

# Formal Verification of Usage Control Models: A Case Study of UseCON Using TLA+

Antonios Gouglidis

School of Computing and Communications  
Lancaster University  
Lancaster, United Kingdom  
a.gouglidis@lancaster.ac.uk

Christos Grompanopoulos

Department of Mechanical Engineering  
University of Western Macedonia  
Kozani, Greece  
cgrompanopoulos@uowm.gr

Anastasia Mavridou

Institute for Software Integrated Systems  
Vanderbilt University  
Nashville, TN, USA  
anastasia.mavridou@vanderbilt.edu

Usage control models provide an integration of access control, digital rights, and trust management. To achieve this integration, usage control models support additional concepts such as attribute mutability and continuity of decision. However, these concepts may introduce an additional level of complexity to the underlying model, rendering its definition a cumbersome and prone to errors process. Applying a formal verification technique allows for a rigorous analysis of the interactions amongst the components, and thus for formal guarantees in respect of the correctness of a model. In this paper, we elaborate on a case study, where we express the high-level functional model of the UseCON usage control model in the TLA+ formal specification language, and verify its correctness for  $\leq 12$  uses in both of its supporting authorisation models.

## 1 Introduction

Access control systems offer the mechanisms to control and limit the actions or operations that are performed by a user or process – referred to as *subjects* – on a set of system *objects*. Specifically, an authorisation process is required to take a decision for granting or denying a *subject* to access an *object* of the system, based on a set of existing policy rules. Thus, access control systems are considered to be amongst the most critical of security components. Their importance is recently highlighted in NIST’s special publication [4], where verification approaches for various access control models are examined.

New computing paradigms, e.g., the cloud [2], require further investigation of access control systems. This eventually introduces the need for models, such as usage control [3], to cope with complex high-level requirements set by new computing paradigms. Usage control models provide an integration of access control, digital rights, and trust management, which may be applicable in environments such as the cloud. To achieve this integration, usage control models support additional concepts such as attribute mutability and continuity of decision. Attribute mutability is responsible for updating the values (e.g., attribute values) of the participating entities (e.g., *subjects*, *objects*), which results in having a dynamic – usage – control model. Continuity of decision considers that access to an *object* is no longer an instantaneous access or action, but it may last for some time. Consequently, decision factors are evaluated not only before (i.e., pre-authorisation), but also during the exercise of an access (i.e., ongoing-authorisation), and thus evolving the concept of access control to that of usage control.

Temporal logic is applied to provide unambiguous semantics to functions supported by usage models, such as attribute mutability and continuity of decision [6]. Usage control systems are concurrent and characterised by non-determinism due to the potential of a *subject* to arbitrarily request or terminate using an *object* during the operation of a system. As discussed in [10], a number of usage control formal models are limited to the specification of a single use, which is isolated and has no interference with other uses. UseCON [3] is a usage model that provides enhanced expressiveness compared to existing access/usage control models and applicable in new computing paradigms [3]. This is achieved by introducing a new entity entitled *use*, which enables the use of historical information in usage control decision. To ensure the correctness of the UseCON model, we used an existing formal language to verify the correctness of its supported *use* management procedures. For example, ensure out of bound values are not assigned to attribute values, and that authorisation states of the usage model always comply with certain behaviours and follow predefined authorisation transitions (e.g., an authorisation must always be requested before changing into another state). The enhanced version of Temporal Logic of Actions (TLA+) [5] is selected for the formal specification of UseCON – TLA+ has been selected also in formal definitions of other usage models (e.g., UCON [11]). In this paper, we anticipate that the formal specification and verification of UseCON using an automated and error-free model checking technique will provide a comprehensive and unambiguous understanding of the concepts introduced in UseCON and ensure its correctness for a number of uses, depending on the authorisation model (pre-authorisation and ongoing-authorisation).

The remainder of this paper is organised as follows: Section 2 provides a specification of UseCON's high-level requirements in TLA+, and the verification of its system model is described in Section 3. A performance evaluation of the TLC model checker to verify deadlock, safety and liveness properties is given in Section 4. Concluding remarks and future work are provided in Section 5. An overview of UseCON and TLA+ specifications are available in [3].

## 2 Specification of UseCON in TLA+

UseCON is a usage control model that is characterised by extended expressiveness when compared with existing usage-based models (e.g., UCON [11]). Specifically, it may support authorisations that have logical relations between entities (see *direct* and *indirect* entities in [3]), as well as historical information, and can express complicated requests between *subjects* and *objects* through simple policies (see automated management of use entities in [3]). In the following, we provide a formal specification of the UseCON model in TLA+ (see TLA+ specification in [1]). This includes the specification of UseCON's main elements, decision making rules, and procedures for managing *use* elements. Furthermore, we describe two transition systems (*TS*) of UseCON, i.e., one *TS* for the *pre* and one *TS* for the *ongoing* authorisation model, which are used for supporting concurrent operations amongst uses in UseCON.

### 2.1 Main Elements

A *use* in UseCON represents a request to execute an *action* (*a*) from a *subject* (*s*) on an *object* (*o*). All *subjects*, *objects* and *actions* used in the usage control system define the sets of *subjects* (*S*), *objects* (*O*), and *actions* (*A*), respectively. The set of *subjects* *S*, *objects* *O*, and *actions* *A* are defined, as follows:

$$S \triangleq \{s_1, \dots, s_i\}, O \triangleq \{o_1, \dots, o_j\}, A \triangleq \{a_1, \dots, a_k\}$$

The set of entities *E* is defined as the union of *S*, *O* and *A* sets, as follows:

$$E \triangleq S \cup O \cup A$$

The characteristics of an entity are represented through its attribute values. An attribute is a function whose domain is a particular set ( $S$  or  $O$  or  $A$ ) and its range is composed of specific attribute values, as follows:

$$att_i : X \mapsto RangeAtt_i$$

where  $X = S$ , or  $X = O$ , or  $X = A$ . For every system's entity, an identification attribute  $id$  is defined for assigning a unique value to the entity. The  $id$  value remains the same throughout the life-time of the usage control system. Thus, the following invariant is valid for all the behaviours of the usage control system:

$$\forall e_1, e_2 \in E : id[e_1] = id[e_2] \Rightarrow e_1 = e_2$$

In UseCON, the operation of the usage control system does not modify automatically the attribute values of system entities. Thus, attribute values of system entities are proposed to be updated manually by the usage control administration model. Consequently, a system entity is represented with a constant record having as fields the entity's attribute values, as follows:

$$e \triangleq [id(e) = k, att_1(e) = l_1, att_2(e) = l_2, \dots, att_n(e) = l_n]$$

where  $e \in E$  is a system entity and  $att_i(e)$ , with  $i \in 1, \dots, n$ , represents the value  $l_i$  of one of its attributes. The first record field of each entity is an  $id$  attribute.

A *use* request in UseCON results into the creation of a *use*. A *use* is an entity instantiated by UseCON and describes all the usage requests in a system and it is described with attributes. More specifically, every *use* must contain a *use* id attribute having a tuple composed of the attribute values  $sid$ ,  $oid$ ,  $aid$  that describe the identities of the *subject*, *object*, and *action* participating in the instantiated *use*. A special attribute  $st$  is associated to every *use* instance representing the status of the *use*. Its value may be one of the following: '*requested*', '*activated*', '*denied*', '*stopped*', and '*completed*'.

A *use*  $u$  that instantiates a specific *use* request from subject  $s$  to object  $o$  for action  $a$  is represented in the specification of the model with a variable record, having as fields its attribute values as follows:

$$u \triangleq [sid(u) = s.id, oid(u) = o.id, aid(u) = a.id, st(u) = state, att_1(u) = v_1, att_2(u) = v_2, \dots, att_n(u) = v_n]$$

where  $s.id$ ,  $o.id$  and  $a.id$  are the identity values of the subject, the object, and the action, respectively. The state attribute  $st$  gets a value  $state$  which belongs to the following set:  $state \in \{ 'requested', 'activated', 'denied', 'completed', 'stopped' \}$  and  $att_i(u)$ ,  $i = 1, 2, \dots, n$  are *use* attribute values. The set of all system *uses* is  $U$ . During the operation of the usage control system,  $U$  is populated due to the *use* requests. Moreover, as these *use* requests are served by the usage control system, the *use* attribute values are modified. Specifically, the  $U$  is altered when a new *use* is requested or change its progress status (e.g., from '*activated*' to '*stopped*'). Consequently, the only variable utilised in the specification is  $U$ , which is the set containing all the *uses* operated in the UseCON model and it is declared as follows:

#### VARIABLES $U$

In the beginning of a TLA+ specification, several modules may be included for supporting different operators. Our specification includes *Integers* and *FiniteSets* modules, which encompass arithmetic and set-related operators, like *Cardinality*. This is declared as follows:

*EXTENDS Integers, FiniteSets*

## 2.2 Decision Making in UseCON

Policy rules in UseCON provide an enhanced utilisation of information from entity and *use* attribute values. More precisely, the general form of a UseCON policy rule that governs the allowance of a *use* request from a *subject*  $s$  on an *object*  $o$  with an *action*  $a$ , is a boolean valued expression defined as follows:

$$\text{Policy\_Rule}(s, o, a, S, O, A) \triangleq \text{expression}(e_1, \dots, e_n)$$

where  $s, o, a$  are the particular *direct* entities of the *use*. In addition, two or more UseCON policy rules can be combined together with logical operators as follows:

$$p = p_1 \otimes p_2 \otimes \dots \otimes p_n$$

where  $\otimes$  is a logical operator (e.g., AND, OR), and  $p_i$  is a policy rule, where  $i = 1, \dots, n$ .

The parameters  $e_i: i \in 1, \dots, n$  of a policy rule that are utilised for the evaluation of the *expression* may have various origins, and thus lead to the creation of the following categories of policy rules:

- **Direct Policy Rules:** The parameters  $e_i$  in the expression of a direct policy rule are values only from attributes of *direct* entities or constant values. Specifically, all parameters  $e_i$ , are defined by the following formula:

$$e_i \in \{s, o, a\} \text{ or } e_i \triangleq l$$

where  $l$  is a constant value.

- **Indirect Policy Rules:** The expression of an *indirect* policy rule consists of attribute values stemming not only from *direct*, but also from *indirect* entities. In this case a logical relation exists between the two types of entities, which can be retrieved by a *select* expression. The definition of such an expression is as follows:

$$e_i \triangleq \text{CHOOSE } x \in E: \text{select}(x, s, o, a, l)$$

An example of an indirect policy rule could be an expression that evaluates into true or false, if the father of a child has a premium membership. In this example, the child is a *direct* entity, while the father is an *indirect*. The *select* expression should take into account the fact that there is a 'father' attribute in the child entity having the attribute value of her father's identity.

- **Complex Indirect Policy Rules** New computing paradigms introduce complex access control policies, where the usage decision is based on information related not only to a single entity, but with a subset of entities. Such complex policies can be supported in UseCON through complex indirect policy rules. More specifically, a parameter  $e_i$  of a complex indirect policy rule can be, apart from a single (i.e., direct or indirect) attribute value, an aggregation of information. This information is derived from all the entities that satisfy a desired (*select*) predicate. The semantics of a parameter  $e_i$  of a complex indirect policy rule is as follows:

$$e_i \triangleq \text{aggregation}(\{e \in E: \text{select}(e)\})$$

An example of a complex indirect policy rule is one that confirms that the balance sum of all the accounts of a bank customer is over a specific amount. Information that is related with a set of




Figure 1: Use attributes updates during the exercise of a *use* in UseCON

bank accounts, those belonging to the corresponding user, is required for the evaluation of the aforementioned policy rule. Consequently, the *selection* expression defines the subset of the bank accounts that belongs to the specific customer.

### 2.3 Use Attribute Update Procedures

Attribute values are utilised by UseCON policy rules through the usage decision making process. Consequently, the implementation of a high level policy should also cope with the definition of *use* attribute value update procedures because these values determine the outcome of the policy rule. In UseCON the mutation of attribute values only records security related information that are related with the *use*. Moreover, a categorisation of the use attributes according to the nature of information they record is as follows:

- *Induced Attributes*: Information that can be inferred from entities involved in a *use* (i.e., *subject*, *object*, and *action*) that a specific use instantiates. An example of an *induced* use attribute could be the price of an offered service.
- *Observed Attributes*: Information recorded during the exercise of a *use*. Such info cannot be derived directly from the entities involved in the *use*. Examples of *observed* use attributes are the duration of a *use*, etc.

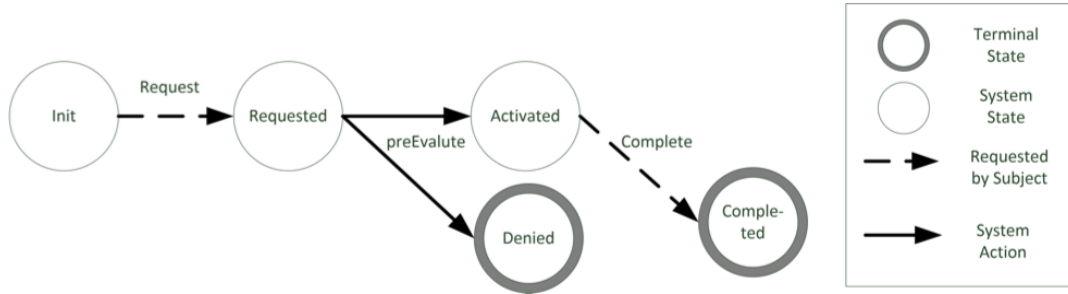
Update procedures for *induced* attributes can be specified during the definition of the implementation of a high level policy by the policy administrator. However, update procedures for *observed* attributes require the existence of a system function that returns the required value. For example, the specification of a policy, which requires recording the system time whenever a *use* is permitted, is as follows:

$$preUpdate \triangleq [u \text{ EXCEPT } !.st = \text{"activated"}, !.allowedtime = SystemTime()]$$

where *SystemTime()* is an internal function provided by the system framework that provides the current time. In TLA+, *EXCEPT* is a special purpose operator representing the modification of a function from a state to the next. In that next state all function values are left unchanged unless stated otherwise<sup>1</sup>.

Attribute update procedures in UseCON are performed whenever a *use* changes its state (e.g., from 'requested' to 'activated'). Therefore, during the execution of a *use*, the times of use attribute update in UseCON are represented in Figure 1.

<sup>1</sup>For a comprehensive definition of EXCEPT operator we refer the reader to [5].

Figure 2: Transition system of a single *use* pre-authorisation UseCON model

## 2.4 Transition Systems

Actions in the UseCON model are categorised to those triggered by a *subject*'s request and those operated automatically by the usage control system. More specifically, for every *use* supervised by the usage control system the following actions can be triggered by a *subject*.

- *Request*: This action performs the transition from the 'init' state of the *use* to 'requested'. Moreover, *request* creates a particular use instance that instantiates the requested *use* and also assigns values to the *id* attribute of the *use*.
- *Complete*: This action changes the state of the *use* from 'activated' to 'completed'.

The actions performed automatically by the usage control system, follows:

- *preEvaluate*: This action is performed by the usage control system only when the allowance of the *use* is governed by a pre-authorisation rule. This action changes the state of the *use* to either 'activated' or 'denied', depending on the outcome of the examined policy rule.
- *onEvaluate*: In case the allowance of a *use* is governed by an ongoing authorisation rule, the *onEvaluate* action is performed by the usage control system. If the particular policy rule is satisfied then the state of the *use* does not change. In case the policy rule is not satisfied, the state of the *use* is changed to 'stopped'.
- *Activate*: This action is performed only when the allowance of the *use* is governed by an ongoing authorization rule. It follows the execution of the *Request* action and changes the state of the *use* from 'requested' to 'activated'.

Any of the previous actions, apart from modifying the *st* attribute value, may also update other *use* attribute values. Such modifications in the use attribute values are considered to be implementation specific, as already described in Section 2.3. Moreover, despite the existence of a great number of updates on attribute values, the execution of any of the previous actions is considered to be atomic<sup>2</sup>, i.e., a single behavioural step. The transition system for a single *use* UseCON system controlled by a pre and an ongoing authorisation policy rule is depicted in Figures 2 and 3, respectively.

In its initial state, no *uses* are exercised in the system, and thus the first state of every behaviour must satisfy the TLA+ predicate *init*, which defines that the set-variable *U* is an empty set:

$$Init \triangleq U = \{\}$$

<sup>2</sup>This constraint is realistic due to the fact that a use is recorded centrally on the policy decision point and there is no need for attribute updates of other entities (*subject* or *object*).

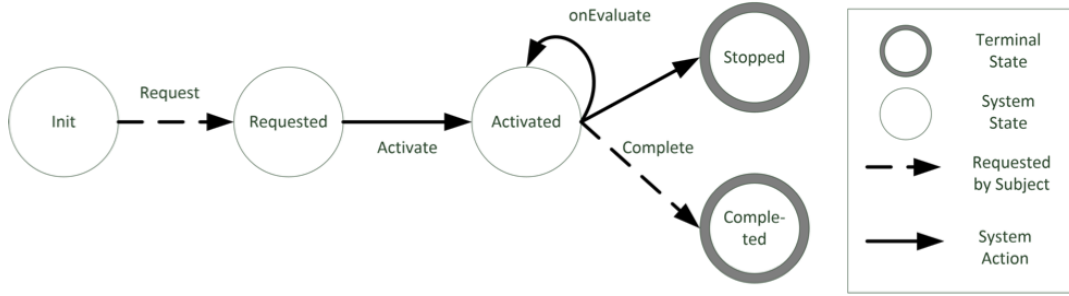


Figure 3: Transition system of a single *use* ongoing authorisation UseCON model

Moreover, according to the time period that a *use* request evaluation is performed, a pre-authorisation and ongoing-authorisation transition systems is created.

### 2.4.1 Pre-authorisation

The possible actions that can be performed on a pre-authorisation UseCON system are the *use* request for a new *use*, the evaluation of an already requested *use*, or the termination of a *use* that is already executed. Therefore, the *next* action, that describes all the possible next states could be a *request*, *evaluate*, or *complete* action, described as follows:

$$Next \triangleq Request \vee preEvaluate \vee Complete$$

More specifically, the *Request* action selects non-deterministically<sup>3</sup> a new *use*  $x$ . This is verified by searching the set of uses  $U$ , and returning one that has not already been requested by the subjects or processed by the usage control system. Consequently, in case that this particular  $x$  exists, the *request* action creates the corresponding *use* instance that instantiates *use*  $x$  and inserts it to the set  $U$ . The definition of the *Request* action is as follows:

$$Request \triangleq \exists u \in (S \times A \times O): ( \\ \wedge \forall x \in U: (x.sid \neq u[1].id \vee x.aid \neq u[2].id \vee x.oid \neq u[3].id) \\ \wedge U' = U \cup \{createUse(u)\})$$

The *preEvaluate* action examines if there are any *uses* that have been requested but have not been processed by the usage control system. More specifically, *preEvaluate* examines if there is any use instance with state attribute value equals to '*requested*'. Consequently, the action evaluates the policy rule that governs the allowance of the *use* that the specific use instance instantiates. Based on the outcome of that policy rule the action modifies the state of the use either to '*activated*' or to '*denied*' with *preUpdate* and

<sup>3</sup>The non-determinism property is implied by the use of  $\exists$  operator. For comprehensive information refer to [5].

*denUpdate* use attribute update procedures respectively as follows:

$$\begin{aligned}
 preEvaluate \triangleq \exists u \in U : & (\wedge u.state = \text{“requested”} \\
 & \wedge IF(PolicyRule) THEN \\
 & \quad U' = (U \setminus \{u\}) \cup \{preUpdate(u)\} \\
 & ELSE \\
 & \quad U' = (U \setminus \{u\}) \cup \{denUpdate(u)\})
 \end{aligned}$$

The *Complete* action simulates a *subject*'s request to terminate the execution of a currently active *use*. If such a use exists, its state attribute value should be equal to 'activated'. Consequently, the *completed* action modifies the state attribute value from 'activated' to 'completed' with the *comUpdate* use attribute update procedure. The *Complete* action is defined as follows:

$$\begin{aligned}
 Complete \triangleq \exists u \in U : & (\wedge u.state = \text{“activated”} \\
 & \wedge U' = (U \setminus \{u\}) \cup \{comUpdate(u)\})
 \end{aligned}$$

*CreateUse* is the procedure that creates a use instance that has the identities of the *direct subject*, *object*, *action*. The definition of the *createUse* procedure follows:

$$\begin{aligned}
 createUse(x) \triangleq [sid(x) = x[1].id, aid(x) = x[2].id, oid(x) = x[3].id, \\
 state(x) = requested, att(x) = k]
 \end{aligned}$$

All the other use attribute update procedures perform a dual role. Firstly, they alter the *state* use attribute to the desired value (e.g., *preUpdate* to 'activated' or *comUpdate* to 'completed'). Secondly, the use attribute update procedures modify the values of other use attributes according to the requirements of the usage control system that they are called to describe. For example, a possible *preUpdate* procedure can be the following:

$$preUpdate \triangleq [u EXCEPT !.st = \text{“activated”}, !.att = value]$$

The same definitions apply to the *denUpdate*, *comUpdate* procedures.

## 2.4.2 Ongoing-authorisation

The transition system of the ongoing-authorisation UseCON model is differentiated from the pre-authorisation model in a number of ways. Firstly, in an ongoing model, a *use* that is requested is permitted to be activated without the evaluation of any policy rule. Secondly, at a given time interval <sup>4</sup>, an ongoing action is executed *onEvaluate*. Thus, the specification of the *Next* action on an ongoing authorisation model has the following definitions:

$$Next \triangleq Request \vee Activate \vee onEvaluate \vee Complete$$

---

<sup>4</sup>The determination of the exact interval is left open as an implementation issue.



The *Activate* action searches for the existence of a use with state attribute value equal to 'requested' and consequently updates it to 'activated' by executing the *preUpdate* procedure.

$$\begin{aligned} \text{Activate} \triangleq \exists u \in U : (\wedge u.\text{state} = \text{"requested"} \\ \wedge U' = (\{U \setminus \{u\}\} \cup \{\text{preUpdate}(u)\}) \end{aligned}$$

Whereas *onEvaluate* action evaluates an ongoing policy rule, based on this result it either leaves use to 'activate' state, but it can possibly update the rest attribute values with *onUpdate* procedure, or modify its state attribute value to 'stopped' with *termUpdate* use attribute update procedure as follows:

$$\begin{aligned} \text{onEvaluate} \triangleq \exists u \in U : (\wedge u.\text{state} = \text{"activated"} \\ \wedge \text{IF}(\text{PolicyRule}) \text{ THEN} \\ U' = (\{U \setminus \{u\}\} \cup \{\text{onUpdate}(u)\}) \\ \text{ELSE} \\ U' = (\{U \setminus \{u\}\} \cup \{\text{stopUpdate}(u)\}) \end{aligned}$$

Moreover, in the case where there is no need for an ongoing use attribute update procedure, the *onUpdate* procedure can be substituted with the *UNCHANGED U* TLA+ operator that leaves the variable *U* unmodified. Here, the semantics of 'request' and 'complete' actions are considered to be the same with the pre-authorization model.

### 3 Model Checking with TLC

*Toolbox* is an Integrated Development Environment (IDE), which is designed for the definition and verification of TLA+ specifications [1]. Specifically, the *toolbox* editor provides functionality for the definition and alteration of TLA+ specifications, and supports syntax highlighting. Additionally, an automatic parser checks the defined specifications for syntax errors and presents them accordingly by marking them in the used modules.

The tool in use for the verification of a TLA+ specification in *toolbox* is the TLC model checker. Specifically, TLC explicitly generates and computes all the possible states of a system. However, many times the specification of a system might contain an infinite number of states. TLC handles such specifications, by choosing a finite model of the system and in turn checks it thoroughly. Specifically, the creation of a system's model in TLC requires the definition of its specifications, properties and values of constant parameters [1]. A specification represents all the behaviours that have to be checked. Moreover, the values assigned to constant parameters are utilized for the instantiation of a specification. TLC can check a model for deadlocks, invariants and properties [1]. A deadlock occurs when the model reaches a state in which its next-state action allows no successor states. An invariant is a predicate that is evaluated on a system state. Consequently, an invariant holds on a system specification, if and only if, every state of all the behaviours of a system satisfy that predicate. Properties are temporal formulas that must be evaluated to true for all the behaviours of the model. TLC has some limitations regarding the handling of a subclass of TLA+ specifications and properties that it can check [5]. A very helpful feature of TLC is the fact that when it identifies an error during the verification process, it provides an error trace viewer that allows the exploration in a structured view of the debugging information. Moreover, TLC supports an arbitrary evaluation of states and action formulas in each step of the trace.

Table 1: Safety and liveness properties in UseCON

	Safety		Liveness	
	Former state	Latter state	Former state	Latter state
<b>Pre</b>	Completed	Any other state	Requested	Activated or Denied
	Activated	Requested or Denied	Requested	Completed or Denied
	Denied	Any other state	Activated	Completed
<b>Ongoing</b>	Completed	Any other state	Requested	Activated
	Activated	Requested	Requested	Completed or Stopped
	Stopped	Any other state	Activated	Completed or Stopped

### 3.1 Use management

One of the fundamental properties that can be verified in a system is that of type correctness. Specifically, type correctness is considered to be an invariant which determines that all the variables of the system are assigned with values originating only from a specific set of values. The UseCON specification uses a single variable  $U$  which corresponds to the set of system uses. The *invariant* property that defines type correctness in the UseCON model, is defined as follows:

$$TypeCorrectness \triangleq U \subseteq Uses$$

where  $Uses$  is the set of all records that have the following form (i.e., all the record fields are assigned with values originating from their domains):

$$[s.id: SubjectIDS, o.id: ObjectIDS, a.id: ActionIDS, st: USTATE, att: ATTDOMAIN]$$

Moreover, the definition of the domains  $SubjectIDS$ ,  $ObjectIDS$ ,  $ActionIDS$  and  $USTATE$  is:

$$SubjectIDS \triangleq \{s.id: s \in S\}, ObjectIDS \triangleq \{o.id: o \in O\}, ActionIDS \triangleq \{a.id: a \in A\}$$

where

$$USTATE \triangleq \{\text{requested, activated, denied, stopped, completed}\}$$

A *use* is capable of recording detailed historical information about the operation of *uses* in the system. Consequently, a valid implementation of the UseCON model, where multiple *use* processes are operating concurrently, depends on a proper management of the *use* instances that represent these *uses*. Specifically, all *use* instances must adhere only to the state transitions depicted in Figures 2 and 3 for the pre-authorisation and ongoing-authorisation models, respectively. Based on that, a number of *safety* and *liveness* properties can be defined. For example, a *safety* property (a faulty state cannot be reached) states that a *use* instance cannot be evaluated as '*requested*' in its *st* attribute, if it has previously been evaluated as '*completed*'. The semantics, expressed in TLA+, which verify the previous property for all the uses of a system are defined by the following temporal formula:

$$Safety \triangleq \Box(\exists u \in U: u.st = \text{"completed"} \implies \Box(u.st \neq \text{"requested"}))$$

Similar *safety* properties can be defined for all the possible prohibited state transitions. The complete set of prohibited state transitions is depicted in Table 1 under the header '*safety*'. Moreover, two examples that illustrate the violation of *safety* properties are depicted in Figure 4. Specifically, the transition of the

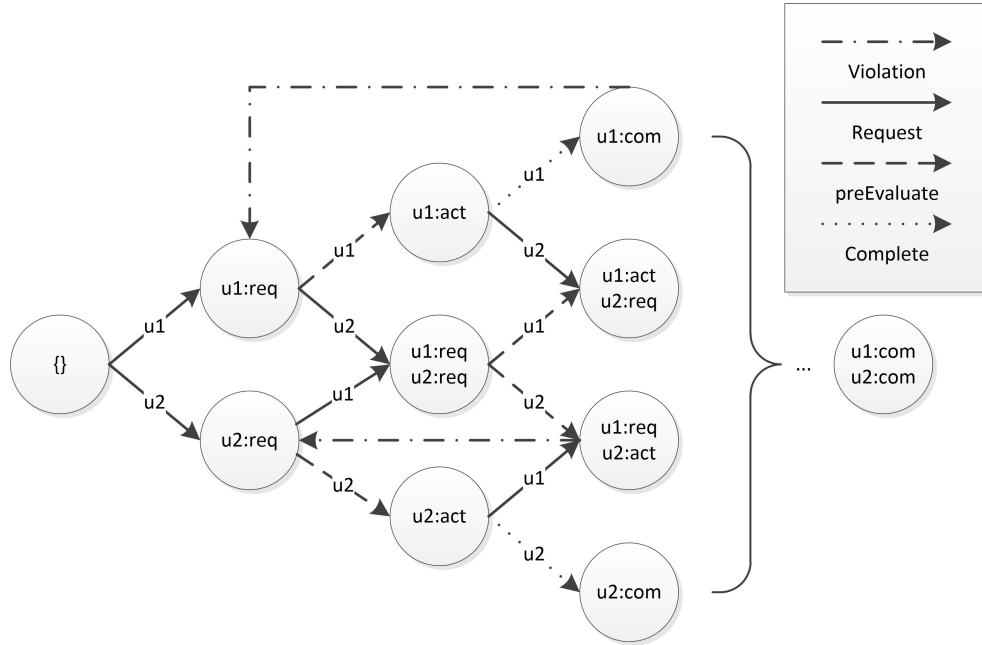


Figure 4: Violation of safety properties

*use* instance  $u1$  from ‘completed’ into ‘requested’ results in a violation of a *safety* property. A similar violation is considered during the transition of the *use* instance  $u2$  if the state transits from ‘activated’ to ‘requested’. In addition, the definition of *liveness* properties in the UseCON model determine all the valid state transitions regarding any use instance. For example, Figure 2 presents that a *use* instance that has at any given state an *st* attribute value that is evaluated to ‘activated’, then its attributed value must be eventually evaluated as ‘completed’. All the possible state transitions that are eligible to be performed are depicted in Table 1 under the header ‘liveness’. This property is defined in TLA+ as follows:

$$Liveness \triangleq \forall u \in U : u.st = \text{“activated”} \rightsquigarrow u.st = \text{“completed”}$$

## 4 Performance Evaluation

The verification of the examined model was performed on a MacBook Pro (Mid 2014) using TLA+ Tool version 1.5.6 and running TLC version 2.12. The system operates on OS X El Capitan and its hardware specifications are: 2.5GHz Intel Core i7 with 16GB 1600 MHz DDR3.

A first set of results were collected by verifying deadlock-freedom for the model of UseCON. Specifically, in a pre-authorisation UseCON model, we consider all the *uses* of the model to be requested, and therefore, to be in either the state of ‘activated’ or ‘denied’. The *uses* being activated were finally completed. Therefore, the final state in every *use* must be ‘denied’ or ‘completed’. The TLC model checker evaluates all the behaviours of the model and terminates when it reaches to a deadlock, and the actions of the deadlocked behaviour are presented along with the attribute values in each state. Moreover, we performed a verification of the model against safety and liveness properties described in Section 3.1. The verification was finished without raising any errors, but this applies only for the verification of  $\leq 12$  uses for both the pre and ongoing authorisation models. Although higher numbers of *uses* were consid-

Table 2: Performance evaluation

Authorisation model	Uses	Diameter	States found	Distinct states	Deadlock (seconds)	Safety and Liveness (seconds)
<b>Pre</b>	2	6	21	12	1	3
	8	23	277969	16832	2	5
	12	31	45533665	560128	53	138
<b>Ongoing</b>	2	7	33	16	1	3
	8	25	367873	23808	2	7
	12	37	79112449	1224704	118	241

ered, the verification process had to be interrupted due to the excessive amount of time required for its completion given the number of *uses*. The collected verification results (see Table 2) requires a better understanding of the internal procedures TLC is applying to compute the behaviours of the model (i.e., generation of the transition system). Initially, TLC computes the states that verify the *Init* predicate and inserts them into a set  $G$ . For every state  $s \in G$ , TLC computes all the possible states  $t$  that  $s \mapsto t$  can be a step in a behaviour. Specifically, TLC substitutes the values assigned to variables by state  $s$  for the unprimed variables of the *Next* action, and then it computes all the possible assignment of values to the primed variables that makes the *Next* action true. For every state  $t$  found by the former procedure it is added to set  $G$  if it does not already exist. The previous two actions are repeated until no new states can be added in  $G$ . Therefore, the verification results produced by TLC, are: *Diameter* expresses the number of states in the longest path of  $G$  in which no state appears twice; *States found* expresses the number of examined states; *Distinct States* expresses the number of examined distinct states. The verification results produced by TLC for the *pre* and *ongoing* UseCON models are presented in Table 2. An additional column presents the actual running time of the TLC model checker in seconds.

## 5 Conclusion and Future Work

The application of model checking as a technique resulted in formally verifying the use management functions of UseCON. The trace of the deadlock errors verifies that the defined specifications of the system operate correctly. An advantage of using model checking techniques for the verification of usage control models is the provision of formal guarantees with regards to the correctness of the model, without requiring an implementation of it. Nevertheless, known issues of model checking, i.e., state explosion problem, prevented the timely verification of *uses* when these increase in number. This resulted in providing formal guarantees for the correctness of UseCON for  $\leq 12$  uses for both the pre-authorisation and ongoing-authorisation models, as depicted in Table 2.

In the future, we aim to demonstrate the verification of complex policy rules supported by UseCON, and specifically investigate the verification of ongoing policy rules. This is of interest since a *use* request might lead to a policy violation in other concurrent *uses*. Thus, any *use* request should be followed by an evaluation of all policy rules in all *uses* in the system to avoid conflicts and violations. Finally, we are considering frameworks, such as secBIP [9], that will allow us to compositionally analyse security properties and generate secure-by-construction systems. Also, we would like to explore guaranteeing security properties by-construction for any number of *uses* in our model through the application of architectures from predefined architecture styles [7] that capture properties of specific access and usage control policies.

## References

- [1] Microsoft Corporation: *TLA+ Tools*. Available at <http://research.microsoft.com/en-us/um/people/lamport/tla/tools.html>. (visited on February 2018).
- [2] Antonios Gouglidis, Ioannis Mavridis & Vincent C. Hu (2014): *Security policy verification for multi-domains in cloud systems*. *Int. J. Inf. Sec.* 13(2), pp. 97–111, doi:10.1007/s10207-013-0205-x.
- [3] Christos Grompanopoulos, Antonios Gouglidis & Ioannis Mavridis (2012): *A Use-Based Approach for Enhancing UCON*. In: *Security and Trust Management - 8th International Workshop, STM 2012, Pisa, Italy, September 13-14, 2012, Revised Selected Papers*, pp. 81–96, doi:10.1007/978-3-642-38004-4\_6.
- [4] Vincent C Hu, Rick Kuhn & Dylan Yaga (2017): *Verification and Test Methods for Access Control Policies/Models*. *NIST Special Publication 800-192*, doi:10.6028/NIST.SP.800-192.
- [5] Leslie Lamport (2002): *Specifying Systems, The TLA+ Language and Tools for Hardware and Software Engineers*. Addison-Wesley.
- [6] Aliaksandr Lazouski, Fabio Martinelli & Paolo Mori (2010): *Usage control in computer security: A survey*. *Computer Science Review* 4(2), pp. 81–99, doi:10.1016/j.cosrev.2010.02.002.
- [7] Anastasia Mavridou, Emmanouela Stachtari, Simon Bliudze, Anton Ivanov, Panagiotis Katsaros & Joseph Sifakis (2016): *Architecture-Based Design: A Satellite On-Board Software Case Study*. In: *Formal Aspects of Component Software - 13th International Conference, FACS 2016, Besançon, France, October 19-21, 2016, Revised Selected Papers*, pp. 260–279, doi:10.1007/978-3-319-57666-4\_16.
- [8] Jaehong Park & Ravi S. Sandhu (2004): *The UCON<sub>ABC</sub> usage control model*. *ACM Trans. Inf. Syst. Secur.* 7(1), pp. 128–174, doi:10.1145/984334.984339.
- [9] Najah Ben Said, Takoua Abdellatif, Saddek Bensalem & Marius Bozga (2014): *Model-Driven Information Flow Security for Component-Based Systems*. In: *From Programs to Systems. The Systems perspective in Computing - ETAPS Workshop, FPS 2014, in Honor of Joseph Sifakis, Grenoble, France, April 6, 2014. Proceedings*, pp. 1–20, doi:10.1007/978-3-642-54848-2\_1.
- [10] Xinwen Zhang, Masayuki Nakae, Michael J. Covington & Ravi S. Sandhu (2008): *Toward a Usage-Based Security Framework for Collaborative Computing Systems*. *ACM Trans. Inf. Syst. Secur.* 11(1), pp. 3:1–3:36, doi:10.1145/1330295.1330298.
- [11] Xinwen Zhang, Jaehong Park, Francesco Parisi-Presicce & Ravi S. Sandhu (2004): *A logical specification for usage control*. In: *9th ACM Symposium on Access Control Models and Technologies, SACMAT 2004, Yorktown Heights, New York, USA, June 2-4, 2004, Proceedings*, pp. 1–10, doi:10.1145/990036.990038.