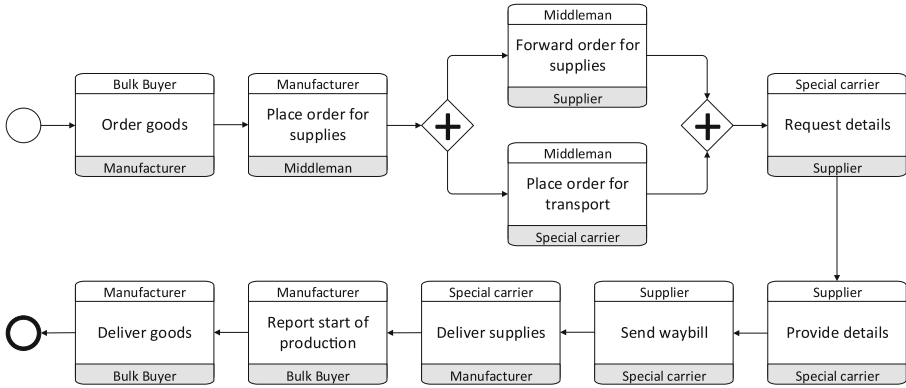


Fig. 3. BPMN choreography diagram of the process in Fig. 1

**Choreography Monitor.** The first way of facilitating collaborative processes is to use a smart contract as C-Monitor, with optional escrow and conditional payment at certain points of the processes. For a new process instance, an instance contract is generated from the factory contract. Initialization includes registering participants and their public keys to roles. The C-Monitor instance contract contains variables for storing the role assignment and for the process execution status. During execution, the involved participants do not interact with each other directly. Instead, they use the monitor to exchange their input/output data payload and, by doing so, advance the state of the collaborative process. Consider the choreography in Fig. 3, which is another representation of the collaborative process from Fig. 1. All tasks are communication tasks between roles. By exchanging the messages through the C-Monitor, it can check conformance with the choreography and track the status. In this way, conformance checking is done implicitly by the C-Monitor, and all transactions (successful or not) are logged in the blockchain. The handling of escrow is described below.

**Mediator.** Similar to the C-Monitor, the mediator is implemented as a smart contract, which is generated from the factory contract. It uses the same components as the C-Monitor. It also implements active components, among others to transform data and receive and send messages and payments.

**Triggers.** The Blockchain is a closed environment, where the deployed smart contracts cannot call external APIs. In our approach, a trigger (or blockchain interface) connects the participants’ internal processes with the blockchain. It monitors the process execution status, logically receives messages from smart contracts and calls external APIs, or receives API calls and logically sends messages to smart contracts accordingly.

Triggers are programs running on full nodes of the blockchain network. In the typical setup, every participant operates its own trigger deployed on a node it

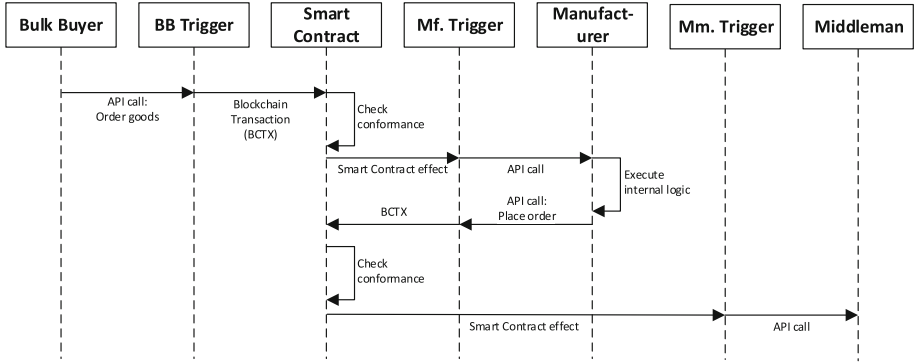


Fig. 4. Sequence diagram for the first two tasks in Fig. 3

controls, and the participant’s systems only communicate with its own trigger. We assume that this situation is given. Since the trigger is required to hold private keys for all participants on whose behalf it operates, a high degree of trust into the individual trigger is required.

When a new process instance is created, the participants register their roles and public keys. The public key corresponds to the account address of a participant. All keys and role assignments are passed to all triggers associated with the process instance, so everyone knows which role is played by whom and can verify messages accordingly. With the private key it holds, the trigger can encrypt or sign a message, allowing the contract and the other participants to verify its messages. In this fashion, it can also create payment transactions.

During the process execution, the trigger is receptive to API calls from its owner, as well as to logical messages from the process instance contract. The interaction between internal process implementations, triggers, and the process instance smart contract is shown in simplified form in Fig. 4. When a trigger’s API is called from its owner, the trigger translates the received message into a blockchain transaction, and sends the transaction to the instance contract. When the trigger receives a logical message from the instance contract, it updates its local state and calls an external API from the private process implementation.

Finally, the trigger takes care of sizable data payloads. For incoming API calls, it moves the data to secure storage, hashes it, and attaches a URI and the hash to the outgoing transaction. For incoming messages from the blockchain, it retrieves the data via its URI, checks if the hash matches, and sends it on to the internal process implementation.

**Encryption and Key Distribution.** All the information on the blockchain is publicly accessible to all nodes within the network. We store two types of information on blockchain, namely the process execution status and the data payload (or its URI/hash). To preserve the privacy of the involved participants, we have the option to encrypt the data payload before inserting it into the

blockchain. However, the process execution status is not encrypted because the C-Monitors and mediators need to process this information. Encrypting the data payload means that mediators cannot perform data transformation at all, but can resort to the source participant's trigger for this task.

We assume the involved participants exchange their public keys off-chain. *Encrypting data payload for all process participants* can be achieved as follows. One participant creates a secret key for the process instance, and distributes it during initial key exchange. When a participant adds data payload to the blockchain, it first symmetrically encrypts this information using the secret key. Thus, the publicly accessible information on blockchain is encrypted, i.e., useless to anyone who has no access to the secret key. The participants involved in the process instance have the secret key and can decrypt the information. *Encrypting data payload between two process participants*, in contrast, may be desired if two participants want to exchange information privately through the process instance. For this case, the sender can asymmetrically encrypt the information using the receiver's public key; only the receiver can decrypt it with its private key.

**Escrow.** The C-Monitor or mediator can also work as an escrow for conditional payment at designated points. Similar to an escrow agent, e.g., in real estate transactions, the smart contract receives money from one or more parties, and only releases the money to other parties once certain criteria are met. For the receivers this has the benefit that they can observe that the money is actually there before doing work; and the sender does not have to pay upfront, trusting it will eventually receive the goods or service in return.

In the running example process, the Manufacturer ( $Mf$ ) needs to pay the Middleman ( $Mm$ ), Supplier ( $S$ ) and Carrier ( $C$ ) when it receives the goods. But  $S$  is unwilling to send the goods without some guarantees that it will get paid. Therefore,  $Mf$  puts the money in escrow, namely an account held by the process instance contract, when ordering the goods. Later, both  $C$  and  $Mf$  confirm the delivery of the goods, which triggers automatic payment from the escrow account to  $Mm$ ,  $S$ , and  $C$ . The smart contract defines under what conditions the money can be transferred and how the money should be transferred. Thus, when a payment function is triggered, the smart contract automatically checks the defined conditions, and transfers the money according to the defined rules. It is, however, of high importance to specify rules that cover all possible scenarios and the respective outcomes: e.g., what shall happen with money in escrow if  $Mf$  and  $C$  disagree about the delivery of the goods or their condition?

**Gas Money.** The computation, data storage, and creation of smart contracts on the blockchain costs crypto-coins. That represents the cost for using the blockchain network, since it is used to pay the miners that execute the smart contracts. Each function call is thus accompanied by cost, but contract creation is relatively much more expensive than a regular function call. For fairness, the

participants in a collaborative process may want to decide on a different split of who pays how much, rather than the implicit split from the process.

## 4 Evaluation

### 4.1 Evaluation Method, Implementation, and Setup

The goal of our evaluation is to assess the feasibility of the approach. To this end, we implemented proof-of-concept prototypes for the translator and the trigger. The translator, written in Java, accepts BPMN 2.0 XML files, which we parse using the source code of the JBoss BPMN2 Modeller (jbpm-bpmn2 6.3.0). The translator's output are files that comply with the Solidity scripting language, version 0.2.0. Our smart contracts are running on go-ethereum 1.3.5, which is the official Golang implementation of the Ethereum protocol. The trigger is written as a Node.js web application, in JavaScript.

We picked three use case processes of different size, two from the literature and one from an industrial prototype. All three could be used directly as C-Monitor, and we extended one to cover the other options, i.e., C-Monitor with escrow and mediator. The key functionality of the blockchain is to accurately record the shared history of the choreography processes. Therefore, we derived the set of permissible execution traces for each process model, which we called the *set of conforming traces*. Furthermore, we randomly modified these traces to obtain a larger *set of not conforming traces* with the following manipulation operators: (i) add an event, (ii) remove an event, or (iii) switch the order of two events, such that the modified trace was different from all correct traces. Then we tested the ability of the smart contracts to discriminate between correct and incorrect traces. For escrow and the mediator data transformation, we ran a smaller number of experiments where we manually verified the effects.

Finally, during the above experiments we collected data that allows us to analyze important qualities. We focused particularly on cost and latency of using the blockchain in our setting, since these are the two non-functional properties that differ most from traditional approaches, such as trusted third parties. We ran experiments on a private blockchain and the public Ethereum blockchain, which allowed us to compare the effects of different options on these qualities.

### 4.2 Use Case Processes

For our evaluation, we used the following three processes.

1. Supply chain choreography: This process is discussed throughout this paper as a running example, see Fig. 3, and adapted from [3]. This process has ten tasks, two gateways and two conforming traces. From the 2 possible conforming traces, we generated 60 randomly manipulated traces. Out of these, 3 were conforming (switched order of parallel tasks) and 57 not.

2. Incident management choreography: This process stems from [11, p.18]. This process has nine tasks, six gateways and four conforming traces. We generated 120 not conforming traces. We implemented it with and without (i) a payment option and (ii) data manipulation in a mediator.
3. Insurance claim handling: This process is taken from the industrial prototype Regorous<sup>2</sup>. Choreographies tend to result in a simplified view of a collaborative process, as can be seen when comparing Figs. 1 and 3. To test the conformance checking feature with a more complex process, we added a third use case which was originally not a choreography. This process has 13 tasks, eight gateways and nine conforming traces. We generated 17 correct and 262 not conforming traces.

### 4.3 Identification of Not Conforming Traces

For this part of the evaluation, we investigate if our implementation accurately identifies the not conforming traces that have been generated for each of the models. The results are shown in Table 2. All log traces were correctly classified. This was our expectation: any other outcome would have pointed at severe issues with our approach or implementation.

**Table 2.** Process use case characteristics and conformance checking results

| Process                                    | Tasks | Gateways | Trace type          | Traces | Correctness |
|--|-------|----------|---------------------|--------|-------------|
| Supply chain process<br>of Fig. 3          | 10    | 2        | Conforming          | 5      | 100 %       |
|  |       |          | Not conforming      | 57     | 100 %       |
| Incident management                        | 9     | 6        | Conforming          | 4      | 100 %       |
|  |       |          | Not conforming      | 120    | 100 %       |
| Incident management<br>with payment        | 9     | 6        | Conforming          | 4      | 100 %       |
|  |       |          | Not conforming      | 19     | 100 %       |
| Incident mgmt. with<br>data transformation | 9     | 6        | Calculation         | 10     | 100 %       |
|  |       |          | String manipulation | 10     | 100 %       |
| Insurance claim                            | 13    | 8        | Conforming          | 17     | 100 %       |
|  |       |          | Not conforming      | 262    | 100 %       |

### 4.4 Analysis of Cost and Latency

In this part of the evaluation, we investigate the cost and latency of involving the blockchain in the process execution, since these are the non-functional properties that are most different from solutions currently used in practice.

<sup>2</sup> <http://www.regorous.com/>. A subset of the authors is involved in this project.

**Cost.** In our experiments on the private blockchain, we executed a total of 7923 transactions, at zero cost. On the public Ethereum blockchain, we ran 32 process instances with a total of 256 transactions. The deployment of the factory contract cost 0.032 Ether, and each run of the Incident Management process, with automatic payments and data transformations, cost on average 0.0347 Ether, or approx. US\$ 0.40 at the time of writing. The data (transactions and contract effects) of the experiment on the public blockchain is publicly viewable from the factory contract’s address, e.g. via Etherscan.<sup>3</sup>

**Latency.** We measure latency as the time taken from when the trigger receives an API call until it sends the response with conformance outcome, transaction hash, block number, etc. A test script iterates over the events in a trace and synchronously calls the trigger for each event. Therefore, the test script sends the next request very soon after receiving a response. This distorts the latency measurement to a degree, since the trigger adds the next transaction to the transaction pool just after the previous block has been mined, and it needs to wait there until mining for the block after the current one is started. Our measurements should thus be regarded as an upper bound, rather than the typical case. A more detailed explanation is given in the technical report [23].


An overview of the latency measurements is shown in Fig. 5<sup>4</sup>. The duration for a block to be mined comes from the complexity of the mining task, which is deliberately designed to be computationally hard. On the public Ethereum blockchain, the *target median time* between blocks at the time of writing is set to around 13 s, with the actual time measured at 14.4 s. On our private blockchain, we can control the complexity mechanism to increase mining time (shown as *Private fast* in Fig. 5) or leave the default implementation in place (*Private uncontrolled*). As can be seen, the variance is high. On the public Ethereum blockchain, the median latency was 23.0 s. In our *private fast* setting we achieved a median latency of 2.8 s, which should be sufficient for many practical deployments. For any application, this *tradeoff* needs to be considered: public blockchains offer much higher trustworthiness in return for higher cost and latency.

## 4.5 Discussion

**Conflict Resolution.** Following up on the *conflict example* from Sect. 2.1, we discuss how conflict resolution can be implemented in our approach. Recall that there was disagreement about the amount of supplies ordered. The blockchain inherently provides an immutable audit trail, thus it is trivial to review the original order and waybill messages – the culprit can be identified through such inspection. Say, the Supplier was at fault, but the Manufacturer paid crypto-coins

<sup>3</sup> <https://etherscan.io/address/0x09890f52cdd5d0743c7d13abe481e705a2706384>.

<sup>4</sup> Note that, instead of the typical error bars with *min* and *max* in box plots, we here show the 1st and the 99th percentile, to reduce the effect of the worst outliers. For *Private uncontrolled*, the *max* was 183 s – almost twice as much as the 99th percentile.



**Fig. 5.** Latency in seconds, using private blockchain with/without speed modification, and public Ethereum blockchain (box plot)

into escrow – how does it get its money back? The conditions for reimbursement from escrow need to be specified in the smart contract, but then they can be invoked at a later time. For instance, the participants may agree upfront that the Manufacturer gets reimbursed only if the Middleman agrees to that; then the Middleman sends a transaction to that effect, and the Manufacturer’s money is transferred back to its account.

**Trust.** Blockchain provides a trustworthy environment, without requiring trust in any single entity. In contrast, in the traditional model participants who do not trust each other need to agree on a third party which is trusted by all. Blockchain can replace this trusted third party. This is of particular interest in cases of cooperation. If multiple parties come together to achieve a joint business goal, but some of the organizations are in competition, it is important that the entity which executes the joint business process is neutral. Say, *Org1*, *Org2*, and *Org3* are in competition, but want to have a joint process to achieve some business goal. However, *Org1* would not accept *Org2* or *Org3* to control the process, and neither of those would accept *Org1*. Using our approach, the blockchain can be used, enabling trustless collaboration as it is not controlled by a single entity. Our translator allows the deployment of business processes on blockchain network without the need to manually implement the corresponding smart contract. *Trust in the deployed bytecode* for a process is established as follows: each participant has access to the process model, translates it to Solidity with our translator, and uses an agreed-upon Solidity compiler. This results in the same bytecode, and each participant can verify that the deployed bytecode has not been manipulated. Finally, the *trigger* allows for seamless integration into service-based message



exchanges. However, each trigger is a fully trusted party, and by default we assume each organization hosts their own trigger.

**Privacy.** Public blockchains do not guarantee any data privacy: anyone can join a public blockchain network without permission, and information on the blockchain is public. Thus, for scenarios like collaborative process execution, a permissioned blockchain may be more appropriate: joining it requires explicit permission. Even with permission management, the information on blockchain is still available to all the participants of the blockchain network. While we propose a method to encrypt the data payload of messages, the process status information is publicly available. As such, if *Org1*'s competitor, *Org4*, knows which account address belongs to which participant, it can infer with whom *Org1* is doing business and how frequently. This can be mitigated by creating a new account address for each process instance: the space of addresses is huge, and account creation trivial. However, this method prevents building a reputation, at least on the blockchain.

**Off-Chain Data Store.** For large data payloads, we propose to store only meta-data with a URI on-chain, and to keep the actual payload off-chain – accessible with the URI. Due to size limits for data storage on current blockchains [14] and associated costs, this solution can be highly advantageous. There are existing solutions that provide a data layer on top of blockchains, such as Factom [14]. Distributed data storage, like IPFS, DHT (Distributed Hash Table), or AWS S3, can also be used in combination with the blockchain to build decentralized applications.

**Threats to Validity.** There are several limitations to our study. To start, we made some assumptions when implementing our evaluation scenario, which bear threats to validity. First, we considered a supply chain scenario in which seconds of latency are typically not an issue. We expect that scenarios in other industries, such as automatic financial trading, would have stronger requirements in terms of latency, which could limit the applicability of our technique. Second, we worked with a network of limited size. A global network might have stronger requirements in terms of minimal block-to-block latency to ensure correct replication. These threats emphasize the need to conduct further application studies in different settings. Furthermore, there are open questions regarding technology acceptance, including management perception and legal issues of using blockchain technology.

## 5 Conclusion

Collaborative process execution is problematic if the participants involved have a lack of trust in each other. In this paper, we propose the use of blockchain and its smart contracts to circumvent the traditional need for a centralized trusted party

in a collaborative process execution. First, we devise a translator to translate process specifications into smart contracts that can be executed on a blockchain. Second, we utilize the computational infrastructure of blockchain to coordinate business processes. Third, to connect the smart contracts on blockchain with external world, we propose and implement the concept of triggers. A trigger converts API calls to blockchain transactions directed at a smart contract, and receives status updates from the contract that it converts to API calls. Triggers can thus act as a bridge between the blockchain and an organization's private process implementations. We ran a large number of experiments to demonstrate the feasibility of this approach, using a private as well as a public blockchain. While latency is low on a private, customized blockchain, the latency on the public blockchain may be considered too high for fast-paced scenarios. Additional benefits of our approach include the option to build escrow and automated payments into the process, and that the blockchain transactions from process executions form an immutable audit trail.

**Acknowledgments.** We thank Chao Li for integrating the trigger prototype with POD-Viz and recording the screencast video.

## References

1. Carminati, B., Ferrari, E., Tran, N.H.: Secure web service composition with untrusted broker. In: IEEE ICWS, pp. 137–144. IEEE (2014)
2. Decker, G., Weske, M.: Interaction-centric modeling of process choreographies. *Inf. Syst.* **36**(2), 292–312 (2011)
3. Fdhila, W., Rinderle-Ma, S., Knuplesch, D., Reichert, M.: Change and compliance in collaborative processes. In: IEEE SCC, pp. 162–169 (2015)
4. Flynn, B.B., Huo, B., Zhao, X.: The impact of supply chain integration on performance: a contingency and configuration approach. *J. Oper. Manag.* **28**(1), 58–71 (2010)
5. Kemme, B., Alonso, G.: Database replication: a tale of research across communities. *Proc. VLDB Endow.* **3**(1–2), 5–12 (2010)
6. Li, G., Muthusamy, V., Jacobsen, H.A.: A distributed service-oriented architecture for business process execution. *ACM TWEB* **4**(1), 2 (2010)
7. Mendling, J., Hafner, M.: From WS-CDL choreography to BPEL process orchestration. *J. Enterp. Inf. Manag.* **21**(5), 525–542 (2008)
8. Mont, M.C., Tomasi, L.: A distributed service, adaptive to trust assessment, based on peer-to-peer e-records replication and storage. In: IEEE FTDCS (2001)
9. Nakamoto, S.: Bitcoin: a peer-to-peer electronic cash system. <https://bitcoin.org/bitcoin.pdf>. Accessed 19 July 2015
10. Narayanan, S., Jayaraman, V., Luo, Y., Swaminathan, J.M.: The antecedents of process integration in business process outsourcing and its effect on firm performance. *J. Oper. Manag.* **29**(1), 3–16 (2011)
11. Object Management Group, June 2010. BPMN 2.0 by Example. [www.omg.org/spec/BPMN/20100601/10-06-02.pdf](http://www.omg.org/spec/BPMN/20100601/10-06-02.pdf). Version 1.0. Accessed 10 Mar 2016
12. Omohundro, S.: Cryptocurrencies, smart contracts, and artificial intelligence. *AI Matters* **1**(2), 19–21 (2014)

13. Panayides, P.M., Lun, Y.V.: The impact of trust on innovativeness and supply chain performance. *J. Prod. Econ.* **122**(1), 35–46 (2009)
14. Snow, P., Deery, B., Lu, J., Johnston, D., Kirby, P.: Business processes secured by immutable audit trails on the blockchain (2014)
15. Squicciarini, A., Paci, F., Bertino, E.: Trust establishment in the formation of virtual organizations. In: *ICDE Workshops, IEEE Computer Society* (2008)
16. Subramanian, S., Thiran, P., Narendra, N., Mostéfaoui, G., Maamar, Z.: On the enhancement of BPEL engines for self-healing composite web services. In: *Proceedings of SAINT Symposium*, pp. 33–39 (2008)
17. Tschorsch, F., Scheuermann, B.: Bitcoin and beyond: a technical survey on decentralized digital currencies. *IACR Cryptology ePrint Archive*, 2015, 464 (2015)
18. van der Aalst, W., ter Hofstede, A.H.M., Kiepuszewski, B., Barros, A.P.: Workflow patterns. *Distrib. Parallel Databases* **14**(1), 5–51 (2003)
19. van der Aalst, W.M.P., Dumas, M., Ouyang, C., Rozinat, A., Verbeek, E.: Conformance checking of service behavior. *ACM Trans. Internet Technol.* **8**(3) (2008)
20. van der Aalst, W.M.P., Weske, M.: The P2P approach to interorganizational workflows. In: *Dittrich, K.R., Geppert, A., Norrie, M. (eds.) CAiSE 2001. LNCS, vol. 2068*, pp. 140–159. Springer, Heidelberg (2001)
21. Viriyasitavat, W., Martin, A.: In the relation of workflow and trust characteristics, and requirements in service workflows. In: *Abd Manaf, A., Zeki, A., Zamani, M., Chuprat, S., El-Qawasmeh, E. (eds.) ICIEIS 2011, Part I. CCIS, vol. 251*, pp. 492–506. Springer, Heidelberg (2011)
22. Weber, I., Haller, J., Mülle, J.: Automated derivation of executable business processes from choreographies in virtual organizations. *Int. J. Bus. Process Integr. Manag. (IJBPM)* **3**(2), 85–95 (2008)
23. Weber, I., Xu, X., Riveret, R., Governatori, G., Ponomarev, A., Mendling, J.: Using blockchain to enable untrusted business process monitoring and execution. Technical report UNSW-CSE-TR-09, University of New South Wales (2016)
24. Zeng, L., Benatallah, B., Ngu, A., Dumas, M., Kalagnanam, J., Chang, H.: QOS-aware middleware for web services composition. *IEEE TSE* **30**(5), 311–327 (2004)
25. Muehlen, M., Recker, J.: How much language is enough? Theoretical and practical use of the business process modeling notation. In: *Bellahsène, Z., Léonard, M. (eds.) CAiSE 2008. LNCS, vol. 5074*, pp. 465–479. Springer, Heidelberg (2008)