



An Action-based Approach to the Formal Specification and Automatic Analysis of Business Processes under Authorization Constraints

Alessandro Armando^{a,b}, Enrico Giunchiglia^{*,a}, Marco Maratea^a, Serena Elisa Ponta^{a,c}

^a*DIST, University of Genova, Genova, Italy*

^b*Security & Trust Unit, Fondazione Bruno Kessler, Trento, Italy*

^c*SAP Research Sophia-Antipolis, Mougins, France*

Abstract

Business processes under authorization control are sets of coordinated activities subject to a security policy stating which agent can access which resource. Their behavior is difficult to predict due to the complex and unexpected interleaving of different execution flows within the process. Serious flaws may thus go undetected and manifest themselves only after deployment. For this reason, business processes are being considered a new, promising application domain for formal methods and model checking techniques in particular. In this paper we show that action-based languages provide a rich and natural framework for the formal specification of and automated reasoning about business processes under authorization constraints. We do this by discussing the application of the action language *C* to the specification of a business process from the banking domain that is representative of an important class of business processes of practical relevance. Furthermore we show that a number of reasoning tasks that arise in this context (namely checking whether the control flow together with the security policy meets the expected security properties, building a security policy for the given business process under given security requirements, and finding an allocation of tasks to agents that guarantees the completion of the business process) can be carried out automatically using the Causal Calculator CCalc. We also compare *C* with the prominent specification language used in model-checking.¹

Key words: Knowledge Representation and Reasoning, Action Languages, Business Processes

1. Introduction

Business processes under authorization control are sets of coordinated activities subject to a security policy stating which agent can access which resource. The order according to which the activities must be executed is given by the workflow, which is usually specified by means of graphical notations, e.g., Petri nets [2], Business Process Modelling Notation (BPMN) [3]. The specification of the policy is usually given in terms of a basic access control model (e.g., the Role-Based Access Control (RBAC) model [4]) possibly enriched with features providing the flexibility required by the application domain (e.g., delegation) and mechanisms that are necessary to meet mandatory regulations (e.g., separation of duty constraints). The behavior of business processes is difficult to predict due to the complex

¹This paper is an extended and thoroughly revised version of [1].

*Corresponding Author. Address: Viale F.Causa 15, 16145, Genova, Italy. Phone: +390103532980. Fax: +390103532948.

Email addresses: armando@dist.unige.it (Alessandro Armando), enrico@dist.unige.it (Enrico Giunchiglia), marco@dist.unige.it (Marco Maratea), serena.ponta@unige.it (Serena Elisa Ponta)

and unexpected interleaving of different execution flows within the process. Serious flaws (either in the control flow or in the security policy, or in both) may thus go undetected and manifest themselves only after deployment. Since these undesirable behaviors are very difficult to spot by simple inspection of the system, or by simulation, business processes are being considered a new, promising application domain for formal methods and model checking techniques in particular. In previous works, e.g., [5, 6], it has been shown that model checking can be profitably used for the automatic analysis of business processes, and a number of techniques have been designed to alleviate the state explosion problem, see, e.g., [7]. Yet, their applicability to business processes appears to be problematic as state-of-the-art model checkers—being geared to the analysis of hardware designs—require the system to be modeled as the composition of independent (yet interacting) sub-components. On the contrary, business processes and the associated security policy are best viewed as a collection of actions subject to a given workflow pattern and a set of independent access control rules.

Reasoning about actions and change is a long standing research area in Artificial Intelligence (AI), and several languages for knowledge representation [8] (e.g., logic programming [9, 10] and action languages [11, 12, 13]) and automated reasoning tools (e.g., SMOELS [14], DLV [15], DLV^K [16], CCALC [17]) have been put forward. In this paper we show that the action-based language *C* [18] provides a rich and natural specification framework for supporting formal declarative specifications of business processes under authorization constraints and that a number of reasoning tasks that arise in this context can be automated by using the Causal Calculator CCALC. The effectiveness of the proposed approach is illustrated by using *C* to specify the Loan Origination Process (LOP), a business process from the banking domain that is representative of an important class of business processes of practical relevance as it features many aspects that frequently occur in practice: non trivial interplay between the control flow and the security policy, sophisticated access-control policies, events and tasks with nondeterministic, conditional and indirect effects. To the best of our knowledge, no other approach to the specification and automatic analysis of business processes encompasses all the above aspects: [19, 20] consider complex security policy but do not take into account the workflow; [21, 22] analyze SoD constraints while considering both the workflow and the security policy but do not support implicit preconditions of tasks, events, roles hierarchy, and delegation; finally [5] considers the workflow and an RBAC security policy enhanced with delegation but does not consider, e.g., nondeterministic, indirect, and conditional effects of tasks, events, and policy exceptions.

More specifically we show that *C* supports

- the separate specification of the workflow and of the associated security policy;
- the formal and declarative specification of a wide range of security policies;
- the specification of a variety of business process features, e.g., events, tasks with nondeterministic, indirect, and conditional effects; and, most importantly,
- the seamless integration of all the above aspects.

Moreover *C* provides modelers with the ability to specify the system incrementally. This is an important feature that is seldom supported by the specification languages of state-of-the-art model-checkers. We substantiate this observation by comparing *C* with SMV, the specification language of the NuSMV model checker [23, 24].

We also show that a variety of automated reasoning tasks occurring in the domain of business processes can be recast as (a sequence of) satisfiability problems of *C* specifications which can be automatically tackled with the aid of CCALC:

1. to establish whether the control flow together with the security policy meets the expected security properties;
2. for a given number of agents, synthesize a security policy for the business process under given security requirements;
3. for a given number of agents and for all execution flows, find (if any) an assignment of activities to agents ensuring the process executability according to the given security policy.




Figure 1: Extended elementary net for the LOP

The rest of the paper is structured as follows. In the next section we provide a brief introduction to business processes under authorization constraints and describe our working example, which will be used throughout the paper. In Section 3 we overview the action language C . In Section 4 we show how a business process under authorization constraints can be modeled in C . In Section 5 we describe the analysis of the LOP carried out with $CCALC$. Then, in Section 6 we compare C and SMV. In Section 7 we discuss the related work and we conclude in Section 8 with some final remarks.

2. Business Processes under Authorization Constraints

Let us consider the Loan Origination Process (LOP) graphically presented in Figure 1. The workflow of the process is represented by means of an extended elementary net (see, e.g., [25]), a simple Petri net [2], extended with conditional arcs between places and transitions.

Let a *fact* be an atomic proposition. A *literal* is either a fact or its negation.

Formally, an *extended elementary net* is a 6-uple $\langle S, T, A, F, \gamma, M_0 \rangle$ where

- S and T are finite sets of *places* and *transitions*, respectively, such that $S \cap T = \emptyset$;

- $A \subseteq (S \times T) \cup (T \times S)$ is the *flow relation*, representing a set of directed *arcs* connecting places and transitions. Places from which arcs run to a transition are called the *input places* of the transition; places to which arcs run from a transition are called the *output places* of the transition;
- F is a set of facts, and the set of literals over F is denoted with \mathcal{L} . A set of literals \mathcal{L} is consistent if and only if \mathcal{L} does not contain a fact and its negation.
- $\gamma: (S \times T) \rightarrow \mathcal{L}$ is a function that associates arcs between places and transitions with literals over F expressing *applicability conditions*;
- $M_0 \subseteq S$ defines the *initial marking* for places.

In Figure 1

- $S = \{p_1, p_2, p_3, p_4, p_5, p_6, p_7, p_8, p_9, p_{10}, p_{11}, p_{12}, p_{12}, p_{13}, p_{14}, p_{15}, p_{16}, p_{17}, p_{18}\}$;
- $T = \{\text{exec}(\text{inputCustData}), \text{exec}(\text{prepareContract}), \text{exec}(\text{intRating}), \text{exec}(\text{extRating}), \text{invoke}(\text{creditBureau}), \text{exec}(\text{approve}), \text{exec}(\text{sign}), \text{exec}(\text{createAccount}), \text{beginFlow1}, \text{beginFlow2}, \text{beginFlow3}, \text{endFlow1}, \text{endFlow2}, \text{endFlow3}, \text{nop1}, \text{nop2}, \text{nop3}\}$;
- $A = \{(p_1, \text{exec}(\text{inputCustData})), (p_2, \text{beginFlow1}), (p_3, \text{exec}(\text{prepareContract})), (p_4, \text{beginFlow2}), (p_5, \text{endFlow1}), (p_6, \text{exec}(\text{intRating})), (p_7, \text{invoke}(\text{creditBureau})), (p_7, \text{nop1}), (p_8, \text{endFlow2}), (p_9, \text{exec}(\text{extRating})), (p_9, \text{nop2}), (p_{10}, \text{endFlow2}), (p_{11}, \text{endFlow1}), (p_{12}, \text{exec}(\text{approve})), (p_{13}, \text{nop3}), (p_{13}, \text{beginFlow3}), (p_{14}, \text{exec}(\text{sign})), (p_{15}, \text{exec}(\text{createAccount})), (p_{16}, \text{endFlow3}), (p_{17}, \text{endFlow3}), (\text{exec}(\text{inputCustData}), p_2), (\text{beginFlow1}, p_3), (\text{beginFlow1}, p_4), (\text{exec}(\text{prepareContract}), p_5), (\text{beginFlow2}, p_6), (\text{beginFlow2}, p_7), (\text{exec}(\text{intRating}), p_8), (\text{nop1}, p_{10}), (\text{invoke}(\text{creditBureau}), p_9), (\text{nop2}, p_7), (\text{exec}(\text{extRating}), p_{10}), (\text{endFlow2}, p_{11}), (\text{endFlow1}, p_{12}), (\text{exec}(\text{approve}), p_{13}), (\text{nop3}, p_{18}), (\text{beginFlow3}, p_{14}), (\text{beginFlow3}, p_{15}), (\text{exec}(\text{sign}), p_{16}), (\text{exec}(\text{createAccount}), p_{17}), (\text{endFlow3}, p_{18})\}$;
- $F = \{\text{highValue}, \text{interrupted}, \text{productApproved}\}$;
- $\gamma(p_7, \text{invoke}(\text{creditBureau})) = \text{highValue}, \gamma(p_7, \text{nop1}) = \neg \text{highValue}, \gamma(p_9, \text{nop2}) = \text{interrupted}, \gamma(p_9, \text{exec}(\text{extRating})) = \neg \text{interrupted}, \gamma(p_{13}, \text{beginFlow3}) = \text{productApproved}, \gamma(p_{13}, \text{nop3}) = \neg \text{productApproved}$;
- $M_0 = \{p_1\}$.

A state of the extended elementary net is a pair (M, L) , where M is a marking and $L \subseteq \mathcal{L}$ is a maximally consistent set of literals, (i.e., a truth-value assignment to the facts in F). A *marking* represents an execution state of the extended elementary net, initially set to M_0 . Given a marking M , a set of literal L and a place p , if $p \in M$ then we say that the place p *contains a token*, otherwise we say that p is *empty*. Starting from M_0 , a transition of the extended elementary net can *fire* if there is a token in every input place of the transition and the conditions associated with the arcs between

the input places and the transition hold in L . As a consequence, a new marking is reached where, for a transition that has fired, each output place contains a token, while the token of each input place is eliminated, and L is updated by the effects of the transition.

Figure 1 represents the LOP as an extended elementary net. Notice that in our model we have some artificially added transitions to begin/end a flow of business activities (i.e., transitions *beginFlow1*, *beginFlow2*, *beginFlow3*, *endFlow1*, *endFlow2*, and *endFlow3*) and to skip certain operations according to particular conditions (i.e., transitions *nop1*, *nop2*, and *nop3*). Such transitions are thus inserted for workflow modeling purposes, and are triggered as soon as their preconditions hold.

In our example, the process starts with the input of the customer's data (*inputCustData*). Afterwards a contract for the current customer is prepared (*prepareContract*) while the customer's rating evaluation takes place concurrently. By means of the rating evaluation the bank establishes if the customer is suitable to receive the loan. In our model, the execution follows different paths: if the amount of the requested loan is not high (\neg *highValue*), then an internal rating suffices (*intRating*); otherwise, an external rating (*extRating*) is executed concurrently by invoking a Credit Bureau, a third-party financial institution. As soon as it is ascertained that the external rating is needed, a request to obtain the credit information about the current customer is sent to the Credit Bureau (*invoke(creditBureau)*). By using this information, the external rating evaluation is performed by executing the task *extRating*. Notice that the invocation of the Credit Bureau and the execution of the task *extRating* must be performed by the same agent. In case there is a forbidden access to the information sent by the Credit Bureau to the bank, i.e., an agent who is not the director and different from the one who sent the request has accessed the information exchanged, then the execution of the task is prevented and the rating evaluation is interrupted. In case of interruption, the director of the bank must re-invoke the Credit Bureau and execute the task *extRating*. Thus, the loop in Figure 1 can be executed at most once.² The loan request must then be approved (*approve*) by the bank. As soon as the customer and the bank have reached an agreement, the contract is signed (*sign*) and an account for the customer is created concurrently (*createAccount*). Notice that the execution of a task may affect the state of the process. In particular

- task *inputCustData* may modify the state of the execution by issuing statements about the type of customer (i.e., if it is industrial, *isIndustrial*) and the amount of the loan (i.e., if it is high, *highValue*),
- task *intRating* may issue statements about the evaluation of the customer, i.e., if the internal rating of the customer is positive, and
- tasks *approve* may issue a statement asserting if the proposed product is suitable or not for the customer.

Moreover a task may perform further operations affecting the state of the process only if some conditions hold. In particular, in case the customer is industrial, during the evaluation of the internal rating the bank also establishes if it is important to deserve a particular care to the customer, i.e., if the customer is industrial the task *intRating* may issue a statement asserting if the customer has a high profile (*highProfileIndCust*). Finally, different statements may determine new process features, e.g., the *lowRisk* condition is used to denote a situation in which the internal rating is positive and the amount of the loan is not high.

An agent can execute a task only if she has the required permissions. As it is common in the business domain, the security policy of the LOP relies on an access control model based on RBAC enhanced with delegations and separation of duty constraints. According to the RBAC model [4], in order to perform a task an agent must be assigned a role enabled to execute the task and the agent must be also active in that role. The roles used in our case study are *director*, *manager*, *supervisor*, *postprocessor*, and *preprocessor*. Roles can be organized hierarchically. In our case study, a *director* is more senior than a *manager* and a *supervisor* is more senior than a *postprocessor*. Senior roles inherit the permissions to perform tasks assigned to more junior roles. Thus, an agent can execute a task if her role

- is directly assigned the required permissions; or
- is more senior than a role owning such permissions.

²Note that this does not limit the specification capabilities of our approach: the fact that such loop can be executed at most once is, by chance, the result of the separate specification of the workflow (which allows for multiple executions of the loop) with a *particular* security policy adopted.

Table 1: Permission assignment for the LOP

Task	Role
<i>inputCustData</i>	<i>preprocessor</i>
<i>prepareContract</i>	<i>postprocessor</i>
<i>intRating</i>	if (<i>isIndustrial</i>) then <i>postprocessor</i> else <i>preprocessor</i>
<i>extRating</i>	if (<i>interrupted</i>) then <i>director</i> else <i>supervisor</i>
<i>approve</i>	if (<i>lowRisk</i>) then <i>manager</i> else <i>director</i>
<i>sign</i>	if (<i>highProfileIndCust</i>) then <i>director</i> else <i>manager</i>
<i>createAccount</i>	<i>postprocessor</i>

In our case study we consider permission assignments subject to the following requirements

- task *inputCustData* is assigned to role *preprocessor*;
- tasks *prepareContract* and *intRating* cannot be assigned to roles *director* or *manager*;
- task *createAccount* cannot be assigned to roles *director*, *manager*, or *preprocessor*;
- tasks *approve* and *sign* cannot be assigned to roles *preprocessor*, *postprocessor*, or *supervisor*;
- if a customer is industrial, then role *preprocessor* cannot be enabled to perform the task *intRating*;
- if the process has been interrupted, the task *extRating* has to be performed by a *director*;
- if the process has not been interrupted, the task *extRating* cannot be assigned to roles *director*, or *manager*;
- if the risk of the loan is not low, then role *manager* cannot be enabled to perform the task *approve*;
- if the industrial customer has a high profile, then role *manager* cannot be enabled to perform the task *sign*.

Table 1 shows a possible permission assignment for the LOP satisfying the requirements presented above. Notice that the invocation of the Credit Bureau does not appear in Table 1 as this activity has to be executed by the same agent of the task *extRating* and, as a consequence, uses its permission assignment.

A *user assignment* is a relation that associates agents and roles. We consider a static assignment of agents to roles subject to the following requirements

- there must be only one director,
- the director must not be assigned to any other role,
- an agent must not be assigned to roles hierarchically related, e.g., an agent cannot be assigned to the roles *supervisor* and *postprocessor*, and
- an agent can be assigned to two different roles at most.

Notice that the requirements we have specified above, for both permission and user assignments, correspond to a *class* of security policies, given, e.g., the use of conditional statements in the permission assignment.

A possible user assignment for the LOP is given by the following assignments of agents to roles: *davide*, the director, *maria* and *marco*, managers, *pierPaolo*, who can act both as preprocessing clerk and as postprocessing clerk, *pierSilvio*, who can act both as preprocessing clerk and as supervisor, *pietro*, postprocessing clerk, and *stefano*, supervisor.

Our RBAC access control model is enhanced with delegation that represents a typical flexibility requirement. Following the idea of conditional delegation presented in [26], we consider *delegation rules* of the form

$$\langle PreConds, ARole, DRole, Task \rangle$$

Table 2: Delegation rules of the LOP

Name	Type	Delegation Rule
D1	transfer	$\langle \neg isIndustrial, supervisor, postprocessor, extRating \rangle$
D2	grant	$\langle intRatingPositive, manager, supervisor, approve \rangle$
D3	grant	$\langle \neg highValue, director, supervisor, sign \rangle$

Table 3: Critical tasks of each object-based SoD of the LOP

Name	Object	Critical Tasks
C1	customer's data	$inputCustData, prepareContract, intRating, extRating$
C2	rating report	$intRating, extRating$
C3	contract	$prepareContract, approve, sign$

where $ARole$ and $DRole$ are roles, $Task$ is a task, and $PreConds$ is a set of conditions that must hold for the delegation to be applicable. A delegation rule states that if $PreConds$ holds and $ARole$ is authorized to perform $Task$ according to the permission assignments, then $ARole$ can delegate $DRole$ to execute $Task$. We also distinguish between *grant* and *transfer* delegation [27]. Considering a grant delegation rule, both agents involved are then allowed to perform the task being delegated; considering a transfer delegation rule the ability to perform the task is transferred and the agent who executed the delegation cannot perform the delegated task anymore. Notice that our delegation rules express task delegation rather than role delegation. In fact, the delegated agent does not acquire a new role but she only obtains the permission to perform $Task$ by means of $ARole$. Also notice that in [26] three different kinds of conditions are considered, i.e., temporal, value, and workload delegation conditions, however in our work $PreConds$ can only be related to the value of the attributes of the process, i.e., value delegation conditions. In our case study we consider the delegation rules in Table 2.

We also consider the ability of the director to disable an agent from performing a task overriding the security policy in use. As a result, an agent can execute a task if she is granted the permission by means of the RBAC model or by means of delegation unless the director explicitly disables her from performing the task.

Finally, the RBAC model of our case study is enhanced with a mechanism that is necessary to satisfy separation of duty (SoD) constraints. SoD constraints are used for internal control and amount to requiring that some critical tasks are executed by different agents (see [5] for a survey on SoD). In this paper we focus on a relaxed form of *object-based SoD* according to which an agent can access the same object through different roles as long as she does not perform all the tasks accessing that object. For each object involved in the LOP, we define the corresponding critical tasks consisting of all and only the tasks accessing the object. We then assume that an agent cannot execute all the critical tasks associated to each object. Such associations are presented in Table 3.

3. The Action Language \mathcal{C}

Action languages are high level formalisms for expressing *actions* and how they affect the world described with a set of atomic formulas called *fluents*. Thus, the signature σ of the language is partitioned into the *fluent symbols* σ^{fl} and the *action symbols* σ^{act} . Intuitively, *actions* are a subset of the interpretations of σ^{act} while *states* are a subset of the interpretations of σ^{fl} . A formula in σ is a propositional combination of atoms. Each action language differs from the others for the constructs used to characterize how actions affect states.

\mathcal{C} [18] is an expressive propositional action language allowing for two kinds of propositions: *static laws* of the form

$$\text{caused } F \text{ if } G \quad (1)$$

and *dynamic laws* of the form

$$\text{caused } F \text{ if } G \text{ after } H, \quad (2)$$

where F and G are *fluent formulas* (i.e., formulas in σ^{fl}) and H is an *action formula* (i.e., a formula in σ). In a proposition of either kind, the formula F will be called the *head* of the law.

Table 4: Abbreviations for causal laws.

Abbreviation	Expanded Form	Informal Meaning
nonexecutable H' if F	caused \perp after $H' \wedge F$	$\neg F$ is a precondition of H'
H' causes F if G	caused F if \top after $G \wedge H'$	F is true after H' is executed in a state in which G is true
H' may cause F if G	caused F if F after $H' \wedge G$	F is true by default after H' is executed in a state in which G is true
default F	caused F if F	F is true by default
caused F if G unless Q	caused F if $G \wedge \neg Q$, default $\neg Q$	F is true after G in a state in which Q is false
constraint F	caused \perp if $\neg F$	F must be true
inertial F	caused F if F after F , caused $\neg F$ if $\neg F$ after $\neg F$	F is inertial
exogenous F	default F , default $\neg F$	F is exogenous

An *action description* is a set of propositions. Consider an action description D . A *state* is an interpretation of σ^{fl} that satisfies $G \supset F$ for every static law (1) in D . A *transition* is any triple $\langle s, a, s' \rangle$ where s, s' are states and a is an action; s is the *initial* state of the transition, and s' is its *resulting* state. A formula F is *caused* in a transition $\langle s, a, s' \rangle$ if it is

- the head of a static law (1) from D such that s' satisfies G , or
- the head of a dynamic law (2) from D such that s' satisfies G and $s \cup a$ satisfies H .

A transition $\langle s, a, s' \rangle$ is *causally explained* according to D if its resulting state s' is the only interpretation of σ^{fl} that satisfies all formulas caused in this transition.

The *transition diagram* represented by an action description D is the directed graph which has the states of D as nodes and includes an edge from s to s' labeled a for every transition $\langle s, a, s' \rangle$ that is causally explained according to D . Intuitively, if $\langle s, a, s' \rangle$ is a transition of the diagram, the concurrent execution of the atomic actions satisfied by a causes a transition from the state s to the state s' . Despite the fact that C consists of only two kinds of propositions, several other propositions can be defined as abbreviations of either (1) or (2), modeling, e.g., actions' preconditions, actions' nondeterministic effects, fluents with default values, and inertial fluents. The abbreviations used in this paper are given in Table 4, where F and G are defined as before, and H' is a formula in σ^{act} .

Example (from [18]). Let $\sigma^{fl} = \{P, Q\}$, $\sigma^{act} = \{A\}$, and let D consist of the propositions

$$\begin{aligned} &\text{inertial } P, Q, \\ &\text{caused } P \text{ after } Q \wedge A. \end{aligned} \quad (3)$$

The second line of (3) tells us that P is made true by the execution of A if the precondition Q is satisfied (as, for instance, in the familiar shooting example which corresponds to *Dead* as P , *Loaded* as Q , and *Shoot* as A). Preconditions are specific conditions holding in a state they are executed. According to Table 4, (3) is a shorthand for

$$\begin{aligned} &\text{caused } P \text{ if } P \text{ after } P, \\ &\text{caused } \neg P \text{ if } \neg P \text{ after } \neg P, \\ &\text{caused } Q \text{ if } Q \text{ after } Q, \\ &\text{caused } \neg Q \text{ if } \neg Q \text{ after } \neg Q, \\ &\text{caused } P \text{ if True after } Q \wedge A. \end{aligned} \quad (4)$$




Figure 2: Transition diagram for action description (3)

In this action description, there are 4 states ($PQ, P\bar{Q}, \bar{P}Q, \bar{P}\bar{Q}$)³ and 2 actions (A and \bar{A}). Consequently, there are $4 \times 2 \times 4 = 32$ transitions. Out of these, 8 transitions are causally explained: $\langle \bar{P}Q, A, PQ \rangle$ and the transitions of the form $\langle s, a, s \rangle$ where $s \neq \bar{P}Q$ or $a \neq A$.

Figure 2 shows the corresponding transition diagram. We can see from it that each of the actions A, \bar{A} can be executed in any state in exactly one way. To check that the transition $\langle \bar{P}Q, A, PQ \rangle$ is causally explained, note that the formulas caused in this transition are the heads Q, P of the 3rd and 5th propositions in (4), and that the resulting state PQ of the transition is the only interpretation that satisfies both heads. We can also see that $\langle \bar{P}Q, \bar{A}, PQ \rangle$ is not causally explained: the only formula caused in this transition is the head Q of the 3rd proposition in (4), and Q is satisfied by more than one interpretation.

4. Formal Modeling of Business Processes with \mathcal{C}

We now present a list of desired features for modeling business processes under authorization constraints, and we show the natural correspondence between the features, exemplified by using the LOP as a working example, and the \mathcal{C} constructs we use. This section clearly shows that \mathcal{C} is a well-suited modeling language for our domain.

The set of fluents involved consists of the fluents listed in Table 5. All the fluents in Table 5 but $pa(r, t)$, $lowRisk$, $granted(a, r, t)$, $activated(a, r)$, $accessed(a)$, and $disabled(a, t)$ are inertial. The set of actions involved is listed in Table 6. All the fluents in Table 5 and actions in Table 6 have to be suitably instantiated for all tasks and roles in Table 1 and the set of available agents.

4.1. Internal and External Events

Events are effects of activities which are not explicitly represented in the model. We distinguish between internal and external events

- internal events are effects of procedures implemented within the process but not explicitly represented in the model;
- external events are effects of activities which are not part of the process but whose effects have an impact on the process itself.

Examples in the LOP are, respectively,

- the activation of a role, which is modeled as an internal event making sure that an agent is active in a role only if she is assigned to that role; and
- the access to some information by an agent not allowed to, which is modeled as an external event making sure that once an agent has accessed the information, a trace of this event remains in the process.

³We represent a propositional interpretation by listing the literals that are satisfied by it. \bar{L} is the literal complementary to L .

Table 5: Fluents and their informal meaning

Fluent	Meaning
<i>activated(a, r)</i>	agent <i>a</i> is playing role <i>r</i>
<i>accessed(a)</i>	agent <i>a</i> has accessed the information sent by the Credit Bureau
<i>delegated(a, r, t)</i>	agent <i>a</i> is delegated by an agent in role <i>r</i> to perform task <i>t</i>
<i>disabled(a, t)</i>	agent <i>a</i> cannot perform task <i>t</i>
<i>pa(r, t)</i>	role <i>r</i> has the permission to perform task <i>t</i>
<i>ua(a, r)</i>	agent <i>a</i> is assigned to role <i>r</i>
<i>granted(a, r, t)</i>	agent <i>a</i> obtained by means of role <i>r</i> the permission to perform task <i>t</i>
<i>executed(a, t)</i>	agent <i>a</i> has executed task <i>t</i>
<i>invoked(a, r, e)</i>	agent <i>a</i> has invoked entity <i>e</i> obtaining the authorization from role <i>r</i>
<i>senior(r₁, r₂)</i>	role <i>r₁</i> is more senior than or is as senior as <i>r₂</i>
<i>lowRisk</i>	the risk associated with the loan is low
<i>highValue</i>	the loan amount is high wrt the financial status of the customer
<i>isIndustrial</i>	the customer is industrial
<i>highProfileIndCust</i>	the industrial customer has a high profile
<i>intRatingPositive</i>	the customer's internal rating is positive
<i>productApproved</i>	both the customer and the bank agree on the contract
<i>interrupted</i>	the execution of the process is interrupted
<i>p₁, ..., p₁₈</i>	the places of the extended elementary net (cf. Fig. 1)

Table 6: Actions and their informal meaning

Actions	Meaning
<i>exec(a, r, t)</i>	agent <i>a</i> executes task <i>t</i> by means of role <i>r</i>
<i>invoke(a, r, e)</i>	agent <i>a</i> invokes entity <i>e</i> obtaining the authorization from role <i>r</i>
<i>disable(a1, a2, t)</i>	agent <i>a1</i> disables agent <i>a2</i> from performing task <i>t</i>
<i>d1(a1, a2)</i>	agent <i>a1</i> delegates agent <i>a2</i> by means of delegation rule <i>d1</i>
<i>d2(a1, a2)</i>	agent <i>a1</i> delegates agent <i>a2</i> by means of delegation rule <i>d2</i>
<i>d3(a1, a2)</i>	agent <i>a1</i> delegates agent <i>a2</i> by means of delegation rule <i>d3</i>
<i>beginFlow1</i>	dummy activity to begin the first flow of business activities in Fig. 1
<i>endFlow1</i>	dummy activity to end the first flow of business activities in Fig. 1
<i>beginFlow2</i>	dummy activity to begin the second flow of business activities in Fig. 1
<i>endFlow2</i>	dummy activity to end the second flow of business activities in Fig. 1
<i>nop1, nop2, nop3</i>	dummy activities to skip operations

Internal and external events are modeled in C as exogenous fluents. Exogenous fluents are fluents that can change their value during the transition from one state to another (unless there is some other law constraining their values). In the case of the LOP, two exogenous fluents are used:

exogenous *activated(a, r), accessed(a)*.

The former is subject to the law

caused \neg *activated(a, r)* **if** \neg *ua(a, r)*

to state that an agent cannot be active in a role she is not assigned to. The latter is subject to the law

caused *accessed(a)* **after** *accessed(a)*

to state that the access to the information exchanged in the process by an agent is not reversible.

4.2. Execution of Tasks and Invocation of Entities

The execution of a task can take place only if its preconditions are met. Preconditions of each task contain information about

- the input places of the corresponding transition in the extended elementary net;
- the condition associated with the arc from the input places, if any; and
- the authorized agent that will perform the task by means of the fluent $granted(a, r, t)$.

In C , preconditions can be expressed by the law

$$\mathbf{nonexecutable} \ H \ \mathbf{if} \ \neg F \quad (5)$$

where H is an action formula and F is a formula expressing the preconditions of H . Consider a task t , and let I be the set of input places of the transition, and I' the set of conditions on the arcs from the input places (if any). Task preconditions have the following pattern:

$$\mathbf{nonexecutable} \ exec(a, r, t) \ \mathbf{if} \ \neg(\bigwedge_{p \in I} p \wedge \bigwedge_{f \in I'} f \wedge granted(a, r, t)).$$

An example about $intRating$ is the following

$$\mathbf{nonexecutable} \ exec(a, r, intRating) \ \mathbf{if} \ \neg(p_6 \wedge granted(a, r, intRating)).$$

Note that for the task $extRating$ we do not need to add the fluent $granted(a, r, extRating)$ in its preconditions as it is implicit in $invoked(a, r, creditBureau)$ because, as presented in Section 2, the task $extRating$ must be executed by the same agent who performed $invoke(a, r, creditBureau)$, i.e.,

$$\mathbf{nonexecutable} \ exec(a, r, extRating) \ \mathbf{if} \ \neg(p_9 \wedge \neg interrupted \wedge invoked(a, r, creditBureau))$$

The execution of tasks has deterministic effects that affect the process each time the task is performed. Consider a task t , and let I (resp. O) be the set of input (resp. output) places of the related transition. The specifications of the deterministic effects of a task in C have the following pattern:

$$exec(a, r, t) \ \mathbf{causes} \ \bigwedge_{p \in I} \neg p \wedge \bigwedge_{p' \in O} p' \wedge executed(a, t)$$

As an example, the deterministic effects of the execution of task $intRating$ are expressed in C with

$$exec(a, r, intRating) \ \mathbf{causes} \ p_8 \wedge \neg p_6 \wedge executed(a, intRating)$$

Some effects of a task can be nondeterministic, i.e., they can, or cannot, be true in the resulting state. As an example, the execution of task $intRating$ has the nondeterministic effect of switching to true the internal rating of a customer initially set to false. The corresponding C construct is

$$exec(a, r, intRating) \ \mathbf{may\ cause} \ intRatingPositive.$$

The execution of a task can also have conditional effects, i.e., effects which can take place or not depending on the state in which the task is executed. As an example, if the customer is industrial the task $intRating$ does not only evaluate if the internal rating is positive but also establishes if the customer has a high profile, i.e., if it is convenient for the bank to reserve particular care to this customer. The result of this evaluation is nondeterministic and conditional as it only takes place if the customer is industrial. This can be expressed in C with

$$exec(a, r, intRating) \ \mathbf{may\ cause} \ highProfileIndCust \ \mathbf{if} \ isIndustrial.$$

The execution of a task can also have some indirect effects in addition to those explicitly stated, resulting from the interaction among action effects and/or static laws. As an example, the execution of task *inputCustData*, that has the nondeterministic effect of stating if the customer is industrial, has an indirect effect on the permission assignment of the RBAC model. In fact, according to the permission assignments of Table 1, the permission to execute task *intRating* is given to role *postprocessor* if the customer is industrial. This can be expressed in C with the law

caused $pa(postprocessor, intRating)$ **if** $isIndustrial$

where $pa(postprocessor, intRating)$ can be an indirect effect of the execution of *inputCustData* through its nondeterministic effect $isIndustrial$. Notice that $pa(postprocessor, intRating)$ is a *statically determined fluent*, i.e., a fluent whose value is determined by means of static laws only. Also notice that we set the value of $pa(postprocessor, intRating)$ to be false by default by means of

default $\neg pa(postprocessor, intRating)$.

Laws related to an invocation of an entity have the same patterns as for the execution of a task. As an example the invocation of the Credit Bureau can be expressed in C by

nonexecutable $invoke(a, r, creditBureau)$ **if** $\neg(p_7 \wedge highValue \wedge granted(a, r, extRating))$ (6)

where $invoke(a, r, creditBureau)$ is the action symbol, p_7 is the input place as shown in Figure 1, $highValue$ represents the condition associated with the arc between place p_7 and the current transition, and $granted(a, r, extRating)$ expresses the fact that the invocation must be performed by an agent granted the right to execute the task *extRating* as presented in Section 2.

Deterministic effects of the invocation of an entity can be expressed in C in the same way as for the invocation of tasks. As an example, the deterministic effects of the invocation of the Credit Bureau are expressed by

$invoke(a, r, creditBureau)$ **causes** $p_9 \wedge \neg p_7 \wedge invoked(a, creditBureau)$.

According to (6), in our working example the invocation of the Credit Bureau cannot be executed if the preconditions do not hold; however, in case the preconditions hold, the invocation is not forced. The fact that the in our working example the invocation of the Credit Bureau is triggered as soon as its preconditions hold can be expressed in C with

nonexecutable $\neg invoke(a, r, creditBureau)$ **if** $p_7 \wedge highValue \wedge granted(a, r, extRating)$.

4.3. Delegation of Tasks

As in Section 4.2, the preconditions for delegation express the requirements for the delegation to take place and are expressed in C by a formula of the form (5). The set of preconditions for the delegation of tasks through the delegation rules presented in Section 2 has a common pattern. In fact, a delegation of a task always requires that

- an agent is assigned to the role $ARole$ of the rule,
- there exists a permission assignment between $Arole$ and $Task$,
- an agent is assigned to the role $DRole$ of the rule,
- $PreConds$ are satisfied,

regardless the type of delegation. The preconditions of the rules in Table 2 in C have the following pattern

nonexecutable $d(a_1, a_2)$ **if** $\neg(ua(a_1, Arole) \wedge ua(a_2, Drole) \wedge pa(Arole, Task) \wedge Preconds)$

for all $a_1 \neq a_2$.

For example the preconditions of rule D1 are expressed with

nonexecutable $d1(a_1, a_2)$ **if** $\neg(ua(a_1, supervisor) \wedge ua(a_2, postprocessor) \wedge pa(supervisor, extRating) \wedge \neg isIndustrial)$

Delegation can also have implicit preconditions. In fact, the effects of a delegation may interact with different parts of the business process; as a result, a delegation can be performed not only if its preconditions are fulfilled but also if its execution does not violate existing dependencies. In C implicit preconditions can be expressed by means of static laws. As an example, delegation has the effect of granting the execution of a task t to an agent a by means of a role r , $granted(a, r, t)$. It appears reasonable that a delegated agent must not be disabled from executing the delegated task. In the LOP case study this can be achieved with

caused $\neg granted(a, r, t)$ **if** $disabled(a, t)$ (7)

to state that $\neg disabled(a, t)$ is an implicit precondition for the delegation of task t to agent a . Notice that this is not equivalent to add $\neg disabled(a, t)$ to the preconditions of the delegation as it corresponds to the fact that the agent is not disabled in the state before the delegation. On the other hand, (7) requires $\neg granted(a, r, t)$ whenever $disabled(a, t)$ holds.

Delegation of tasks has the deterministic effect of granting the execution of the tasks to the delegated agents. Moreover, in case of transfer delegation the agent who performed the delegation is disabled from performing the delegated task in the future. Given $Type$ to be the type of delegation, $Type \in \{transfer, grant\}$, the deterministic effects can be expressed in C using the following pattern

$d(a_1, a_2)$ **causes** $granted(a_1, Arole, Task) \wedge delegated(a_2, Arole, task) \wedge_{Type=transfer} disabled(a_1, Task)$ (8)

for all $a_1 \neq a_2$.

As an example, the deterministic effects of the transfer delegation rule D1 are expressed by

$d1(a_1, a_2)$ **causes** $granted(a_1, supervisor, extRating) \wedge delegated(a_2, supervisor, extRating) \wedge disabled(a_1, extRating)$ (9)

while the deterministic effects of the grant delegation rule D2 are expressed with

$d2(a_1, a_2)$ **causes** $granted(a_1, manager, approve) \wedge delegated(a_2, manager, approve)$.

Notice that the fluent $disabled(a, extRating)$ as effect of rule D1 in (8) represents the means of transfer delegation: the agent performing the delegation is disabled from performing the task in the future.

4.4. Internal Dependencies

Different elements within a business process may be characterized by dependencies. For example results of evaluations performed during the process may influence the value assumed by other fluents describing the state of the process.

In C such dependencies can be expressed by means of static laws. As an example, the inertial fluent $interrupted$ depends on the value of other fluents as follows

caused $interrupted$ **if** $p_9 \wedge invoked(a_1, r, creditBureau) \wedge accessed(a_2)$

for all $a_1 \neq a_2$ and $r \neq director$, and

caused $\neg interrupted$ **if** $p_9 \wedge invoked(a, director, creditBureau)$.

As another example, the statically determined fluent $lowRisk$ has a default value false expressed with

default $\neg lowRisk$

and depends on the value of $highValue$ and $intRatingPositive$ through

caused $lowRisk$ **if** $\neg highValue \wedge intRatingPositive$

to state that the risk of a loan is low if its amount is not high and the internal rating is positive.

4.5. Workflow

As we have already seen in Section 2, the workflow of a business process manages the invocation of entities and the execution of tasks: in our work we model it by adding “dummy” activities to begin/end a flow of business activities and to skip certain operations.

In C these activities can be expressed by dynamic laws. The preconditions of dummy activities only deal with input places and conditions associated to inner arcs, if any. Moreover, being activities that do not require to be performed by an agent and do not perform any task, dummy activities are triggered as soon as their preconditions hold. Consider an activity e , and let I (resp. O) be the set of input (resp. output) places of the related transition, and I' be the set of conditions on the arcs from the input places (if any). Such activities follow the pattern

$$\begin{aligned} \text{nonexecutable } e \text{ if } & \neg(\wedge_{p \in I} p \wedge \wedge_{f \in I'} f), \\ \text{nonexecutable } \neg e \text{ if } & \wedge_{p \in I} p \wedge \wedge_{f \in I'} f, \\ e \text{ causes } & \wedge_{p \in I} \neg p \wedge \wedge_{p' \in O} p'. \end{aligned}$$

As an example, *nop1* is expressed by

$$\begin{aligned} \text{nonexecutable } \textit{nop1} \text{ if } & \neg(p_7 \wedge \neg \textit{highValue}), \\ \text{nonexecutable } \neg \textit{nop1} \text{ if } & p_7 \wedge \neg \textit{highValue}, \\ \textit{nop1} \text{ causes } & p_{10} \wedge \neg p_7. \end{aligned}$$

4.6. Separation of Duty Constraints

SoD represents an important aspect of business processes used to prevent frauds by stating that conflicting tasks must be performed by different agents. With reference to the LOP, SoD constraints in Table 3 can be expressed in C by static laws constraining each set of critical tasks to be executed by at least two different agents. For example, constraint C3 is expressed by

$$\text{constraint } \neg(\textit{executed}(a, r_1, \textit{prepareContract}) \wedge \textit{executed}(a, r_2, \textit{approve}) \wedge \textit{executed}(a, r_3, \textit{sign})).$$

4.7. Security Policy

The security policy establishes which agent can perform which task. As described in Section 2, we consider a security policy based on an RBAC model enhanced with delegation and SoD constraints. We model the permission of an agent a to perform a task t with role r using the fluent *granted*(a, r, t). Then, the RBAC model enhanced with delegation can be expressed with the law

$$\text{caused } \textit{granted}(a, r, t) \text{ if } ua(a, r) \wedge \textit{activated}(a, r) \wedge pa(r_1, t) \wedge \textit{senior}(r, r_1) \quad (10)$$

and

$$\text{caused } \neg \textit{granted}(a, r, t) \text{ if } \neg ua(a, r) \vee \neg \textit{activated}(a, r) \vee \left(\bigwedge_{r_1 \in \mathcal{R}} (\neg pa(r_1, t) \vee \neg \textit{senior}(r, r_1)) \right) \text{ unless } \textit{delegated}(a, r, t) \quad (11)$$

where \mathcal{R} is the set of roles involved in the business process. Notice that (11) is made defeasible by “unless *delegated*(a, r, t)” to express the fact that an agent can be granted the execution of a task even if she does not fulfill the RBAC model assuming she has been delegated to execute it. Also notice that in the overall specification of a business process an agent can be granted the execution of a task only if the SoD constraints specified as presented in Section 4.6 are satisfied.

As it appears from (10), the RBAC model relies both on the user assignment, *ua*(a, r), and the permission assignment, *pa*(r, t). A specific, static, user assignment can be expressed in C by initializing the inertial fluents *ua*(a, r) in the initial state as long as no action affects their values. In this way, the assignment of agents to roles will not change

during the process execution. As an example, the static assignment of agents *davide* and *maria* given in Section 2 can be modeled as follows

$$\begin{aligned} & ua(davide, director), \\ & \neg ua(davide, r_1), \\ & ua(maria, manager), \\ & \neg ua(maria, r_2) \end{aligned}$$

for all $r_1 \neq director$ and $r_2 \neq manager$.

A specific permission assignment can be expressed in C by defining the fluents $pa(r, t)$ as statically determined and specifying static causal laws for them. As an example, an excerpt of the permission assignment in Table 1 can be expressed in C by

$$\begin{aligned} & \mathbf{default} \quad pa(preprocessor, inputCustData), \\ & \mathbf{default} \quad \neg pa(postprocessor, inputCustData), \\ & \mathbf{caused} \quad pa(preprocessor, intRating) \mathbf{if} \quad \neg isIndustrial, \\ & \mathbf{default} \quad \neg pa(preprocessor, intRating), \\ & \mathbf{caused} \quad pa(postprocessor, intRating) \mathbf{if} \quad isIndustrial, \\ & \mathbf{default} \quad \neg pa(postprocessor, intRating). \end{aligned}$$

Notice that the assignment of permissions to roles can change during the process execution according to some conditions.

The partial specifications on user and permission assignments can, in general, be modeled by means of the set of requirements they have to ensure. As an example, for the specific requirements on a user assignment defined in Section 2, the following C laws can be used

$$\begin{aligned} & \mathbf{constraint} \quad \neg(ua(a_1, director) \wedge ua(a_2, director)), \\ & \mathbf{constraint} \quad \neg(ua(a, director) \wedge ua(a, r)), \\ & \mathbf{constraint} \quad \neg(ua(a, supervisor) \wedge ua(a, postprocessor)), \\ & \mathbf{constraint} \quad \neg(ua(a, r_1) \wedge ua(a, r_2) \wedge ua(a, r_3)) \end{aligned}$$

for all $a_1 \neq a_2$, $r_1 \neq r_2$, $r_2 \neq r_3$, $r_1 \neq r_3$, and $r \neq director$.

Analogously, an excerpt of the requirements for the permission assignments defined in Section 2 can be expressed in C with

$$\begin{aligned} & \mathbf{constraint} \quad \neg isIndustrial \vee \neg pa(preprocessor, intRating), \\ & \mathbf{constraint} \quad \neg highProfileIndCust \vee \neg pa(manager, sign). \end{aligned}$$

Notice that in this case the fluents $pa(r, t)$ are defined as exogenous.

4.8. Exceptions

A business process may also be characterized by exceptions to its normal behavior. In the presented scenario the security policy is based on an RBAC model enhanced with delegation. However, we can have exceptions to the RBAC, e.g., in case an important agent (e.g., a director) decides to do so. To model this behavior in C we enhance the model with a law

$$disable(a_1, a_2, t) \mathbf{causes} \quad disabled(a_2, t) \wedge \bigwedge_{r \in \mathcal{R}} (\neg granted(a_2, r, t)) \mathbf{if} \quad ua(a_1, director) \wedge \neg disabled(a_2, t) \quad (12)$$

where $a_1 \neq a_2$, and modifying (10) as follows

$$\mathbf{caused} \quad granted(a, r, t) \mathbf{if} \quad ua(a, r) \wedge activated(a, r) \wedge pa(r_1, t) \wedge senior(r, r_1) \mathbf{unless} \quad disabled(a, t)$$

to make the policy defeasible, i.e., an agent can be granted the execution of a task unless a director disables her from executing the task.

5. Experiments

Given a set of causal laws expressed in C and a query (i.e., a Boolean formulas built out of time-indexed facts), CCALC automatically checks whether there exist paths in the transition system specified by the action description that satisfies the query. The length of the path considered is determined by fixing the *maxstep* variable in CCALC. In particular, CCALC

1. produces a description of the transition system in the form of a set of clauses (where a clause is a disjunction of literals): assuming the head of each causal law is a literal (as in the LOP), this process is done in polynomial time through literal completion [28];
2. produces *maxstep* – 1 copies of the clauses, each copy corresponding to a time step as in planning as satisfiability [29];
3. converts each query into a corresponding set of clauses; the clauses generated so far are such that there exists a 1-1 correspondence between the paths of length *maxstep* of the transition system satisfying the query and the assignments satisfying the set of clauses;
4. calls a SAT solver to determine an assignment satisfying the set of clauses (if any);
5. if the SAT solver returns a satisfying assignment, then the corresponding path is returned to the user.

Notice that there is no 1-1 correspondence between the paths in the extended elementary net of the LOP and the paths in the transition diagram corresponding to the formalization of the LOP as an action description. Indeed, in our formalization we have transitions from one state to the state itself because no action is performed. These loops, called *stuttering steps* (cf., page 17 in [30]), can be practically useful because they allow us to have paths in the transition diagram with length *maxstep* corresponding to paths in the extended elementary net with length \leq *maxstep*. (Because of the security policy, the loop in the LOP can be executed at most once, and thus *maxstep*=14.) Considering that only p_1 is true in the initial states of the transition diagram, we take into account all the possible paths in the extended elementary net leaving from p_1 . Other conditions on paths are imposed depending on the specific problem we consider, detailed in the following subsections. The complete specification is available at <http://www.ai-lab.it/serena/jcss.txt>.

In the rest of this section we describe how we have used CCALC to (i) to establish whether the control flow together with the security policy meets the expected security properties, (ii) to synthesize a security policy for the business process under given security requirements, and (iii) to find a resource allocation plan ensuring the process executability according to the given security policy.

5.1. Verification of Security Properties

The security policy of a business process manages the access of agents to tasks, and should ensure that undesirable behaviors, e.g. frauds, do not occur. In this paper the security policy is given in terms of an RBAC model enriched with delegation rules and object-based SoD constraints. Because of the complexity of the resulting specifications, and the interplay with the security policy, it may not be trivial to establish if other desirable security properties hold (e.g., because entailed by the already enforced security policy) or if it is necessary to revise the model and/or the security policy in order.

In our first experiment we fed CCALC with the specification of the LOP given in Section 4 featuring a security policy characterized by

- a specific RBAC model with the permission assignment of Table 1 and the user assignment presented in Section 2;
- the delegation rules in Table 2; and
- the SoD constraints in Table 3,

and we considered the following property:

“If the process terminates successfully, then no single agent has performed all the tasks *intRating*, *extRating* if *highValue*, *approve*, and *sign*.”

Thus, we used CCALC to determine whether there exists a path leading to a state in which the same agent performs the tasks *intRating*, *extRating* if *highValue*, *approve*, and *sign*, i.e., where the following property is satisfied:

$$p_{18} \wedge \text{productApproved} \wedge \text{executed}(a, \text{intRating}) \wedge (\neg \text{highValue} \vee \text{executed}(a, \text{extRating})) \wedge \text{executed}(a, \text{approve}) \wedge \text{executed}(a, \text{sign}) \quad (13)$$

CCALC found the trace, an excerpt of it is reported in Figure 3. The trace shows that (13) is satisfied, i.e., the violation occurs, if the loan amount is not high, i.e., $\neg \text{highValue}$, and thus the external rating is not evaluated. The perpetrator for the violation is *stefano*, who executes *intRating* as *supervisor* and can nevertheless execute *approve* and *sign* by means of delegation. In fact a manager, *marco*, delegates him to approve the document by means of the delegation rule D2 and the director delegates him to sign the contract by means of the delegation rule D3. By inspecting the intermediate states of the trace it is easy to conclude that the violation occurs if the internal rating is positive and the customer is industrial with a high profile.

To avoid this violation we restricted the applicability conditions of delegation rule D2 by conjoining it with the fact *highValue*. In fact, when the loan has not an high value, the security policy is less restrictive and the application of D2 must be prevented. CCALC does not find any violation in the specification modified in this way.

The verification of SoD properties over a version of the LOP is carried out by using NuSMV in [5]. However the case study considered is rather different, e.g., the workflow (which involves different tasks) structure and the delegation model are different and no exceptions to the RBAC security policy with delegation are considered. Moreover, the specification of the case study and the experimental analysis performed in [5] are not mature, e.g. though an RBAC policy is considered, the assignment between agents and tasks, via roles, is static. Thus this work can be compared with ours wrt the overall approach and the expressiveness of the language used but it is not possible a comparison of the analysis results.

5.2. Synthesis of the Permission Assignment

In our second experiment we synthesize the permission assignment for an RBAC model for the LOP given the requirements for the user and the permission assignment presented in Section 2. Note that the approach we use is general as it relies on the idea of partially defining an aspect of the business process, i.e., only defining a set of requirements we want it to satisfy, and automatically synthesizing a specific instance that ensures the process executability and satisfies the requirements. As an example, in this experiment we consider the problem of synthesizing the permission assignment but the approach could be used as well to synthesize the user assignment.

We fed CCALC with the specification of the LOP given in Section 4 featuring a security policy characterized by

- a generic RBAC model with the requirements for the user and permission assignments given in Section 2; and
- the SoD constraints in Table 3,

and we considered the problem of finding a permission assignment that ensures a successful process execution. The process ends successfully if it reaches a state in which both p_{18} and *productApproved* hold regardless of the value of *intRatingPositive*, *isIndustrial*, *highValue*, and *highProfileIndCust* and the fact that the process has been interrupted. We have therefore run CCALC against all possible scenarios requiring $p_{18} \wedge \text{productApproved}$ and we extracted the permission assignment from the execution traces returned by the tool. Additionally, we looked for a permission assignment such that the process can be executed involving the minimal number of agents.

SoD constraints require at least two agents for performing the critical tasks. CCALC reports that it is not possible to find an execution path for all the scenarios with two agents. Thus, we considered three agents. In this case CCALC was able to find an execution trace for all the scenarios. We then extracted from the execution traces returned by the tool the permission assignment through the values of the fluents $pa(r, t)$. The permission assignment we extracted is reported in Table 7.

5.3. Resource Allocation Plan

Our last experiment was to determine a resource allocation plan ensuring a successful process completion according to the security policy.

We fed CCALC with the specification of the LOP given in Section 4 featuring a security policy characterized by

```

...
4: granted(stefano,supervisor,intRating) isIndustrial p5 p6
   p7 ...
ACTIONS: exec(stefano,supervisor,intRating) nop1 ...

5: isIndustrial intrRatingPositive highProfileIndCust lowrisk
   executed(stefano,intRating) p5 p8 p10 ...
...
7: isIndustrial intrRatingPositive highProfileIndCust
   ua(marco,manager) ua(stefano,supervisor) lowrisk
   executed(stefano,intRating) p12 pa(manager,approve) ...
ACTIONS: d2(marco,stefano) ...

8: granted(stefano,manager,approve) executed(stefano,intRating)
   delegated(stefano,manager,approve) p12 lowrisk ...
ACTIONS: exec(stefano,manager,approve) ...

9: isIndustrial intrRatingPositive highProfileIndCust
   productApproved executed(stefano,intRating) lowrisk
   executed(stefano,approve) p13 ...
...
10: isIndustrial intrRatingPositive highProfileIndCust lowrisk
   productApproved ua(davide,director) ua(stefano,supervisor)
   executed(stefano,intRating) executed(stefano,approve)
   p14 p15 pa(director,sign) ...
ACTIONS: d3(davide,stefano) ...

11: granted(stefano,director,sign) isIndustrial lowrisk
   intrRatingPositive highProfileIndCust productApproved
   executed(stefano,intRating) executed(stefano,approve)
   delegated(stefano,director,sign) p14 p15 ...
ACTIONS: exec(stefano,director,sign) ...

12: isIndustrial intrRatingPositive highProfileIndCust lowrisk
   executed(stefano,intRating) executed(stefano,approve)
   executed(stefano,sign) productApproved p15 p16 ...
...
14: isIndustrial intrRatingPositive highProfileIndCust
   productApproved executed(stefano,intRating) lowrisk
   executed(stefano,approve) executed(stefano,sign) p18 ...

```

Figure 3: An excerpt of the trace violating property (13) found by CCALC

- the specific permission assignment in Table 1;
- a generic user assignment with the requirements given in Section 2;
- the delegation rules in Table 2 with the applicability conditions of D2 restricted according to Section 5.1; and
- the SoD constraints in Table 3,

and we considered the problem of finding for all possible scenarios a resource allocation plan that ensures a successful process completion, i.e., $p_{18} \wedge productApproved$. Notice that as a specific assignment of agents to roles is not given, a possible agent-role assignment can also be obtained from the execution trace returned by CCALC.

As in Section 5.2, the process can take place in different scenarios characterized by different nondeterministic effects. As a result, we run CCALC against all possible successful process completions and we extracted a resource

Table 7: Permission assignment for the LOP automatically synthesized

Task	Role
<i>inputCustData</i>	<i>preprocessor</i>
<i>prepareContract</i>	<i>preprocessor</i>
<i>intRating</i>	if (<i>isIndustrial</i>) then <i>postprocessor</i> else <i>preprocessor</i>
<i>extRating</i>	if (<i>interrupted</i>) then <i>director</i> else <i>preprocessor</i>
<i>approve</i>	<i>director</i>
<i>sign</i>	<i>director</i>
<i>createAccount</i>	<i>postprocessor</i>

Table 8: Resource Allocation Plan for all scenarios where the process is interrupted

Agent	Roles	Tasks
a_1	<i>preprocessor, supervisor</i>	<i>inputCustData, intRating, prepareContract, createAccount</i>
a_2	<i>director</i>	<i>extRating, approve, sign</i>

allocation plan (i.e., an assignment of agents to tasks) from the execution trace returned by the tool. We considered the further requirement of finding resource allocation plans involving the minimal number of agents.

Considering two agents, say a_1 and a_2 , CCALC found an execution trace for all the scenarios where

- the process is interrupted; or
- *intRatingPositive* and *isIndustrial* hold while *highValue* and *highProfileIndCust* do not; or
- *intRatingPositive* holds while *highValue*, *isIndustrial*, and *highProfileIndCust* do not.

The resource allocations found by CCALC for each of these scenarios are reported in Tables 8, 9, and 10, respectively. However CCALC could not find an unique execution trace for all other scenarios involving only two agents.

Thus, we considered three agents and CCALC found an execution trace for all the remaining scenarios. The resource allocation plans extracted from the execution traces are reported in Tables 11, 12, 13, 14. In this way we obtained assignments of agents to roles for every scenario in which the process can be executed. Furthermore, we found that the minimal number of agents depends on the specific scenario and, for each scenario, we also obtained the assignments of agents to roles that allow them to complete the process.

Table 9: Resource allocation plan for all scenarios where \neg *highValue*, *intRatingPositive*, *isIndustrial*, and \neg *highProfileIndCust*

Agent	Roles	Tasks
a_1	<i>manager, supervisor</i>	<i>intRating, approve, sign</i>
a_2	<i>preprocessor, supervisor</i>	<i>inputCustData, prepareContract, createAccount</i>

As far as computational aspects of our approach are concerned, we herewith briefly mention what are the CPU times that CCALC spent for solving the three tasks of interest. On a Pentium IV 1.6GHz machine with 3GB of RAM, CCALC takes from 30 to 50 seconds to ground the specifications in input and create the corresponding set of clauses. A satisfying assignment is then found in negligible time by running a modern, e.g., MINISAT, SAT solver on ϕ .

6. A Comparison between C and the SMV Language

In this section we analyze the languages C and SMV, which are supported by CCALC and NuSMV respectively. The comparison focuses on the ability of the two languages to manage changes and updates of the specifications.

Here we provide a very brief introduction of the SMV language focusing on the aspects that are relevant for our comparison. A detail account of the language can be found in [31]. In a SMV specification the system under design

Table 10: Resource allocation plan for all scenarios where $\neg highValue$, $intRatingPositive$, $\neg isIndustrial$, and $\neg highProfileIndCust$

Agent	Roles	Tasks
a_1	<i>manager, postprocessor</i>	<i>prepareContract, createAccount</i>
a_2	<i>preprocessor, manager</i>	<i>inputCustData, intRating, approve, sign</i>

Table 11: Resource allocation plan for all scenarios where $highValue$ and $\neg highProfileIndCust$

Agent	Roles	Tasks
a_1	<i>preprocessor, supervisor</i>	<i>inputCustData, intRating, prepareContract</i>
a_2	<i>manager, supervisor</i>	<i>extRating, createAccount, sign</i>
a_3	<i>director</i>	<i>approve</i>

is represented as a Kripke structure defined by state variables and transitions. State variables are defined using the keyword **VAR** and, for the scope of our analysis, are boolean. We also consider input variables (defined using the keyword **IVAR**) that are used to represent values given as input to the system and are also used to label transitions. SMV supports two styles for declaring transitions: the assignment style and the constraint style. The *assignment style* is based on assignments of initial and next values of each variable using the keyword **ASSIGN**. (If no assignment is specified for a variable, then the value of the variable evolves nondeterministically.) The *constraint style* allows to specify conditions over variable values on the initial states, the states, and the transitions by using the constructs **INIT**, **INVAR** e **TRANS**, respectively. In our comparison we use the constraint style as it is more flexible. Moreover, specifications in assignment style can be easily rewritten as an equivalent specification in constraint style.

C and SMV differ in fundamental way. In C , if there is no cause for a fact, the fact can be neither true nor false (and thus the formula corresponding to the specification is unsatisfiable). In SMV declared facts are exogenous. The different semantics given to the facts declared has a considerable impact on the way a model can be specified to obtain the same behavior in both languages. As an example consider a simple scenario (1) where an agent is granted the execution of a task if she is a potential owner of the task (*potentialOwner*) or delegated to perform it (*delegated*), while she is not granted the execution if she is not authorized to execute it (*unassigned*).

This scenario can be described in C by the specification in Figure 4 where either *potentialOwner*, or *delegated*, or

- 1 **simpleFluent** *potentialOwner, delegated, unassigned, granted*,
- 2 **exogenous** *potentialOwner, delegated, unassigned*,
- 3 **caused** *granted* **if** *potentialOwner*,
- 4 **caused** *granted* **if** *delegated*,
- 5 **caused** $\neg granted$ **if** *unassigned*,

Figure 4: C specification of scenario (1)

unassigned have to hold in order to provide a cause for *granted* to be either true or false. To specify the same behavior in SMV it is necessary to explicitly state the dependencies that exist between the facts *potentialOwner*, *delegated*, and *unassigned* as they all determine the value of *granted* as shown in Figure 5. A straightforward consequence is that while modifications to the C specification of Figure 4 can be done incrementally, this is not the case for the SMV specification. As an example we can imagine to extend the previous scenario by introducing a mechanism that according to some conditions (abstracted away in our example) disables an agent, i.e. prevents the execution of tasks, namely a scenario (2). In C this can be specified as shown in Figure 6 by adding to the specification in Figure 4 an action *disable* that causes *granted* to be false, i.e. incrementally adding the lines 3 and 7 in Figure 6. In SMV this new scenario can be expressed as shown in Figure 7. It is clear that the SMV specification is not incremental and the fact that the action *disable* is executed if neither *potentialOwner*, *delegated*, nor *unassigned* hold must be explicit (line 12 of Figure 7).

Many behaviors that can be obtained implicitly in C specifications must be made explicit in SMV. As a further

Table 12: Resource allocation plan for all scenarios where *highValue* and *highProfileIndCust*

Agent	Roles	Tasks
a_1	<i>preprocessor, postprocessor</i>	<i>inputCustData, intRating, createAccount</i>
a_2	<i>preprocessor, supervisor</i>	<i>extRating, prepareContract</i>
a_3	<i>director</i>	<i>approve, sign</i>

Table 13: Resource allocation plan for all scenarios where \neg *highValue* and *highProfileIndCust*

Agent	Roles	Tasks
a_1	<i>preprocessor, supervisor</i>	<i>inputCustData, sign</i>
a_2	<i>manager, supervisor</i>	<i>intRating, prepareContract, createAccount</i>
a_3	<i>director</i>	<i>approve</i>

1 **VAR**2 *potentialOwner* : *boolean*,3 *delegated* : *boolean*,4 *unassigned* : *boolean*,5 *granted* : *boolean*,6 **INVAR** *potentialOwner* \rightarrow *granted*,7 **INVAR** *delegated* \rightarrow *granted*,8 **INVAR** *unassigned* $\rightarrow!$ *granted*,9 **INVAR** *potentialOwner* | *delegated* | *unassigned* .

Figure 5: SMV specification of scenario (1)

1 **simpleFluent** *potentialOwner, delegated, unassigned, granted*,2 **exogenous** *potentialOwner, delegated, unassigned*,3 **exogenousAction** *disable*4 **caused** *granted* **if** *potentialOwner*,5 **caused** *granted* **if** *delegated*,6 **caused** \neg *granted* **if** *unassigned*,7 *disable* **causes** \neg *granted* .

Figure 6: C specification of scenario (2)

example consider a scenario (3) where agents are not granted the execution of tasks by default unless they are assigned this duty by, e.g., an administrator, or they are delegated. This scenario can be specified in *C* as shown in Figure 8. In this example, *assigned* is exogenous while *granted* has a negative default value unless *assigned* or *delegate* cause it to be true. This behavior can be specified in SMV as in Figure 9 by explicitly stating when the value of *granted* is negative (line 8 in Figure 9). Imagine that we now want to modify the scenario by allowing an agent granted the execution of a task to retain this authorization, namely a scenario (4). In *C* this can be simply done by incrementally adding **inertial** *granted* to the specification in Figure 8. The main difference of the new specification is that after *granted* has become true, either because *delegate* has been executed or *assigned* was true, it can retain the true value other than become false because of the default. To apply the same modification in SMV the specification changes considerably in the way *granted* is specified (lines 8 and 9 in Figure 10).

Similar scenarios occur also in the LOP case study presented in Section 2. As an example we can consider the specification of the authorization to execute task *extRating* with respect to delegation. In this case it is again possible to observe that SMV requires to explicitly state all the mechanism which can be left implicit in *C*. In particular, in the *C* specification presented in Section 4, *delegated*(*a, r, t*) prevents the authorization to execute to be denied when an agent is delegated. More in detail, *delegated*(*a, r, t*) is an inertial fact which is set to true after D1 is executed and has

Table 14: Resource allocation plan for all scenarios where $\neg highValue$, $\neg intRatingPositive$, and $\neg highProfileIndCust$

Agent	Roles	Tasks
a_1	<i>preprocessor, postprocessor</i>	<i>inputCustData, createAccount</i>
a_2	<i>preprocessor, postprocessor</i>	<i>intRating, prepareContract</i>
a_3	<i>director</i>	<i>approve, sign</i>

- 1 **VAR**
- 2 *potentialOwner : boolean,*
- 3 *delegated : boolean,*
- 4 *unassigned : boolean,*
- 5 *granted : boolean,*
- 6 **IVAR**
- 7 *disable : boolean,*
- 8 **INVAR** *potentialOwner \rightarrow granted,*
- 9 **INVAR** *delegated \rightarrow granted,*
- 10 **INVAR** *unassigned \rightarrow !granted,*
- 11 **TRANS** *disable \rightarrow next(granted) = 0,*
- 12 **INVAR** *!(potentialOwner | delegated | unassigned) \rightarrow disable .*

Figure 7: SMV specification of scenario (2)

- 1 **simpleFluent** *granted, assigned*
- 2 **exogenous** *assigned*
- 3 **default** \neg *granted*
- 4 **exogenousAction** *delegate*
- 5 **caused** *granted if assigned*
- 6 *delegate causes granted .*

Figure 8: SMV specification of scenario (3)

- 1 **VAR**
- 2 *granted : boolean;*
- 3 *assigned : boolean;*
- 4 **IVAR**
- 5 *delegate : boolean;*
- 6 **INVAR** *assigned \rightarrow granted;*
- 7 **TRANS** *delegate \rightarrow next(granted) = 1;*
- 8 **TRANS** *!assigned & !delegate \rightarrow !granted;*

Figure 9: C specification of scenario (3)

a false default value through the **unless** abbreviation in (11), and thus its behavior resembles the one of specification in Figure 8 after the modification according to the scenario (4). As a consequence, its translation in SMV recalls the specification in Figure 10. Notice that SMV is a propositional language and thus each *C* predicate has to be specified

```

1 VAR
2 granted : boolean;
3 assigned : boolean;
4 IVAR
5 delegate : boolean;
6 INVAR assigned → granted;
7 TRANS delegate → next(granted) = 1;
8 TRANS granted → next(granted) = 1 | next(granted) = 0;
9 TRANS !granted & !assigned & !delegate → !granted;

```

Figure 10: SMV specification of scenario (4)

for each ground instance, e.g.

```

INIT !delegated(pierS ilvio, pierPaolo, supervisor, extRating);
TRANS d1_pierS ilvio_pierPaolo_supervisor_extRating →
next(delegated_pierS ilvio_pierPaolo_supervisor_extRating) = 1;
TRANS delegated_pierS ilvio_pierPaolo_supervisor_extRating →
next(delegated_pierS ilvio_pierPaolo_supervisor_extRating) = 1 |
next(delegated_pierS ilvio_pierPaolo_supervisor_extRating) = 0;
TRANS !delegated_pierS ilvio_pierPaolo_supervisor_extRating &
!d1_pierS ilvio_pierPaolo_supervisor_extRating →
next(delegated_pierS ilvio_pierPaolo_supervisor_extRating) = 0;

```

where the ground predicates, e.g. $delegated(pierS\ ilvio, pierPaolo, supervisor, extRating)$, stand for the corresponding propositional letters.

7. Related Work

A preliminary version of this work is [1]. Here we have significantly extended [1] by providing: (i) a mapping from elementary net and security policy into C, which also clearly shows how the specification of the workflow and the security policy can be kept separate; (ii) an extended experimental analysis; (iii) a comparison between C and SMV, which highlights some of the main features of the language C; (iv) an exhaustive analysis of the related works that covers all main aspects of our work; and (v) a detailed description of the reasoning tasks we analyze.

7.1. Action-based Specification Languages and Reasoning Tools

Action languages [11] serve for describing changes that are caused by performing actions. As predecessor of C, action language \mathcal{A} , i.e., the propositional fragment of ADL, was presented in [32] to enhance the expressive power of the STRIPS planning language [33] by allowing conditional effects. Action language \mathcal{B} extended \mathcal{A} with static laws. In action language C nondeterministic actions and the concurrent execution of actions are more conveniently described than in \mathcal{B} .⁴ Moreover, with C we are free to decide for each fluent whether or not to postulate inertia for it. Given it is based on the theory of causal explanation proposed in [35], C distinguishes between asserting that a fact “simply” holds, and the stronger assertion that “it is caused”. C+ [28] further extended C by allowing fluents to be multi-valued, other than some other (minor) changes. The language C+ has been then extended in several directions, e.g., to include “additive fluents” [36], to represent numeric-valued fluents [37] and to include the possibility of referring to other action descriptions in the definition of a new action domain [38]. The applicability of action language C/C+ spans from the representation of “classical” AI problems [39], to planning [40], multi-agent domains [41] and robotics [42], coupled with the reasoning capabilities of C_{CALC}. C has been also used as “basic language” to

⁴In [34] is presented an alternative extension of the language \mathcal{A} to enable concurrent execution of actions.

compare [43] and update [44] action domain descriptions. As far as other knowledge representation languages are concerned, we herewith mention three more languages. The language \mathcal{K} [10] is a declarative planning language based on principles and methods of logic programming. Its distinguished feature is that \mathcal{K} describes transitions between “states of knowledge” rather than between “states of the world” used in language $C/C+$. Like $C/C+$, it uses a notion of causation. Correspondences between (fragments of) \mathcal{K} and C are presented in the mentioned paper. Temporal Action Logics Language TAL [12] uses a surface language $\mathcal{L}(ND)$ (Narrative Description Language) to provide a high-level notation of “narratives”, i.e., collections of statements. There is the possibility to describe both the static and dynamic aspects of a narrative. A narrative described in $\mathcal{L}(ND)$ is translated, after some steps, into a logically equivalent first order theory. Dedicated algorithms are then used to reason about narratives. Finally, in [45] Concurrent Transaction Logic ($CT\mathcal{R}$) is proposed, and used as a language for specifying, analyzing and scheduling of workflows.

Action language C and system $CCALC$ have been already used in the context of business processes [46, 47]. In [46] the research objective is to describe, simulate, compose, verify, and develop Web services. However, this work does not consider (complex) security policies to be modeled or verified. On the contrary, our approach takes into account security policies and focuses on modeling and reasoning about business processes with the objectives of verifying security properties, synthesising policies, and finding resource allocation plans. In [47] the author considers activities with duration and the cost of a workflow execution but, differently from our approach, he does not take into account complex security policies (i.e., agents are statically assigned to tasks) and mandatory requirements, e.g., SoD. This means that it does not consider delegations either. The model is simpler than the one we consider, e.g., indirect effects are not considered.

7.2. Automatic Verification of Business Processes under Authorization Constraints

The use of model checking for the automatic analysis of business processes has been put forward and investigated in [5]. The paper shows that business processes with RBAC policies and delegation can be formally specified as transition systems and that SoD properties can be formally expressed as Linear Temporal Logic formulas specifying the allowed behaviors of the transition systems. The viability of the approach is shown through its application to a version of the LOP and the NuSMV [24] model checker is used to carry out the verification. However the LOP version considered in [5] is much simpler and does not feature many of the aspects we consider, as already underlined in Section 5.1. Our approach, by using an action language, allows for a more natural and concise modeling of the business process and of the associated security policy. Moreover, our approach allows for the separate specification of the workflow and of the security policy while this is not the case for the approach presented in [5], where even small changes in the workflow or in the security policy may affect the specification of the whole transition system. Our approach therefore considerably simplifies the specification process, thereby reducing the probability of introducing bugs in the specification.

A formal framework that uses the SAL model checker [48] to analyze SoD properties as well as to synthesize a resource allocation plan for a business process has been presented in [6]. However, differently from our approach, this framework does not offer a natural modeling and RBAC is the only access control model supported (with tasks rigidly associated with specific roles). Moreover, this work does not take into account the global state of the process and assumes an interleaving semantics. This is not the case in our approach as it accounts for a global state that can be affected by the execution of the tasks as well as for multiple actions to be executed simultaneously.

An approach to the combined modeling of business workflows with RBAC models is presented in [7]. The paper proposes an extended finite state machine model that allows for the model checking of SoD properties by using the model checker SPIN [49]. It considers a simple RBAC model only based on previous activation (or non-activation) of roles and it does not take into account delegation. A number of techniques to alleviate the state explosion problem are presented.

Some computational techniques for analyzing SoD by integrating workflows of the enterprise processes into the RBAC framework is presented in [21]. It proposes an algorithm for generating mutually exclusive roles (MER) to enforce SoD and verification algorithms to check if a role authorization together with a user-role assignment satisfy static and dynamic SoD constraints. The algorithm proposed to check for dynamic SoD also returns a set of user-role activations that guarantees to satisfy the SoD constraints. However the approach does not allow for the separate specification of the workflow and the security policy as the workflow in input to the algorithms captures aspects related to the policy, e.g. a conditional permission assignment has to be captured by duplicating the task and adding an if condition to the workflow structure. Moreover it does not consider hierarchy of roles that by allowing senior roles to

perform tasks assigned to junior roles would considerably affect the proposed approach. Also notice that MER is a way to enforce a subset of SoD that, e.g., is not sufficient for the kind of SoD constraints we consider in our work. Finally, [21] does not consider the problem of synthesizing the permission assignment relation while this is the case in our approach.

The paper [22] proposes to model workflows and security policies in a notation based on Colored Petri nets, with an automatic translation from the process model into a specification language and the usage of SPIN to verify SoD properties. However, no provision is made for the assignment of an agent to multiple roles, role hierarchy, delegation, and the global state of the process.

Other approaches that use (Colored) Petri Nets in this context are [51, 52]. The first paper introduced a model for SoD in workflows that are specified with Petri nets, and allows for simulating and analyzing workflows and security rules at build time; rules are given as facts of a logic program and expressed in propositional logic. However it does not take into account, e.g., role hierarchy or conditions in the permission assignment and the reasoning task of synthesizing the permission assignment. [52], which is much more recent than [51], presents a formal technique to model and analyze RBAC using Colored Petri nets (CP-nets). The resulting CP-net model can be composed with context-specific aspects of the application of interest, e.g. the workflow, however the approach focuses on the access control policy and does not allow for the separate specification of the workflow and the access control policy in a modular way. Moreover the use of CP-nets does not allow the use of advanced features as, e.g., implicit preconditions of tasks. A graphical representation of the CP-nets models can be provided by using CPN tools.⁵ However the approach does consider the problems of synthesizing permission assignments and resource allocation plans.

A security validation approach for business processes that employs state-of-the-art model checking techniques and makes them usable by business analysts is presented in [59]. The paper considers business processes expressed in BPMN enhanced with an application-dependent security policy, e.g. delegation of tasks, agent substitution, and fully automates their translation into a formal model suitable to formal analysis while offering graphical user interfaces to define security properties and easy-to-understand feedback for the business analysts. However the approach does not consider some business processes features which are relevant for an accurate security analysis, e.g., non-deterministic, indirect and conditional effects of tasks. This is due to the fact that these aspects are not directly available in the industrially suited languages, as BPMN, as their focus is on the definition of the procedural behavior of the process, e.g. they define as output of a task all the data that can be affected but they do not specify how and when their values can be affected. It would thus be of great interest to (i) evaluate extensions to the languages used in the industrial environments in order to support the declarative specification of advanced, relevant features and (ii) automate the process of transformation of the extended languages into *C*. As an example, BPMN could be enhanced with annotations supporting a declarative, *C*-like specification of, e.g., non deterministic, indirect and conditional effects of tasks. The resulting model could then be automatically translated into a *C* specification as the one presented in Section 4. Notice that the automation of the translation would make it language- and application-dependent, e.g. [59] relies on BPMN and the industrial environment SAP Netweaver BPM [60]. We leave these research directions for future work.

7.3. Frameworks for Security Policies

An approach based on model checking for the analysis and synthesis of complex security policies is presented in [19] and further developed in [20]. In particular, they consider fine-grained policies where permissions are the ability of agents to access resources to read or write. They propose the RW (Read and Write) access control formalism based on propositional logic and define a machine-readable language to express policies modelled in the RW formalism and properties to be verified against the model. The RW model checking algorithm and a related tool which implements the algorithm are presented. The proposed framework is expressive and allows for the specification of administrative policies either. However, there is no way to express mutual exclusions between the values of different variables, e.g., it is not possible to express a rule of the form “an agent cannot be a student and a lecturer at the same time”, and there is no way to express inheritance between roles. Moreover, such works do not take into account the natural interactions between the workflow and the security policy while this is the case in our approach.

In [53, 54, 55, 56] rule-based policy specifications are presented and evaluated. In particular, [56] first presents a broad view of different policy languages, and then a system for specifying and enforcing security and private policy,

⁵See, e.g., <http://wiki.daimi.au.dk/cpntools/cpntools.wiki>.

called PROTUNE [55]. However the focus of these works is on policies languages and enforcement from a run-time point of view, usually in the context of the Semantic Web, i.e., they evaluate the existing languages and propose a system for specifying and cooperatively enforcing security and privacy policies in open environments such as the Web, where parties may get in touch without being previously known to each other. On the contrary our work focuses on security policies in the context of business processes and thus wrt the execution of tasks within a workflow and on a security analysis to be performed prior to the deployment of the system.

8. Conclusions

The design and verification of business processes under authorization constraints is a time consuming and error-prone activity. Moreover, due to the complexity that business processes subject to a security policy may reach, it can be difficult to verify even basic properties such as the executability of the process wrt the available resources by manual inspection only, or by simulation.

In this paper we have presented an action-based approach to the formal specification and automatic analysis of business processes under authorization constraints. With our approach we have been able to both greatly simplify the specification activity, and allow for the separate specification of the workflow and of the associated security policy, while retaining the ability to perform a fully automatic analysis of the specifications by using the Causal Calculator CCALC. The experiments we have presented indicate that our approach can be profitably used to execute a number of reasoning tasks particularly important from the application viewpoint, such as verify the process executability, synthesize a permission assignment and identify resource allocation plans complying with the security policy.

Acknowledgements. This work is partially supported by the FP7-ICT-2007-1 Project no. 216471, “AVANTSSAR: Automated Validation of Trust and Security of Service-oriented Architectures” (www.avantssar.eu) and by the project “Automated Security Analysis of Identity and Access Management Systems (SIAM)” funded by Provincia Autonoma di Trento in the context of the “Team 2009 - Incoming” COFUND action of the European Commission (FP7).

References

- [1] A. Armando, E. Giunchiglia, S. E. Ponta, Formal Specification and Automatic Analysis of Business Processes under Authorization Constraints: An Action-based Approach, in: S. Fischer-Hübner, C. Lambrinouidakis, G. Pernul (Eds.), Proc. of 6th International Conference on Trust, Privacy & Security in Digital Business (TrustBus'09), vol. 5695 of *Lecture Notes in Computer Science*, Springer, 63–72, 2009.
- [2] J. L. Peterson, Petri Net Theory and the Modeling of Systems, Prentice Hall PTR, 1981.
- [3] OMG, Business Process Modeling Notation (BPMN), <http://www.omg.org/spec/BPMN/2.0>, 2009.
- [4] R. S. Sandhu, E. J. Coyne, H. L. Feinstein, C. E. Youman, Role-based access control models, *Computer* 29 (2) (1996) 38–47.
- [5] A. Schaad, V. Lotz, K. Sohr, A model-checking approach to analysing organisational controls in a loan origination process, in: D. F. Ferraiolo, I. Ray (Eds.), Proc. of 11th ACM Symposium on Access Control Models and T (SACMAT 2006), ACM, 139–149, 2006.
- [6] A. Cerone, Z. Xiangpeng, P. Krishnan, Modelling and Resource Allocation Planning of BPEL Workflows under Security Constraints, Tech. Rep. 336, UNU-IIST, <http://www.iist.unu.edu/>, 2006.
- [7] A. Dury, S. Boroday, A. Petrenko, V. Lotz, Formal Verification of Business Workflows and Role Based Access Control Systems, in: L. Peñalver, O. A. Dini, J. Mulholland, O. Nieto-Taladriz (Eds.), Proc. of First International Conference on Emerging Security Information, Systems and Technologies, (SECURWARE 2007), 201–210, 2007.
- [8] C. Baral, Knowledge Representation, Reasoning and Declarative Problem Solving, Cambridge University Press, 2003.
- [9] M. Gelfond, V. Lifschitz, Representing Action and Change by Logic Programs, *Journal of Logic Programming* 17 (2/3&4) (1993) 301–321.
- [10] T. Eiter, W. Faber, N. Leone, G. Pfeifer, A. Polleres, A logic programming approach to knowledge-state planning: Semantics and complexity, *ACM Transaction on Computational Logic* 5 (2) (2004) 206–263.
- [11] M. Gelfond, V. Lifschitz, Action Languages, *Electronic Transaction on Artificial Intelligence* 2 (1998) 193–210.
- [12] P. Doherty, J. Gustafsson, L. Karlsson, J. Kvarnström, TAL: Temporal Action Logics Language Specification and Tutorial, *Electronic Transaction on Artificial Intelligence* 2 (1998) 273–306.
- [13] T. Eiter, W. Faber, N. Leone, G. Pfeifer, A. Polleres, A logic programming approach to knowledge-state planning: Semantics and complexity, *ACM Transactions on Computational Logic* 5 (2) (2004) 206–263.
- [14] P. Simons, I. Niemelä, T. Sooinen, Extending and implementing the stable model semantics, *Artificial Intelligence* 138 (2002) 181–234.
- [15] N. Leone, G. Pfeifer, W. Faber, T. Eiter, G. Gottlob, S. Perri, F. Scarcello, The DLV system for knowledge representation and reasoning, *ACM Transactions on Computational Logic* 7 (3) (2006) 499–562.
- [16] T. Eiter, W. Faber, N. Leone, G. Pfeifer, A. Polleres, A logic programming approach to knowledge-state planning, II: The DLV^K system, *Artificial Intelligence* 144 (1-2) (2003) 157–211.
- [17] Texas Action Group at Austin, The Causal Calculator, available at <http://www.cs.utexas.edu/users/tag/cc/>, 2009.
- [18] E. Giunchiglia, V. Lifschitz, An action language based on causal explanation: Preliminary report, in: Proc. of Fifteenth National Conference on Artificial Intelligence (AAAI 1998), AAAI Press, 623–630, 1998.

- [19] D. P. Guelev, M. Ryan, P.-Y. Schobbens, Model-Checking Access Control Policies, in: K. Zhang, Y. Zheng (Eds.), Proc. of 7th International Conference on Information Security (ISC 2004), vol. 3225 of *Lecture Notes in Computer Science*, Springer, 219–230, 2004.
- [20] N. Zhang, M. Ryan, D. P. Guelev, Synthesising verified access control systems through model checking, *Journal of Computer Security* 16 (1) (2008) 1–61.
- [21] R. Hewett, P. Kijsanayothin, A. Thipse, Security Analysis of Role-based Separation of Duty with Workflows, in: Proceedings of the The Third International Conference on Availability, Reliability and Security (ARES 2008), IEEE Computer Society, 765–770, 2008.
- [22] C. Wolter, P. Miseldine, C. Meinel, Verification of Business Process Entailment Constraints Using SPIN, in: F. Massacci, S. T. R. Jr., N. Zannone (Eds.), Proc. of First International Symposium Engineering Secure Software and Systems (ESSoS 2009), vol. 5429 of *Lecture Notes in Computer Science*, Springer, 1–15, 2009.
- [23] A. Cimatti, E. Clarke, F. Giunchiglia, M. Roveri, NuSMV: A New Symbolic Model Verifier, LNCS 1633 ISSN 0302-9743, URL <http://nusmv.irst.itc.it/index.html>.
- [24] A. Cimatti, E. M. Clarke, E. Giunchiglia, F. Giunchiglia, M. Pistore, M. Roveri, R. Sebastiani, A. Tacchella, NuSMV 2: An OpenSource Tool for Symbolic Model Checking, in: E. Brinksma, K. G. Larsen (Eds.), 14th International Conference on Computer Aided Verification (CAV 2002), vol. 2404 of *Lecture Notes in Computer Science*, Springer, 359–364, 2002.
- [25] S. Frau, R. Gorrieri, C. Ferigato, Petri Net Security Checker: Structural Non-interference at Work, in: P. Degano, J. D. Guttman, F. Martinelli (Eds.), Formal Aspects in Security and Trust, vol. 5491 of *Lecture Notes in Computer Science*, Springer, 210–225, 2008.
- [26] V. Atluri, J. Warner, Supporting conditional delegation in secure workflow management systems, in: E. Ferrari, G.-J. Ahn (Eds.), Proc. of 10th ACM Symposium on Access Control Models and Technologies (SACMAT 2005), ACM, 49–58, 2005.
- [27] J. Crampton, H. Khambhammettu, Delegation in Role-Based Access Control, in: Proc. of 13th European Symposium on Research in Computer Security (ESORICS 2006), vol. 5283 of *Lecture Notes in Computer Science*, Springer, 174–191, 2006.
- [28] E. Giunchiglia, J. Lee, V. Lifschitz, N. McCain, H. Turner, Nonmonotonic causal theories, *Artificial Intelligence* 153 (1-2) (2004) 49–104.
- [29] H. Kautz, B. Selman, Planning as satisfiability, in: B. Neumann (Ed.), Proc. of 10th European Conference on Artificial Intelligence (ECAI), IOS Press, 359–363, 1992.
- [30] L. Lamport, Specifying Systems, The TLA+ Language and Tools for Hardware and Software Engineers, Addison-Wesley, 2002.
- [31] R. Cavada, A. Cimatti, C. A. Jochim, G. Keighren, E. Olivetti, M. Pistore, M. Roveri, A. Tchaltev, NuSMV 2.5 User Manual, available at <http://nusmv.fbk.eu>, 2002.
- [32] E. Pednault, Formulating multi-agents, dynamic world problems in the classical planning framework, in: M. Georgeff, A. Lansky (Eds.), Proc. of Reasoning about Actions and Plans, Morgan Kaufmann, 47–82, 1987.
- [33] R. Fikes, N. J. Nilsson, STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving, *Artificial Intelligence* 2 (3-4) (1971) 189–208.
- [34] C. Baral, M. Gelfond, Reasoning About Effects of Concurrent Actions, *Journal of Logic Programming* 31 (1-3) (1997) 85–117.
- [35] N. McCain, H. Turner, Causal Theories of Action and Change, in: Proc. of Fourteenth National Conference on Artificial Intelligence (AAAI 1997), AAAI Press, 460–465, 1997.
- [36] J. Lee, V. Lifschitz, Describing Additive Fluents in Action Language C+, in: G. Gottlob, T. Walsh (Eds.), Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI-03), Morgan Kaufmann, 1079–1084, 2003.
- [37] E. Erdem, A. Gabaldon, Representing Action Domains with Numeric-Valued Fluents, in: M. Fisher, W. van der Hoek, B. Konev, A. Lisitsa (Eds.), Proc. of 10th European Conference on Logics in Artificial Intelligence (JELIA 2006), vol. 4160 of *Lecture Notes in Computer Science*, Springer, 151–163, 2006.
- [38] V. Lifschitz, W. Ren, A Modular Action Description Language, in: Proc. Twenty-First National Conference on Artificial Intelligence (AAAI 2006), AAAI Press, 2006.
- [39] V. Akman, S. T. Erdogan, J. Lee, V. Lifschitz, H. Turner, Representing the Zoo World and the Traffic World in the language of the Causal Calculator, *Artificial Intelligence* 153 (1-2) (2004) 105–140.
- [40] C. Castellini, E. Giunchiglia, A. Tacchella, SAT-based planning in complex domains: Concurrency, constraints and nondeterminism, *Artificial Intelligence* 147 (1-2) (2003) 85–117.
- [41] C. Baral, T. Son, E. Pontelli, Modeling Multi-Agent Domains in an Action Languages: an Empirical Study using C, in: E. Erdem, F. Lin, T. Schaub (Eds.), Proc. of 10th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2009), vol. 5753 of *Lecture Notes in Computer Science*, Springer, 409–415, 2009.
- [42] O. Caldiran, K. Haspalamutgil, A. Ok, C. Palaz, E. Erdem, V. Patoglu, Bridging the Gap between High-Level Reasoning and Low-Level Control, in: E. Erdem, F. Lin, T. Schaub (Eds.), Proc. of 10th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2009), vol. 5753 of *Lecture Notes in Computer Science*, Springer, 342–354, 2009.
- [43] T. Eiter, E. Erdem, M. Fink, J. Senko, Comparing action descriptions based on semantic preferences, *Annals of Mathematics and Artificial Intelligence* 50 (3-4) (2007) 273–304.
- [44] T. Eiter, E. Erdem, M. Fink, J. Senko, Updating Action Domain Descriptions, in: L. P. Kaelbling, A. Saffiotti (Eds.), Proc. of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI 2005), Professional Book Center, 418–423, 2005.
- [45] H. Davulcu, M. Kifer, C. R. Ramakrishnan, I. V. Ramakrishnan, Logic Based Modeling and Analysis of Workflows, in: Proc. of Seventeenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS 1998), ACM Press, 25–33, 1998.
- [46] A. Chirichiello, Two Formal Approaches for Web Services: Process Algebras & Action Languages., Ph.D. thesis, "Sapienza" University of Roma., 2008.
- [47] P. Koksall, N. K. Cicekli, I. H. Toroslu, Specification of workflow processes using the action description language C, in: AAAI Spring 2001 Symposium Series: Answer Set Programming, 103–109, 2001.
- [48] L. M. de Moura, S. Owre, H. Rueß, J. M. Rushby, N. Shankar, M. Sorea, A. Tiwari, SAL 2, in: R. Alur, D. Peled (Eds.), Proc. of 16th International Conference on Computer Aided Verification (CAV 2004), vol. 3114 of *Lecture Notes in Computer Science*, Springer, 496–500, 2004.
- [49] G. J. Holzmann, The Model Checker SPIN, *Software Engineering* 23 (5) (1997) 279–295.
- [50] A. Herzberg, Y. Mass, J. Mihaeli, D. Naor, Y. Ravid, Access Control Meets Public Key Infrastructure, Or: Assigning Roles to Strangers, in:

- IEEE Symposium on Security and Privacy, 2–14, 2000.
- [51] K. Knorr, H. Weidner, Analyzing Separation of Duties in Petri Net Workflows, in: V. I. Gorodetski, V. A. Skormin, L. J. Popyack (Eds.), International Workshop on Information Assurance in Computer Networks: Methods, Models, and Architectures for Network Security (MMM-ACNS 2001), vol. 2052 of *Lecture Notes in Computer Science*, Springer, 102–114, 2001.
 - [52] H. Rakkay, H. Boucheneb, Security Analysis of Role Based Access Control Models Using Colored Petri Nets and CPNtools, *Transactions on Computational Science* (4), Special Issue on Security in Computing 5430 (2009) 149–176.
 - [53] P. A. Bonatti, D. Olmedilla, Rule-Based Policy Representation and Reasoning for the Semantic Web, in: G. Antoniou, U. Aßmann, C. Baroglio, S. Decker, N. Henze, P.-L. Patranjan, R. Tolksdorf (Eds.), Reasoning Web, Third International Summer School 2007, Tutorial Lectures, vol. 4636 of *Lecture Notes in Computer Science*, Springer, 240–268, 2007.
 - [54] G. Antoniou, M. Baldoni, P. A. Bonatti, W. Nejdl, D. Olmedilla, Rule-based Policy Specification, in: T. Yu, S. Jajodia (Eds.), Secure Data Management in Decentralized Systems, vol. 33 of *Advances in Information Security*, Springer, 169–216, 2007.
 - [55] P. A. Bonatti, J. L. D. Coi, D. Olmedilla, L. Sauro, Policy-Driven Negotiations and Explanations: Exploiting Logic-Programming for Trust Management, Privacy & Security, in: Proc. of the 24th International Conference on Logic Programming (ICLP 2008), vol. 5366 of *Lecture Notes in Computer Science*, Springer, 779–784, 2008.
 - [56] P. A. Bonatti, J. L. D. Coi, D. Olmedilla, L. Sauro, Rule-Based Policy Representations and Reasoning, in: Semantic Techniques for the Web, The REVERSE Perspective, vol. 5500 of *Lecture Notes in Computer Science*, Springer, 201–232, 2009.
 - [57] K. E. Seamons, M. Winslett, T. Yu, B. Smith, E. Child, J. Jacobson, H. Mills, L. Yu, Requirements for Policy Languages for Trust Negotiation, in: Proc. of the 3rd International Workshop on Policies for Distributed Systems and Networks (POLICY 2002), IEEE Computer Society, 68–79, 2002.
 - [58] G. Tonti, J. M. Bradshaw, R. Jeffers, R. Montanari, N. Suri, A. Uszok, Semantic Web Languages for Policy Representation and Reasoning: A Comparison of KAoS, Rei, and Ponder, in: D. Fensel, K. P. Sycara, J. Mylopoulos (Eds.), Proc. of the 2nd International Semantic Web Conference, vol. 2870 of *Lecture Notes in Computer Science*, Springer, 419–437, 2003.
 - [59] W. Arsac, L. Compagna, G. Pellegrino, S. E. Ponta, Security Validation of Business Processes via Model-checking, Submitted to 26th Annual Computer Security Applications Conference (ACSAC 2010), 2010.
 - [60] S. Karch, L. Heilig, SAP NetWeaver, Galileo Press, 1. Aufl. edn., ISBN 3-89842-608-4, URL http://gso.gbv.de/DB=2.1/CMD?ACT=SRCHA&SRT=YOP&IKT=1016&TRM=ppn+397693370&sourceid=fbw_bibsonomy, 2004.